



AIOPS

PERSPECTIVES

ISSUE II 2021

digitalisationworld.com



AIOPS - Where to next?

AIOPS has a crucial role to play in any
Digital Operations Management strategy

[STAND OUT FROM THE CROWD]

Front Cover Promotion Package

The front cover of Digitalisation World is the most prestigious position in the magazine. It allows for an image and a technology story connected to your company to be showcased in front of a global audience. Digitalisation World magazine is distributed to over 37,000 readers interested in the data centre, cloud, storage, security, automation, hardware and software sectors - covering the entire spectrum of technology and industries impacted by digital transformation.

What do you get?

- Cover image with your branding
- 2 - 3 pages of editorial within magazine
- Full page advertisement
- Magazine with your image hosted on www.digitalisationworld.com for six weeks
- 468 x 60 pixels banner for 6 weeks
- Low res PDF of your cover and article for web promotion
- 20 minute Zoom interview

Price on application

Contact us today for further information on:

+44 (0)1923 690205

jackie.cannon@angelbc.com

Editor's View

By Phil Alsop



AIOps – where to next?

AS DIGITAL OPERATIONS MANAGEMENT becomes the possible catch-all title for all things Ops-related, it will be fascinating to see how AIOps – both the technology and the vendor market – develops in the coming months.

Every now and then, working as an editor in the IT industry, I bump into a new technology or trend that, for no particular reason, grabs my attention like no other. Whisper it quietly but, long before VMware appeared on the horizon, I was fascinated by the storage virtualisation technology of a company called Datacore. Happy to report, they still exist and, of course, virtualisation has taken off in a major way. And then there was InfiniBand. What a great technology, with a major future, until Compaq abandoned it, and it all but vanished, and now finds a niche within many a server.

And so we come to AIOps. I won't tell you my success rate when it comes to liking and following a particular technology or trend and how it fares subsequently, but I like to think that AIOps in some form or other will continue to play a significant role in the digital business era. Yes, it might become just one part of the relatively newly defined Digital Operations Management (DOM) space, but I think there are enough vendors and enough end user interest to ensure that AIOps continues to thrive for quite some time yet. After all, the advantages it brings, both in its own right, and in the way in which it works alongside various other technologies, such as DevOps, Value Stream Management and Digital Experience Management, mean that it stands up to the closest end user scrutiny.



AIOPS
PERSPECTIVES
ONLINE SUMMIT

I hope that this publication, and our AIOps Perspectives 2 virtual event <https://aiopspectives.com> will contribute in some way to spreading the AIOps message. Thanks to all the companies who have contributed content – we've got a great mix of articles which reflect the state of the market right now – pure-play AIOps, subjects adjacent to AIOps and, of course, some focusing on AI in broader terms.

In talking with an intelligent automation vendor recently, I was rapidly jumping between his references to AI, then IA, then AI, and it came to me that there was the chance to update the words of that famous children's song – Old Macdonald had a farm – or at least the chorus, with AI, IA, O (the O being for observability!). I haven't progressed far on the verses just yet, other than:

- Young, IT-literate Macdonald had a server farm, AI, IA, O
- And on that farm he had some DevOps, AI, IA, O
- Not sure what noises to add at this point!

Anyhow, apologies for the digression, I hope you enjoy the publication, and please do make sure to attend our virtual event in June.



Editor		
Philip Alsop	+44 (0)7786 084559	philip.alsop@angelbc.com
Sales Manager		
Peter Davies	+44 (0)2476 718970	peter.davies@angelbc.com
Sales Manager		
Jessica Harrison	+44 (0)2476 718970	jessica.harrison@angelbc.com
Director of Logistics		
Sharon Cowley	+44 (0)1923 690200	sharon.cowley@angelbc.com
Design & Production Manager		
Mitch Gaynor	+44 (0)1923 690214	mitch.gaynor@angelbc.com
Publisher		
Jackie Cannon	+44 (0)1923 690215	jackie.cannon@angelbc.com

Circulation & Subscriptions
+44 (0)1923 690214 circ@angelbc.com

Directors
Stephen Whitehurst: Chairman
Scott Adams: Chief Technical Officer
Sukhi Bhadal: Chief Executive Officer

Published by:
Angel Business Communications Ltd, 6 Bow Court, Burnsall Road, Coventry CV5 6SP
T: +44 (0)2476 718970 E: info@angelbc.com



AIOps Solutions is published 4 times a year on a controlled circulation basis in Europe, Middle East and Africa only. Subscription rates on request. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd. © Copyright 2021. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. ISSN 2396-9016 (Online)

66 Why intelligent automation is an absolute must

At a time when competitive pressure has never been more intense, good customer service has never been more important.

By Kay Knoche,
DACH Lead for Decisioning
& AI Solutions at
Pegasystems



AIOps

- 06 How to get started with AIOps
- 08 AIOps: Why network monitoring plays an even more important role
- 10 How AIOps is driving digital transformation
- 16 Implementing AIOps in 5 simple steps
- 20 Observability vs. Monitoring for DevOps professionals
- 24 AIOps: Moving beyond the trigger-driven alerts model with deep learning
- 26 Seven KPIs for AIOps
- 28 The importance of visibility in AIOps for digital transformation
- 32 It's all about scale
- 34 The move to SaaS



- 36 Observability vs. Monitoring: What's the difference?
- 40 Lift-and-Shift cloud migrations
- 44 The data-driven IT operations organization
- 46 Reducing operational costs and improving productivity in a post-Covid world
- 52 Tackling cloud complexity with the right tools
- 54 Telecom: Enabling automaton everywhere through AI and analytics
- 56 The importance of AIOps within value stream management

AIOps+

- 58 The future of service management in the DevOps era
- 60 Weeding your digital garden to provide outstanding software experiences
- 62 It's now or never for more perfect software
- 68 Accelerate your automation journey with a Centre of Excellence
- 70 Understanding legacy infrastructure by looking at the Egyptian pyramids
- 74 The key to unlocking optimum digital employee experience: Proactivity
- 76 Three tips for creating outstanding digital customer experiences



80 Why is APM important? Breaking down the benefits

82 How to thrive in the enterprise AI era

AIOPS + +

84 How to train your AI algorithm

86 What is an AI strategy and how do you build one?

90 The ROI of data-driven development: improving how teams work

93 Using the greater data ecosystem to drive great decision making



NEWS

12 Study says Digital Operations Management is essential

13 ScienceLogic are in the money

14 Raygun launches APM for Node.js

14 Entuity innovations simplify network performance and diagnostic monitoring

14 Federator 4.4 arrives

22 Visibility across the entire software stack

23 Neteo acquires Stackify

24 Supporting the digital transformation journey

30 Improving cloud application performance

31 Logz.io launches Prometheus-as-a-Service

31 LogicMonitor expand global partner network

42 Honeycomb launches refinery

42 GAVS Technologies launches the GAVS Healthcare Technology Institute

43 Lightstep introduces 'Change Intelligence'

43 Graylog launches Graylog cloud

50 Survey examines digital workspace deployment and performance monitoring

50 Datadog completes Sscreen acquisition

51 Poor visibility and silos impact business outcomes

64 CloudFabrix launches AIOps 2.0

65 Monitoring and optimising Google Cloud infrastructures

65 TET Partners with Centreon

72 Catchpoint releases enhanced version of the WebPageTest API

73 Prioritise, predict and act with BMC Helix

78 Companies consider work from anywhere

79 BigPanda launches Automatic Incident Triage

79 Anodot and Coralogix partner

88 U-Haul speeds its digital transformation with Dynatrace

88 Micro Focus recognizes Gazprombank

89 Datadog launches Windows server monitoring of live traffic

89 Park 'N Fly speeds up development cycles by 29% with Dynatrace



How to get started with AIOps

It can be stressful for IT Ops to manage, and AIOps is one solution coming to help IT improve system reliability and customer satisfaction while reducing some of the manual work.

BY ISSAC SACOLICK, ON BEHALF OF IVANTI



“Houston, we have a problem.”

This is exactly what people in IT Operations think whenever a series of monitoring alerts go off simultaneously. Within five minutes, they receive the invite to the bridge call and start to read out what each monitor is reporting. The team reviews incidents raised in Cherwell, network alerts from Nagios, system alerts in LogicMonitor, log files in Splunk, and Jenkins deployments to identify potential causes and decide a course of action.

Fifteen minutes into the call and business leaders join to get status and remind everyone of the expected service levels on business-critical applications. Business leaders have higher expectations on system reliability and performance, especially on customer-facing applications and critical workflows.

It can be stressful for IT Ops to manage, and AIOps is one solution coming to help IT improve system reliability and customer satisfaction while reducing some of the manual work.

What is AIOps?

- Proactive IT leaders look to apply AIOps capabilities to reduce complexities, enhance employee experiences, and improve service levels.

AIOps refers to applying AI and machine learning capabilities to support IT operations. A must-have outcome of AIOps helps IT correlate multiple monitoring alerts into a single, time-sequenced incident that's easier to review and faster to resolve. It might show that a Continuous Integration (CI)/Continuous Delivery (CD) deployment triggered database failures followed by application errors and group them into a single incident. An incident manager seeing this sequence can quickly deduce the root cause, consult the development team on the recent changes, and determine the required steps to restore service.

AIOps in incident management and processing data from multiple monitoring tools is one use case of platform intelligence. Applying AI and machine

learning in IT Operations also includes:

- Discovery and Dependency Mapping (DDM) automations to capture hybrid cloud infrastructure changes, maintain the CMDB, and capture dependencies between systems.
- Virtual agents use Natural Language Processing (NLP) to help end-users search and access the service catalog.
- Sentiment analysis applied to end-user feedback can trigger a follow-up by the IT service desk when there is a negative customer response.
- Machine learning categorization of requests can improve mapping requests to the correct services and route them to the right team.

With several machine learning capabilities available, IT Ops leaders should consider the following steps in getting started with AIOps.

1. Configure the AIOps Data Sources

Machine learning algorithms require clean data sources, so the first steps to enable AIOps capabilities are to connect data sources and iterate through their configurations.

IT teams should start by configuring DDM to capture systems, network, and application data from data centers, private clouds, and public clouds, including AWS and Azure. DDM should update the CMDB regularly, and IT Ops should then map systems information to business services and service levels. ITSM practices like incident management and change management have a lot more context with a DDM powered CMDB. Incidents and change tickets already capture what happened, and with an integrated DDM powered CMDB, the tickets can also include where they happened. The added context helps IT resolve issues faster and enables analytics on repeat problem areas.

Secondly, connect all the system, network, and application monitoring tools to a central AIOps solution. This solution should reduce the noise in reviewing multiple monitoring tools by correlating multiple alerts to one manageable incident. Connecting the data sources also starts capturing historical data for predictive analytics and anomaly detection once there is sufficient data to train machine learning algorithms.

2. Discover ITSM Pain Points and Opportunities

Once IT connects DDM to the CMDB and aggregates monitoring data, it's time to put the data to use and improve KPIs. IT should seek opportunities to improve customer satisfaction, mean time to resolve issues (MTTR), and system reliability.

Proactive IT operation teams look at these improvements strategically by following the steps below:

- Itemize and prioritize IT operational pain points from incident management, request management, and

services with business impacts.

- Review the data and insights in DDM and monitoring insight tools and perform a data discovery exercise. Identify, "What does the data tell you" and what improvement opportunities the data suggests.
- Align pain points with the opportunities to help prioritize focus areas.
- Determine stakeholders that can steer process improvements and articulate success criteria.
- Identify which metrics or KPIs demonstrate whether improvements meet the selected success criteria.

These steps assure that IT gains a business partner on the implementation and invests effort in areas that deliver the greatest business impact.

3. Use Agile Process to Implement Solutions

This process of identifying stakeholders, priorities, and KPIs also helps define an ongoing process of improvements. As teams use data-driven insights to identify opportunities and partner with stakeholders on pain points, a backlog of process improvement projects should emerge. The backlog is exactly what proactive IT leaders need when implementing AIOps capabilities. These leaders should form agile teams in IT Operations to implement improvements iteratively.

Why an agile process? Priorities are likely to change based on business needs, opportunities, and risks. The team might focus on end-user computing applications in the first sprints to support hybrid working environments. The group might then shift to work with an application development team that's modernizing applications and migrating them to the cloud.

In both these scenarios, topology maps, infrastructure visualizations, and other DDM tools help the team identify implementation opportunities.

For example, let's say end-users open incident tickets for slow application response time during the afternoon. The DDM's topology report can help IT Ops identify the bottleneck by comparing application flows between poor and normal performance periods. The team may then choose to adjust cloud elasticity parameters to ramp up infrastructure in the bottlenecking areas ahead of peak periods. After the modifications, the team monitors performance and incidents to validate whether the change addressed the issue.

Using an AIOps solution, including a DDM subset of a solution, enables this data-driven cycle of process improvement. Machine learning in DDM and AIOps reduces the complexity of working with multiple data sources or stale data. Instead, it enables IT to focus on the customer, understand pain points, implement solutions, and validate results. The result is a proactive IT Operations team that's constantly improving and delivering stronger system performance to business stakeholders.

AIOps: Why network monitoring plays an even more important role

AIOps, or Artificial Intelligence for IT Operations is paving the way for the digital transformation of businesses.

BY SRINI BANGALORE, CONSULTANT, HEAL SOFTWARE INC.

AIOps APPLIES Machine Learning techniques and Big Data Analytics to large collections of data gathered from various IT operations tools in order to identify issues in real time and proactively fix them. The corporate network is a key resource that enables transactions between enterprise users and applications. So, let's try to understand the role of network monitoring in implementing a robust AIOps strategy for an enterprise.

The IT operations team's challenge

Over the last few years, we have seen a significant change in how enterprise applications and services are deployed and delivered. Legacy architectures

relied heavily on enterprise IT infrastructure and networks to host, deliver, and run key business-centric applications. Enterprises had complete control over the performance of the network, the infrastructure, and the applications. Today, with the migration of applications and services to hybrid networks (i.e., on-premises and cloud), and with an increased adoption of distributed microservices based architectures, enterprises are struggling to ensure optimum performance of services across infrastructure over which they have little control.

Although many of the business-centric applications are offloaded to external service providers, the IT



Operations team finds it ever more challenging today to ensure an uninterrupted availability of these applications. Imagine the intense pressure on the IT Manager when a video conference session during an earnings call ends abruptly because of a problem somewhere: in the corporate network or in the service provider's network! It is therefore imperative that IT managers have end-to-end visibility across private and public networks, and across the various distributed applications that are collaborating to deliver a seamless service to users. Essentially an IT Operations team must be able to:

- Monitor the performance of applications and the infrastructure (physical and virtual) over which the applications are running.
- Monitor the network to quickly identify and locate bottlenecks and performance issues related to the network infrastructure that connects enterprise users to external, cloud-based services.

How does AIOps impact application performance?

According to Gartner Research, application performance monitoring (APM) involves using software and/or hardware components to monitor three key aspects:

- Digital Experience Monitoring (DEM), which deals with monitoring how end-to-end application availability impacts overall experience.
- Application discovery, tracing, and diagnostics (ADTD), which deals with the discovery of applications and visualization of their topology, logging of transactions between applications, and diagnostics related to application components.
- Artificial Intelligence for IT Operations (AIOps), which has emerged from the adoption of Artificial Intelligence (AI) and Machine Learning (ML) techniques to analyze application behavior. Preventive Healing – the next frontier for this cutting-edge technology – involves coming up with lead signals or early warnings, flagging anomalies, predicting bottlenecks in application performance, and taking remedial actions to mitigate or avert problems even before they can occur.

Monitoring the network – easier said than done!

Traditional network management systems relied primarily on managing devices using Command Line Interfaces (CLI) and protocols like Simple Network Management Protocol (SNMP). However, these approaches are insufficient when dealing with today's networks. With advances in networking and communication technologies, there is a proliferation of high-speed networks that use a variety of physical and wireless media to deliver different types of traffic such as voice, video, and data.

Furthermore, with the growing appeal for cloud computing, corporate networks are required to be connected to the public Internet – exposing corporate networks to unprecedented security risks. It is a

well-known fact that whenever end-to-end service performance is degraded, the network is the first to be blamed, although in fact, the problem may be related to a poorly performing application or to an abnormal user-triggered activity. In such cases, without full visibility into the network, the IT Operations team is usually the last to hear about a problem – long after a user has experienced service degradation – and the IT response is at best, reactive in nature. A more proactive approach will help mitigate problems and ensure optimum service levels.

What is NPMD and Why is it Important?

As noted earlier, the biggest challenge for IT is to keep the network up and running all the time. IT Managers lack the granular network visibility that is required to help them understand what applications and services are running across corporate-, public- and third-party networks and how to effectively detect, troubleshoot and resolve performance problems when they surface. What is even more challenging is the fact that cloud-based architectures make a prolific use of microservices residing in containers that live on distributed networks, which are themselves virtual and dynamic in nature. As a matter of fact, network virtualization adds an unprecedented level of complexity to managing present day networks.


Thankfully, the simple management techniques of the past have been augmented by a plethora of new network protocols and technologies in recent times to provide unprecedented visibility into the network and its performance. The traditional network management function has taken a broader scope and is widely known today as Network Performance Monitoring and Diagnostics (NPMD). NPMD includes the monitoring of physical network infrastructure and servers, measurement of network- and application-level traffic flows, and detection/remediation of problems across the network.

NPMD tools can process data from a variety of sources comprising network-device-generated logs, health metrics and traffic statistics (such as application-specific bandwidth usage, network latency, packet loss, etc.), to provide real-time and historical views of traffic patterns and predict network behavior.

Additionally, NPMD tools can integrate with AIOps to identify the root causes of application performance issues. Thus, NPMD tools essentially provide the means to understand how network performance impacts application performance and the resulting digital experience for the end-user.

The bottom line

It is no surprise that APM and NPMD go hand-in-hand and are essential components of a good IT strategy for today's complex networks. It is essential to have good NPMD tools in place and to ensure that they integrate well with AIOps software so that end-to-end service experience can be guaranteed.



How AIOPs is driving digital transformation

In the current digital age, any successful organisation will understand that business growth requires managing dynamic infrastructures and embracing digital transformation.

BY ALI SIDDIQUI, CHIEF PRODUCT OFFICER, BMC SOFTWARE



AIOPs is a key example of this – applying artificial intelligence (AI), machine learning (ML), and big data analysis to automate and ultimately enhance IT operations. Once implemented, there is a significant improvement in speed, agility, and efficiency for the business, with any anomalies being identified and often fixed in real-time. As IT teams continue to face a surge in operational data volumes and complexities due to multi-cloud and remote working environments,

the adoption of AIOPs has never been more important for organisations.

The key benefits of AIOPs

There are a variety of key benefits that come with implementing an AIOPs solution, which provide insight into all layers of the IT environment. By combining the power of machine learning with IT, AIOPs separates the true problems which may urgently

need addressing, from the false alarms which can often appear in extremely complex environments. Artificial intelligence and machine learning are both incredibly valuable tools when it comes to analysing performance data and identifying patterns which help IT teams take precautionary maintenance measures. Where some customers already have knowledge on identifying patterns, organisations may find it helpful to have a more modern approach with no-code capabilities, which can capture that user-defined knowledge. AIOps not only provides a more effective approach to automatically identifying patterns and accelerating root cause isolation, but in environments where there is so much data and complex dependencies, it is the only suitable approach for managing and reducing noise. An intelligent AIOps solution automates manual processes meaning organisations will begin to see an increase in employee satisfaction, productivity and customer retention, as well as significantly saving time and resources.

Making the switch

To start on a successful AIOps strategy, it is essential for organisations to assess what the business is set to gain from the initiative. For example, this could be addressing current needs such as more effective event noise reduction or faster probable cause analysis to help reduce issues. Additionally, it would be used for identifying areas of IT Ops friction that need to be addressed.

The next aspects to address are the people, processes, and tools. IT organisations often work across complex, hybrid environments, meaning it is more expensive to use traditional approaches, and workload can rapidly become unmanageable. When it comes to selecting the right tools, it is critical for teams get the appropriate vendor, which can address the specific use cases and assist the business in achieving its goals. The value of the solution will increase with a vast amount of data, enabling it to

observe and analyse in real-time. It is important for an AIOps solution to have an open approach which can be integrated with existing IT tools and data sources, to support agility and collaboration across all functions.

Embracing digital transformation

Over the next couple of years, AIOps is expected to become significantly more mainstream, as organisations continue embracing digital transformation and becoming an autonomous digital enterprise (ADE). In an ADE, automation is a complementary business function which works alongside humans to execute tasks faster, free employees from manual tasks and lower costs, while improving customer satisfaction. With the rapid adoption of technologies such as IoT and 5G, data volumes will continue increasing and become harder to manage.

The demand for AIOps will therefore increase, as businesses attempt to meet performance demands. There will also be a stronger link between DevOps and IT Ops, which will use insights from AIOps to ensure that performance management of applications is at the forefront. As the adoption of cloud-based apps and services is already in motion for many organisations, AIOps will start to play a central role in managing these. The rapid improvement in technology along with algorithms becoming mainstream, mean the mass adoption of AIOps is inevitable.

As data volumes and the adoption of technologies such as IoT and 5G increases, AIOps will become essential to any data-driven business. AIOps enables organisations to use predictive problem solving approaches which results in more intelligent business decisions. With the solution, there is reduced disruptions and time spent on manual tasks, which gives employees the chance to focus on driving operational excellence and helping the business evolve into an ADE.



Study says Digital Operations Management is essential

OpsRamp has completed a study on how IT operations leaders are adapting their processes and tools to help their organizations achieve a smooth economic recovery.

The OpsRamp 2021 State of Digital Operations Management study, which was conducted in March by a third party, includes input from 132 IT operations directors or above in U.S. companies with at least 500 employees and \$5 million in annual technology spending. The year 2020 was characterized by frequent lockdowns, strict social distancing norms, and unexpected pivots in business strategy. Enterprise IT departments shifted to digital-first, cloud-centric solutions to stabilize operations and maintain revenue streams.

As the pandemic wanes and the world economy begins to reopen, enterprise IT leaders are leaning further into digital transformation. According to OpsRamp's study, IT leaders are prioritizing tools, technologies and methods that unlock data, enhance visibility and mitigate risk, such as digital operations management platforms. Following are three key takeaways from OpsRamp's research:

The Emergence of Digital Operations Management Platforms

A digital operations management platform brings together diverse data sources, optimizes organizational processes and uses analytics to drive innovation and agility across the organization. Given the volume, velocity, and variety of system events that occur daily, IT needs a digital operations management platform that can reduce repetitive alerts, identify the proximate cause(s) for an IT outage, and ensure compelling user experiences using AI/ML.

A majority of US technology leaders surveyed by OpsRamp see value in deploying a digital operations management platform with native capabilities for hybrid, multi-cloud and cloud native monitoring, while supporting intelligent incident management and

automated remediation, and which integrates with their existing tools in a single place. Seventy-three percent of respondents expect to roll out a digital operations management platform in 2021, while 21% expressed interest in purchasing a modern platform if no budgetary constraints were holding them back.

Why Invest in a Digital Operations Management Platform

While tech practitioners have long tried to find a "single pane of glass" for IT operations, they've usually ended up with a single glass of pain due to the duct-taped nature of legacy IT Operations management suites. Modern digital operations management platforms are different. They combine data, analytics and workflows to provide IT leaders with relevant and timely insights into their organization's hybrid IT ecosystem.

OpsRamp's study found that IT leaders plan to roll out digital operations management platforms for the following reasons:

- Deploying automation and AIOps (51%) so that IT can use machine learning algorithms and process automation to reliably cut down alert noise, drive faster root cause identification, and handle repetitive activities.
- Driving agility and faster resolution (46%) with a single source of truth for incident response so that IT has the right situational context.
- Saving time and money (41%) by eliminating the need to switch between tools and by swapping out costly, legacy IT management tools.

Key Requirements for a Modern IT Operations Solution

The demands on modern IT operations solutions have never been higher.



Enterprise IT teams need technology platforms that can capture real-time telemetry for a wide variety of infrastructure environments, offer the ability to scale with demand, provide chargeback and capacity recommendations to optimize infrastructure, and deliver algorithmic insights for troubleshooting and repairing IT services.

Given the fallout of the SolarWinds sunburst supply chain attack, the OpsRamp survey shows that platform security (59%), which is the ability to withstand sophisticated attacks, is the most critical attribute of a modern solution. The next two top capabilities include hybrid infrastructure management (52%) for controlling the chaos of distributed architectures and flexible price offerings (37%) that offer clear and compelling business value at a competitive price point.

"The study validates what we're hearing from customers," said Sheen Khoury, Chief Revenue Officer at OpsRamp. "The economic tumult caused by the pandemic forced IT leaders to reevaluate their toolkits and processes. On the way out are legacy IT operations tools that limit an organization's ability to support digital transformation and deliver outstanding customer experiences. Technology leaders are looking to invest in digital operations management platforms that can help drive faster time-to-market with proactive insights and enable rapid incident resolution with workflow automation."

ScienceLogic is in the money

SCIENCELOGIC has raised \$105 million in growth financing. Silver Lake Waterman led the company's Series E round with participation from existing investors Goldman Sachs, Intel Capital and NewView Capital.

The investment will support the company's continued innovation in the AIOps market and further broaden ScienceLogic's position within the \$30+ billion IT Operations Management software market.

Companies across the globe are adopting cloud strategies and tools more than ever amid the COVID-19 pandemic, with massive demand for digital services and scalability. We believe ScienceLogic is poised to meet that demand as organizations look for hybrid-cloud toolsets from trusted partners to securely harness business agility, cost savings and innovation.

"More than ever, IT Operations Management has taken root as a front-office priority supporting mission-critical digital experiences that define the way we live, work and play. As large enterprises shift workloads to the cloud while managing on-prem resources, new tools are paramount to deliver service visibility and faster incident resolutions made better by advanced AI/ML technologies," said Dave Link, Founder & CEO of ScienceLogic. "What we're witnessing

is a major investment cycle away from legacy monitoring tools and toward AIOps platforms."

The funding is intended to accelerate ScienceLogic's product development and engineering leadership, supporting the company's broader expansion plans and the reach of its flagship SL1 digital infrastructure monitoring platform.

Funds are expected to be allocated toward recruitment efforts and product investments aimed at cloud-native technologies including microservices and container solutions, AI/machine learning, and hybrid cloud operations that transforms digital experiences and enhance security.

"The ScienceLogic team has built a leading platform to monitor mission-critical infrastructure and applications and is at the center of some of the largest, most complex IT environments at the forefront of digital transformation," said Shawn O'Neill, Managing Director and Group Head of Silver Lake Waterman. "Dave Link and the leadership team have a long track record of building value and trust with customers and we look forward to partnering with the team and helping drive further adoption".

Some of the largest global enterprises, federal agencies and managed service providers trust ScienceLogic's modern



platform to ensure availability of their applications and business operations across hybrid-cloud and multi-cloud deployments. The scalable SL1 monitoring platform helps IT operations teams ingest hyperscale data volume in real-time across disparate hybrid-cloud architectures, while its patented discovery and automation technology improves agility, accelerates incident response and drives productivity by strengthening application health, resolution time and user experience.

The funding news comes after a recent announcement from the company to expand its product development, engineering and sales and marketing staff. ScienceLogic's AIOps market leadership also was recently highlighted by the Forrester Wave, which included ScienceLogic as one of only three firms highlighted as AIOps Leaders – honors that further cement the company's forward momentum.

Raygun launches APM for Node.js

RAYGUN has announced the latest in their APM language support: Raygun APM for Node.js.

Node.js is the technology of choice for some of the largest companies in the world, and remains one of the most popular with over 50% of developers having used Node.js in 2020. As companies respond to the changing demands of customers in our current climate, it has never been more important for scaling companies to better manage their online experiences. In today's world, application quality is inextricably linked to business success — undetected slowdowns could cost a company millions of dollars.

Raygun APM helps developers detect problems before they affect customers with actionable, real-time insights into application quality metrics like Apdex score, slowest requests and methods, active issues, and more. Because they have

access to insights and diagnostics, developers can pinpoint issues faster, and some companies reduce the time spent resolving issues by 90%.

Raygun's points of difference in the APM market are that they offer more code-level detail, better data visualization, and a usage-based billing model that compliments modern infrastructures.

Product Manager at Raygun, Bianca Grizhar, says that: "Raygun APM presents an exciting opportunity for the Node.js community. We built Raygun with the modern development team front-of-mind — it's easy to use, yet highly configurable for flexibility with technologies like Docker. We're proud to offer one of the best tools on the market so that developers can proactively identify performance problems before they have a chance to negatively impact a business."

Entuity innovations simplify network performance and diagnostic monitoring

ENTUITY SOFTWARE'S Version 19 release greatly enhances the award-winning enterprise monitoring and management solution. With its innovative all-in-one approach, Entuity V19 offers new physical storage and server monitoring and OS monitoring of Windows and Linux servers while adding enhanced Discovery capabilities and Asset Management features.

The pandemic caused an instant transition to entirely remote workforces, as well as a massive uptick in BYOD devices. Security and software-defined solutions were adopted with intensified frequency, and demand for expanded network access spiked. As organizations have progressively used the network in ways it wasn't planned for, the need for expanded network monitoring solutions became paramount.

To meet this increased demand, Park Place Technologies, a single source for award-winning data center hardware maintenance and a full suite of managed services, has launched Entuity Software Version 19, formerly Entuity Network Analytics. The new Entuity Software™ Network Monitoring tool release provides an innovative all-in-one approach –

ensuring IT teams have the capability to discover and track their network infrastructure changes, take network devices and services under monitoring and management, roll out configuration updates, and monitor server and storage devices through a single dashboard.

"IT teams have seen their network perimeters rapidly expand with virtualization, cloud computing, and software-defined networks over the past few years. That excludes the rapid change the COVID-19 pandemic has also brought to IT infrastructures," said Kathie Lyons, Executive Vice President and General Manager of ParkView and Entuity.

"One key piece of feedback we get from customers is that visibility is a significant challenge for them. Visibility into what's on their network, where it is on their network, and more importantly, how their network is changing. With this latest release, we've focused on bringing visibility of assets and network health to our customers while setting the stage for Server and Storage visibility."

The new software package offers a newly built Asset Management functionality,

enhanced Discovery, new server and storage device monitoring and OS monitoring capabilities on Windows and Linux. These new features build on the award-winning enterprise monitoring and management solution.

The new Entuity Software package offers a single source for customer-discovered-and-managed assets, providing detailed reporting and an at-a-glance view of what is in the network. When new assets are discovered, they will appear flagged for users to quickly evaluate. From this view, users will be able to manage or unmanage assets all in one place.

Entuity can discover server infrastructure and storage assets, providing a list of assets directly to its built-in CMDB, where users can seamlessly decide which assets to take under monitoring management. Entuity monitors all the components of server and storage hardware, as well as OS and Virtual Machine instances. This gives users direct access to the health of their physical and virtual server assets and storage devices alongside its award-winning network performance data and event management capabilities.

Federator 4.4 arrives

PROPHETSTOR Data Services has introduced Federator.ai 4.4. Federator.ai, ProphetStor's Artificial Intelligence for IT Operations (AIOps) platform, provides intelligence for orchestrating Kubernetes container resources on top of virtual machines (VM) or bare metal, allowing users to operate applications without the need to manage the underlying computing resources manually.

Key features in the latest Federator.ai release include Spot instance recommendations for MultiCloud cost analysis and the support of metrics collected by Sysdig, a secure DevOps company that helps organizations secure containers, Kubernetes, and cloud services.

Utilizing metrics from Sysdig, Federator.ai provides AI/Machine Learning-based predictions for containerized application workloads and cluster node resource usages as the basis for resource recommendations to joint Sysdig and ProphetStor customers. DevOps no longer needs to monitor the Kubernetes cluster utilization constantly and manually adjust capacity.

With resource recommendations from Federator.ai, DevOps can efficiently perform intelligent resource planning and significantly reduce over-provisioning costs. The application-aware workload prediction also enables Federator.ai to auto-scale application pods via extension API server, providing fitted

resources for optimal performance. "With the latest release of Federator.ai, Sysdig customers can enjoy the benefits of AI-based intelligent workload predictions using application and cluster node metrics from Sysdig. Joint customers can easily view the prediction results and recommendations from the integrated Sysdig dashboards, all from the same single pane of glass," said Tad Lebeck, EVP of Business Development, ProphetStor. "Also, Federator.ai analyzes evictable application workloads and provides the most cost-effective Spot instance recommendations while maintaining required resources for these applications. This is a tremendous cost-savings for Federator.ai customers."



TRAINING

DEVELOPMENT

SKILLS



BUSINESS COMMUNICATIONS

WEBINARS

**Specialists with 30 year+ pedigree
and in-depth knowledge in these
overlapping sectors:**

Expertise: Moderators, Markets, 30 Years + Pedigree
Reach: Specialist vertical databases
Branding: Message delivery to high level influencers
via various in house established magazines,
web sites, events and social media



Semiconductor (Silicon/Compound)

Publications include: Compound Semiconductor, Silicon Semiconductor, CS China, SiS China



Power Electronics

Publications include:
Power Electronics World



Future Mobility

Publications include: TaaS Technology, TaaS News



Data Centres

Publications include: DCS Europe, DCS UK, SNS International



SmartSolar UK & Ireland

Publications include: Solar and Power Management, Solar UK and Ireland



Sensors

Publications include: Sensor Solutions Magazine, Sensor Solutions International



Digitalisation

Publications include: Digitalisation World, Information Security Solutions, Managed Services



Photonics

Publications include: PIC Magazine, PIC Conference

Expert Moderators

Dedicated technical and time-served experts/editors



MARK ANDREWS

Mark Andrews is technical editor of Silicon Semiconductor, PIC Magazine, Solar+Power Management, and Power Electronics World. His experience focuses on RF and photonic solutions for infrastructure, mobile device, aerospace, aviation and defence industries



PHIL ALSOP

Journalist and editor in the business to business publishing sector for more than 30 years currently focusing on intelligent automation, DevOps, Big Data and analytics, alongside the IT staples of computing, networks and storage



JACKIE CANNON

Director of Solar/IC Publishing, with over 15 years experience of Solar, Silicon and Power Electronics, Jackie can help moderate your webinar, field questions and make the overall experience very professional



DR RICHARD STEVENSON

Dr Richard Stevenson is a seasoned science and technology journalist with valuable experience in industry and academia. For almost a decade, he has been the editor of Compound Semiconductor magazine, as well as the programme manager for the CS International Conference

For more information contact:

Jackie Cannon **T:** 01923 690205 **E:** jackie@angelwebinar.co.uk **W:** www.angelwebinar.co.uk
6 Bow Court, Burnshall Road, Coventry, CV5 6SP. UK
T: +44(0)2476 718 970 **E:** info@angelbc.com **W:** www.angelbc.com



Implementing AIOPs in 5 simple steps

IT Operations teams have faced mounting challenges as the tech stack has both grown in size and complexity.

BY MICHAEL PROCOPIO, PRODUCT MARKETING MANAGER, MICRO FOCUS



NOW HAVING TO OVERSEE data distribution and accessibility, applications, varying interdependent infrastructures, customer touchpoints and more, keeping these plates spinning has become a monumental task. Ticket management alone can place overwhelming pressures on teams, where sifting through the sheer noise can be incredibly resource-intensive. Fortunately, for IT Operations teams, IT management solutions have matured in line with these challenges.

Around 2015, the idea of IT Operations Analytics (ITOA) crystallised. The application of big data analytics to raw IT operations data unlocked the ability to analyse huge amounts of information. This would, in principle, allow for more informed and focused decision making. At the time, however, the tools for doing so were cost-prohibitive and highly specialised.

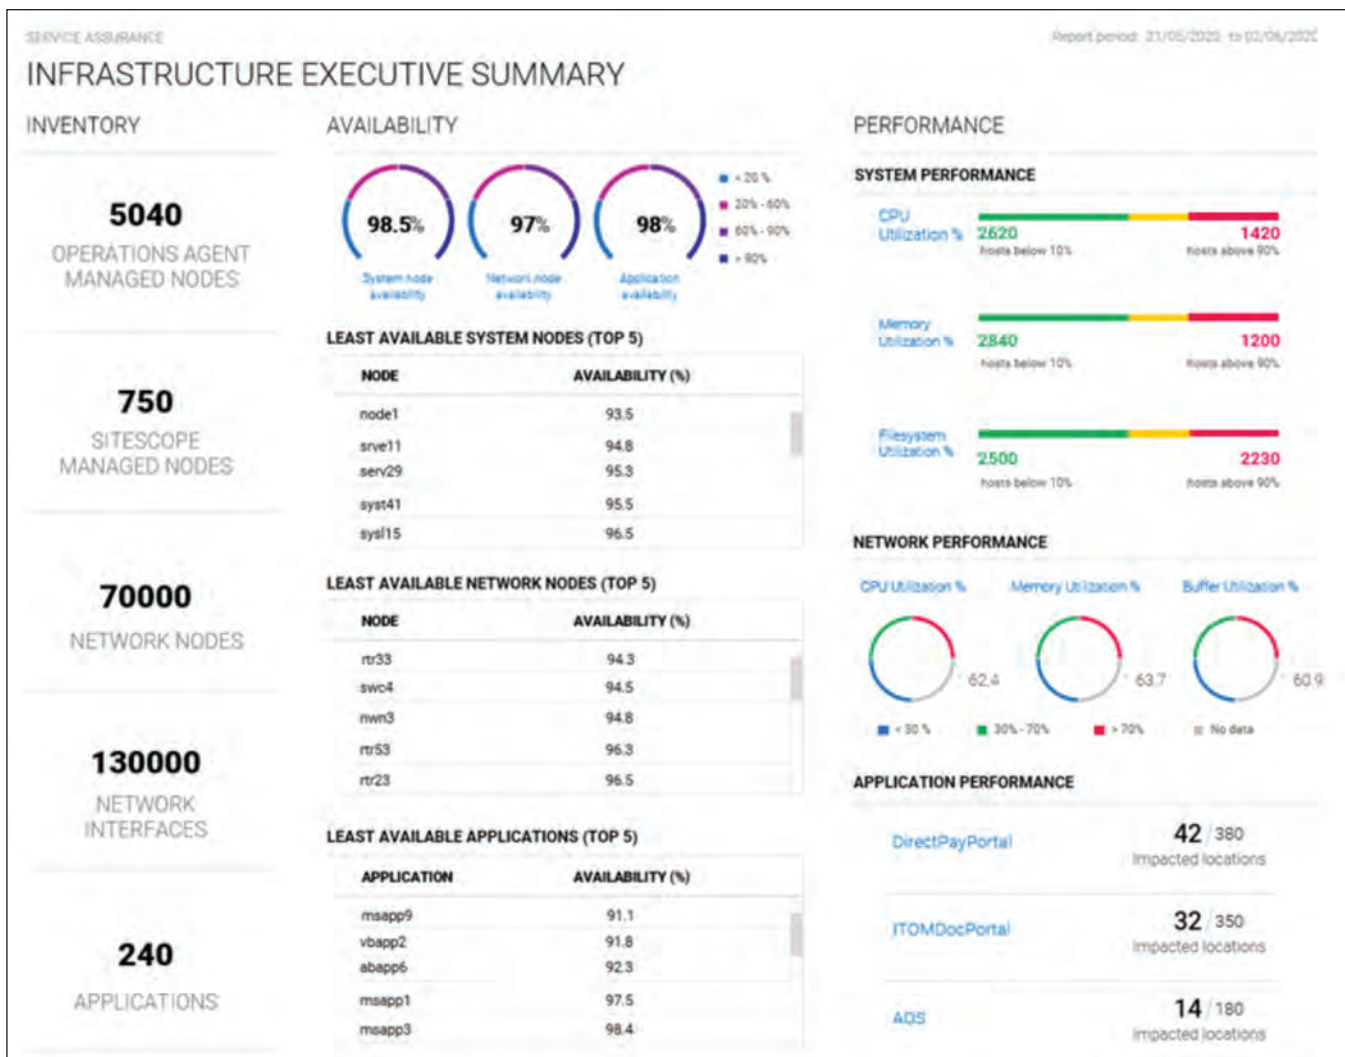
Fast forward a couple of years and deep learning, machine learning and AI are at the top of the Gartner hype cycle and AIOps enters the lexicon of cutting edge CIOs. The idea of combining AI, machine learning and automation with big data analysis on

this ITOA data frames AIOps as the ideal extension of the ITOps team. Although at the time, the “how” was somewhat missing.

Now, with commercialised intelligent systems and white-label solutions widely available, AIOps helps leading enterprises manage their IT estates and streamline workflows. With this market maturation, we’ve seen more technical deployments such as “multi-domain AIOps”, differentiating between domain agnostic vs domain specific solutions, and the concept of an AIOps overlay of existing tools emerge.

At its core principles, however, AIOps always comprises of three tenants: the ability to Observe estate performance, Engage with IT personnel, and Act through automation and remediation capabilities. With these abilities, vast, complex, hybrid environments are able to be managed relatively simply.

And yet, despite its proliferation and market growth, adopting AIOps solutions still sounds like a job for an Oxford data science graduate. It isn’t. With the right



provider and partner, the intelligence should be built in. All IT teams need are the operational know-how and management skillsets so commonly found within the department.

So here's how to get started:

Adopt with purpose

Aligning with business goals and addressing well-known pain points is key. AIOps has the habit of coming across as either marketing-speak or nebulous scientific tinkering rather than an essential pillar of IT operations.

Research from EMA found mature AIOps supports on average 8 different domain specific roles and 11 cross-domain roles. In other words, it can quickly seem unwieldy.

This can make pitching to purse holders extremely challenging on first look. However, identifying readily addressable points of friction as small scale pilots can make executive buy in easier. With these in hand, teams can pre-determine success criteria and give a quantifiable, projected figures or ROI.

Event noise reduction is a common starting point that can have a significant service impact. Large enterprises are inundated with IT events, so employing AIOps to help filter, manage and eventually remediate tickets can provide a relatively simple, yet highly impactful win.

Enable observability

The ideal state of AIOps is for it to be interwoven with the entire IT estate. Every system, subsystem and even end node wants to eventually be connected, brought into the automated analysis and remediation fold.

There are three reasons for this. Firstly, the more data, the more the system can learn. It needs to be able to read and understand patterns, which allows it to become more proactive and independent of the ITOps team – allowing for them to prioritise complex issues that need their attention. Secondly, faults or issues may otherwise lie beyond the system's reach. IT errors typically cause a chain of events, having your system only know half of the picture adds limited value. Lastly, and related to the previous point, issues caused by foundational systems may only come to light further



up the tech stack – remediating this requires system-wide reach in order to find the root cause.

Domain agnostic AI

It is common these days for domain tools to come with built in intelligent capabilities. Yes, these have their immediate uses, but they typically operate within a silo. They fall short of true AIOps. As mentioned earlier, there are often root causes to symptomatic events that are many times removed. In short, single domain, or domain-specific AI has relatively significant constraints compared to their domain agnostic counterparts.

With agnostic systems, you can achieve what is called a first-time fix. This means that the AI system has pointed ITOps straight to the root cause or ideally fully remediated the issue itself. One Micro Focus customer experienced a 118% percent uptick in first-time fixes

after rolling out the appropriate, domain agnostic tool. As you can imagine, this translates to huge time and resource savings.

Centralising the data

A common method of enabling observability and expediting implementation involves creating a central data lake. AIOps requires diverse and wide ranging data sets in order to function effectively. This includes both historical data, to train the system, know what is normal and what is not, how to react and how not to, along with current live data in order to react to, remediate and even ultimately predict issues.

While many organisations have their estate domains covered with some form of intelligent system, this is more often than not siloed. Data lakes resolve and prevent this walling off while providing additional advantages. Firstly it is easy to develop single-pane dashboards that give ITOps teams the same level of observability that the system has (fig. 1). Secondly, it can simplify implementation and roll out; new pilots and projects are simpler to deploy, and scalability is more achievable.

From observe to action

AIOps can often be described in only passive terms, i.e. what it can see and tell you. The next big step in realising its full potential is to add that final Act characteristic.

Embedding AIOps into workflows to automate and remediate is where the true value comes in. The impact this can have is to orders of magnitude. Another MicroFocus customer saved over \$1m in their first year, rising to \$4M per annum by year three with 95% of remediation incidents completed without any human interaction. The start point for any AIOps journey will vary dependent on business priorities, and the order in which these steps are addressed may change but these are the boxes you'll need to check.

While many organisations have their estate domains covered with some form of intelligent system, this is more often than not siloed. Data lakes resolve and prevent this walling off while providing additional advantages. Firstly it is easy to develop single-pane dashboards that give ITOps teams the same level of observability that the system has (fig. 1). Secondly, it can simplify implementation and roll out; new pilots and projects are simpler to deploy, and scalability is more achievable



DIGITALISATION WORLD

ONLINE ROUNDTABLE



- Based around a hot topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion
- Moderated by the editor Phil Alsop, this could also include 3 speakers
- Questions would be prepared and shared in advance
- There would be an opportunity to view and edit out any unflattering bloopers

This event would be publicised for 4 weeks through all our mediums including:

- A banner on the Digitalisation World homepage for 8 weeks
 - 4x weekly dedicated HTMLs
 - 4x news pieces which would also appear on the weekly newsletters
 - Promoted through our social media platforms for 8 weeks (pre and post event)
 - Available as an on-demand asset through all mediums
 - All registered attendees' details would be made available to you
- Cost: £4995**

Contact: Jackie Cannon
jackie.cannon@angelbc.com

Observability vs. Monitoring for DevOps professionals

What precisely are the requirements of a DevOps practitioner, as opposed to an SRE, legacy developer, or operations manager? And do those specific requirements require a different approach to monitoring?



BY WILL CAPPELLI, CHIEF TECHNOLOGY OFFICER EUROPE AND HEAD OF PRODUCT STRATEGY AT MOOGSOFT

What is the relationship between observability and monitoring? If I am monitoring a stack with the latest tools, is there anything else I need to do to make that stack truly observable? The short answer is yes. The type of monitoring required for observability has been singled out by the market as that type of monitoring which is of particular interest to DevOps practitioners and Site Reliability Engineers (SREs). But what is it that makes observability stand out from other kinds of monitoring?

In this post, I'll discuss the requirements of a DevOps practitioner, as opposed to an SRE, legacy developer, or operations manager, and how those specific requirements require a different approach to monitoring: namely, observability-enablement.

What the DevOps community needs

A DevOps practitioner is primarily in the business of rapidly adding new functional components to digital services. The need for speed has forced changes on both service architecture and on the scope of a developer's concern.

With regard to service architecture, the DevOps practitioner requires extreme levels of modularity and, as a corollary, very small modules with very short life-spans.

Furthermore, boundaries between contemporary digital services become blurred and, rather than introduce completely new services, services tend to evolve incrementally (module by module, so to speak.)

With regard to the scope of concern, a DevOps practitioner, unlike a legacy developer, needs to maintain ongoing awareness of what is taking place in the production environment and to be able to step in to prevent or ameliorate performance issues that might affect the continuous flow of new service modules into production. (Hence, the 'Ops' in DevOps.) Now, in order to keep track of the production environments and the array of live digital services, the DevOps team needs some kind of monitoring.

Monitoring, in many ways, replicates the process by which we, as human beings, obtain information about our environments. Signals are generated by objects, events, and fields which are transmitted through space and time, are picked up by various sense organs, and then processed by our brains and nervous systems.

Similarly, monitoring technologies pick up various signals generated as a by-product of the actions taken by digital systems in the production environment. They transmit those signals through various pathways which add structure to the signals



and ultimately deliver them to a human (but, also possibly a robotic) agent who formulates and then executes some kind of response. The agent, in other words, plays the part of the brain and the nervous system, although some elements of nervous system functionality are taken up by the structure adding components along the signal transmission paths.

Why legacy monitoring fails

Legacy monitoring systems are problematic for DevOps practitioners in two ways, however. First, they ingest and process signals at relatively slow rates (usually 10s of signals per minute.) Second, the structure added on the transmission pathways is information-heavy (i.e., it purports to tell the agent receiving the signal quite a lot about what the signal means without extracting that information from the signal, relying instead on pre-fabricated rules and interpretations) and rigid (i.e., the structures can't change without some kind of intervention and, in any case, not while the signals are actually being transmitted).

With DevOps-crafted digital services, instead, signals need to be generated and processed very quickly (1000s or 10,000s per minute). Because the service is in a state of continuous incremental evolution, the interpretive structures which are intended to make sense of the signals need to arise directly from the signals themselves and also must evolve as the signal stream itself evolves. Finally, the signals themselves need to be granular – capable of transmitting information about the states of many distinct modules more or less simultaneously.

Because of the mismatch between the signal transmission technology of legacy monitoring systems and the requirements of DevOps practitioners, the DevOps community was definitely in need of a new type of monitoring. In a way, however, DevOps community influencers misunderstood the critical issues. They conflated the community's needs for a new signal stream with the requirement to understand causal relationships among system events. In fact, while it is true that legacy monitoring systems did not support robust techniques for discovering causality, AIOps systems, particularly those that supported Moogsoft's five-dimensional AIOps model, went a long way towards addressing that gap. The DevOps community, however, was unaware, or uninterested, or maybe even, in some cases, ill-disposed to the idea of AIOps.

Observability is not enough

Initial concerns about the complexities of modern infrastructures led DevOps practitioners to concepts developed under the aegis of Control Theory during the 1960s. This is the concept of a system that generates signals of sufficient quality and quantity to allow signal recipients to understand the causal relations governing the system. In Control Theory, it is the data alone that reveals whether or not a causal

relationship is present. On this basis, the DevOps community came to believe that if the right data streams were provided, causal insights would be forthcoming without any further effort or algorithmic processing.

Three different types of data seemed to match the granularity and dynamism of DevOps-crafted digital services:

1. Metrics
2. Logs, and
3. Traces.

Now, metrics had long been one piece of the legacy monitoring puzzle but, in most cases, a relatively unimportant or supplementary one. Logs had thrust into prominence as a result of the mid-2010s excitement around Big Data and Splunk's successful IPO, but the use of log management databases as real-time monitoring tools was a novel notion.

Tracing was a largely discredited application performance monitoring technique, but the idea of tracing an execution path across microservices rather than multi-tiered app system components seemed compelling. In the end, traces proved difficult to deploy in this context, although some vendors developed highly specialized products for that purpose. But by 2020, the market was full of vendors targeting the DevOps community with observability-oriented monitoring tools majoring in the ingestion and presentation of signals based on metrics and logs.

The good news was that the metric and log data streams could keep pace with the event streams characteristic of DevOps-crafted digital services. The bad news was that the causal insights were not forthcoming. In fact, by turning to more granular, lower-level signals, it became even more difficult for practitioners to figure what caused what amongst the buzzing, booming confusion of signals they were now able to observe.

What comes next

That takes us to where we are now. The DevOps community has the signals it needs, but it still needs the analysis. It still needs the patterns that make those signals tell a story. And the only way in which patterns will emerge is through an automated pattern discovery technology. This technology needs to survey large, high-dimensional data sets in micro-seconds and, then, almost simultaneously, tease out the correlations and the causal patterns that make sense of those data sets.


This is not something that human agents can accomplish, no matter how intelligent, how knowledgeable, or how experienced. Insead, it requires something very much like AIOps, except now targeted at the granular, fast-changing data streams made available by observability technology.

Visibility across the entire software stack

NEW RELIC has introduced **New Relic Explorer**, its reimagined Full-Stack Observability experience that delivers innovative new visualizations and capabilities to give engineers unprecedented visibility into their complete estate. With zero configuration required, New Relic Explorer brings together an organization's telemetry data from across applications and infrastructure to provide an essential live view of an entire software system's health and performance. With this new single source of truth, engineers can quickly discover emerging performance issues and swiftly take action to get systems back to full health before customers or employees are impacted. All existing New Relic customers will be automatically upgraded to this new and powerful experience in the coming weeks.

"IT environments have become increasingly complex and engineering teams use a myriad of tools to see the performance of their entire tech stack, but there's no way to see everything in one place and quickly take action. This leads to fragmented visibility, cumbersome and inefficient workflows, and out-of-control costs," said Bill Staples, President and Chief Product Officer, New Relic. "With this launch, New Relic has delivered true Full-Stack Observability so engineers can finally get their arms around all of their telemetry data to quickly understand what's happening and resolve issues faster before they become problems. "For digital organizations and their IT teams, real-time visibility into critical software infrastructures is crucial in keeping their business running smoothly," said Stephen Elliot, Program Vice President, Management Software and DevOps, IDC. "Customers now demand real-time awareness and observability into IT environments via a single dashboard; providing them with simplicity and speed for identifying, solving and preventing problems before they impact the customer experience."

"Being able to visualize and analyze data with just a few clicks of a button with New Relic Explorer is amazing – definitely a time saver for us," said Paweł Przybyła, Director of Software Development at Metapack. "And I love having the ability to drill down into the 'why' behind changes



taking place. Having quick accessibility to the data we need with New Relic Explorer has been empowering for my team."

In today's complex environments that rely on manual alert configuration, it's become increasingly challenging for organizations to understand the overall health of their entire system at a glance. It's become a challenge to discover where unintentional changes are occurring, what may be contributing to them, and how issues across their system are related. New Relic Explorer enables engineers to uncover blind spots and surface details needed to support a faster, deeper understanding of large distributed software systems.

"With end-to-end observability under the hood in New Relic, it's been simple to instrument our applications as we move them to the cloud," said Greg Gentling, Director of Enterprise Architecture at Viewpoint. "Being able to grab telemetry data and take action before issues arise or customers notice – that's been huge." New Relic Explorer includes the following powerful visualizations and capabilities: **New Relic Lookout:** This industry-first innovation provides an estate-wide, real-time view of any changes in telemetry data, including third party and open source data, automatically drawing attention to where it's needed most.

This is delivered through an intuitive user experience visualization with no configuration or reliance on static, pre-configured alert thresholds required. Zoom in capabilities allow teams to pinpoint correlations, abnormal history and traces immediately. This gives engineering teams the ability to leverage

New Relic's Profiles to quickly uncover blind spots and unknown relationships, and to understand all changes so that issues can be resolved well before they impact the end customer or employee.

"Expanding the simplicity of New Relic Lookout into our stack will improve our ability to provide a seamless end-user experience," said Yang Tang, BEES Director Engineering, Anheuser-Busch InBev. "In current times, providing a smooth digital experience has been a top priority for us as more and more of our customers are engaging with us digitally. With New Relic, we will be alerted as soon as there is an abnormal activity, enabling us to immediately address the issue and ensure there's no interruption for our customers."

"New Relic Navigator: For the first time, customers can explore all entities at a glance in a highly intuitive visualization. The health of each application, service, container, function and host is displayed in red, yellow or green based on alert conditions, with the ability to group and filter based on attributes for a curated view of all the entities that a service (or set of services) encompasses. These traffic light colors display health so users can quickly and easily investigate large numbers of entities while simplifying cross-team collaboration. In addition, relationships between specific applications, hosts, containers or integrations are shown in one view, making it quick and easy to understand which upstream or downstream services are related to an issue. This enables engineers to gain a broader view of the overall health of their systems and troubleshoot cascading failures faster.

Netreo acquires Stackify

NETREO has acquired Stackify, a developer-centric provider of SaaS application performance monitoring (APM) solutions. The acquisition strengthens Netreo's full-stack IT monitoring and AIOps offerings by giving customers improved application performance management, centralized logging, full transaction tracing, deployment tracking capabilities and greater visibility into continuous improvement/continuous deployment (CI/CD) workflows.

Netreo helps technology leaders drive greater visibility, automation, and performance into their IT operations through its full-stack, AIOps-driven IT monitoring platform and Digital Experience Monitoring (DEM) solutions. The addition of Stackify's developer and devops-centric APM capabilities deepens observability capabilities needed to serve the growing \$30+ billion IT Operations Management software market.

"The acceleration of digital transformation and the rise of remote work this past year have made it even more critical for business and technology leaders to have actionable insights on the quality of internal and external digital customer experiences – more intelligently, efficiently and securely," said Jasmin Young, CEO at Netreo. "The combined capabilities of Netreo and Stackify will enable developers, DevOps, IT and business leaders to have AIOps-driven end-to-end observability and rapid incident resolution for their cloud and on-premises IT infrastructure. We are delighted to welcome the Stackify team and community to the Netreo family." Stackify's Retrace is an easy-to-use APM toolset trusted by thousands of customers in more than 50 countries to deliver faster time-to-value and proactively improve application performance by providing the critical insights developers need to deploy better applications faster. The agent-based solution helps both developers (Dev) and operations (Ops) teams quickly troubleshoot problems and identify ways to optimize code by:

- Delivering full visibility of application development operations from pre-development to production;
- Driving improved, seamless

collaboration between Dev and Ops teams through a consolidated, single-pane-of-glass UI;

- Accelerating time to value through full-featured, easy-to-use tools for quickly identifying and resolving production issues; and,
- Integrating code profiling, error tracking and application logs with real-time details on how long code takes to perform certain tasks

Also new to the Netreo portfolio is Prefix, the free programming tool that actively profiles, tests and validates code as it is written and traces issues for easy resolution. Prefix guarantees optimal application performance by helping even the most experienced developers find slow and underperforming SQL queries, hidden exceptions, ORM generated queries and unknown bottlenecks before moving applications into production. "We are thrilled to bring advanced APM capabilities to the Netreo community, and believe our combined product suite delivers best-in-class application and infrastructure monitoring to our customers," said Matt Watson, Founder and CEO at Stackify. "I look forward to working with leadership and technology teams dedicated to delivering easy-to-use IT solutions that advance our customers' business objectives, and to leading the development of exciting new solutions that address customer needs in new, high-growth markets."

Stackify Founder and CEO Matt Watson assumes the role of CTO for Netreo effective immediately and will focus on the integration of core capabilities and the ongoing development of customer-centric technology solutions. Before founding Stackify, the entrepreneur and technologist was the co-founder of Full Scale, a technology services company that helps businesses build teams of software engineers, and co-founder and CTO of VinSolutions, an auto dealership inventory management, website, CRM and internet marketing product that was sold to AutoTrader.com

The Stackify acquisition marks the second acquisition in nine months for Netreo, which also acquired Chicago-based cloud infrastructure monitoring company CloudMonix in June 2020.

Supporting the digital transformation journey

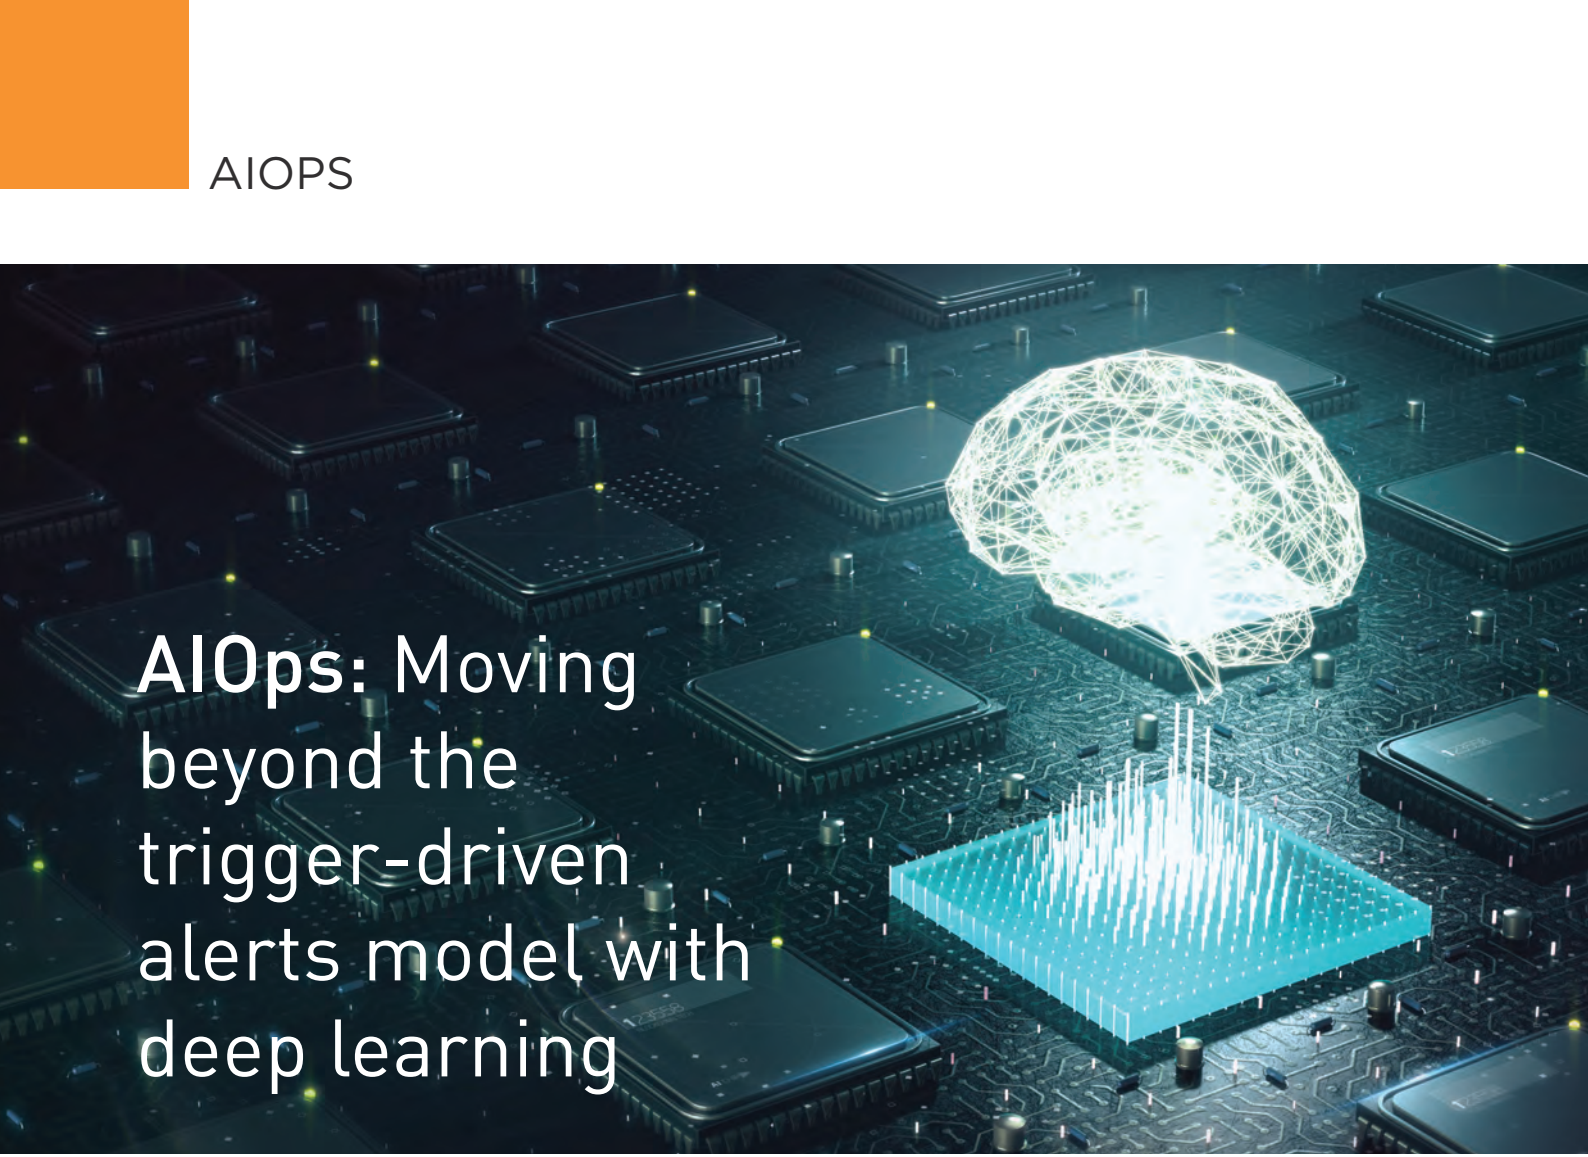
NETSCOUT SYSTEMS has announced a collaboration with Dell Technologies OEM Solutions to deliver solutions designed for medium-to-large-sized enterprises looking for agile, cost-efficient solutions to address their data center and digital transformation needs.

"We share Dell Technologies' expertise to bring enterprise customers the I.T. solutions they need to advance their infrastructures and transform their businesses in a hyper-digital world," stated Michael Szabados, chief operating officer, NETSCOUT. "We provide the visibility required for an exceptional customer experience and the agility to deploy new infrastructures and execute other DX initiatives."

To initiate the collaboration efforts, NETSCOUT certified Dell Technologies' OEM PowerEdge servers and PowerSwitch switches to work seamlessly with its nGenius Packet Flow Operating System (PFOS) and InfiniStreamNG (ISNG) software. NETSCOUT will also provide customers and channel partners with a three-year warranty and maintenance agreement for the bundled PFOS solution, which will be sold through a single SKU for easy ordering.

The two companies will initially work together to sell these solutions in the U.S. to accommodate organizations that have standardized on Dell Technologies platforms in their data centers and offer new customers a simple, reliable, cost-efficient solution.

Plans are in place to extend the solutions offered and expand into international markets soon.



AIOps: Moving beyond the trigger-driven alerts model with deep learning

The standard application of AI to IT operations is to mitigate the “overload of red alerts”: too much noise, not enough signal. This is a valid concern as IT estates grow larger, but the emphasis on individual overload does nothing to address the underlying process issue.

BY SHAUN MCGIRR, AI EVANGELIST, DATAIKU



WITH AIOps, organisations are able to generate more meaningful signals to begin with. The industries leading the way use AIOps to strategically manage vast IT estates and improve the quality of data they are generating for better, more specific insight. And for those on the frontiers, deep learning models have become a key asset, as when applied correctly they can counter the human bias that leads to “tell-me-everything,” trigger-driven alerting.

Telecoms companies face particularly complex IT estates and have already reaped rewards from this application of AIOps to the underlying problem rather than its surface-level manifestation.

To understand this potential value deeply, it's useful to analyse the trigger-driven alerts model, and how large enterprises including telecoms are using AIOps with deep learning to move past it.

The problem with trigger-driven alerts in telecoms

It's time to re-think the problems AI was meant to solve for IT operations. Think about it this way: if every system or service is sending alerts, how do you know which ones deserve attention? This is a tough enough problem to solve, but could easily lead to ever-more complex AI systems built on top of the current alerts ecosystem, which in turn would require alerts.

What we need is to use AIOps to move beyond filtering through thousands of trigger-driven alerts and instead improve the way alerts are generated — it's a bit like saying, if you have trouble locating the needles in the haystack, maybe it's time to get rid of 90% of the hay. This requires proactive improving the quality of the data you generate about service levels, so your operations are running on a higher-octane fuel to begin with.

Telecoms face twin pressures that make this kind of AIOps solution especially relevant: ever-increasing complexity and scale in these companies' global IT operations, yet unignorable customer expectations of increasingly smooth, frictionless experiences across channels.

As enterprises, telecoms can have tens of thousands of discrete IT services running nationally or globally, which need constant monitoring. When companies only had 10 or 20 services to monitor, they put people in a global operations centre with a few dashboards and everything was fine, but there are only so many dashboards a pair of eyes can monitor at once, not to mention all the other potential sources of alerts. These companies know that hiring another 1,000 people to monitor 100,000 services simply isn't realistic, because of the work involved in attending to each service. It's a bit like repainting the Golden Gate Bridge – by the time the crew is done, the next crew is starting on the other side because the work already needs doing again. This is where AIOps can be a lifesaver: but it will require pushing into new techniques such as deep learning.

The Challenges of Using Deep Learning

Ever since deep learning became trendy, companies have been tempted to use it as the proverbial hammer for every nail. Ultimately deep learning models, like any technology, are best applied to particular types of problems (otherwise, they can create more problems than they solve).

Deep learning is a high-octane and high-risk approach in many ways, and time and time again, we see it overcomplicating processes when there may have already been a much more explainable, existing way of doing things. Time series forecasting is a great example of this: models from the sixties worked just fine, but when companies used deep learning to reinvent the wheel with them, they often found worse results at greater cost!

AIOps in Action with Deep Learning

However, there is a place for deep learning – it just needs to be implemented in the right place within AIOps. One of the most advanced AIOps approaches is training a model on the prior history of a given IT service, unlocking value from the vast data assets that IT operations typically collect, especially in unstructured or semi-structured log files.

By training models on the history of service failures, companies can monitor degradation and get early warnings of a likely upcoming failure long before the dashboard turns red, the issue becomes a problem, and the engineers become overloaded with alert fatigue.

It's this highly proactive AIOps-driven approach to creating more relevant signals in the first place that gets telecoms IT departments away from the scattergun approach, and make headway. Using deep

learning for what it's good at, these teams are able to gain enhanced visibility of the very complex processes that may lead to failure in a system, or even the series of events that may lead to customer churn, but without generating alert fatigue.

Deep learning models are able to manage the hundreds, if not thousands of variables that may determine these processes, and excel over other approaches where models quickly become too complex, and where humans can't specify upfront what the model should take into consideration. From a strategic point of view, IT teams know they need to break their IT estates into smaller, more manageable pieces, but may not have appreciated that data scientists could help them get there. If companies can get to this data, organise it and use it in this way, they can use the years of service performance data they have been recording all along, but which has probably only ever been used to generate alerts. With AIOps, that data can be repurposed for something greater: increasing knowledge about the health of the IT estate as a whole.

When companies begin to apply these technologies this way, they can also begin to understand which services are most critical. They can then focus on making those services faster and more resilient, and solve the underlying business problems faster. And, as they get to practice using more complex techniques such as deep learning to improve their service delivery, the cost of each additional deployed model will decrease over time, especially if they can shift the intensive-but-temporary computing workloads required to train deep learning models to a cloud service.

It's about using the technology at hand for what it's good at. The diversity of IT services today is so immense that it is already impossible for any one person or group to manage it single-handedly. With judicious use of deep learning and other advanced techniques within AIOps, companies can manage these vast IT estates across huge networks both on-premise, and in the cloud. Telecoms is only one example of how AIOps is driving efficiencies across billing, customer service, maintenance, and infrastructure.

When companies begin to apply these technologies this way, they can also begin to understand which services are most critical. They can then focus on making those services faster and more resilient, and solve the underlying business problems faster



Seven KPIs for AIOps

IT leaders looking to measure the benefits of AIOps and build key performance indicators (KPIs) for both IT and business audiences should focus on key factors such as uptime, incident response, remediation time and predictive maintenance, so that potential outages affecting employees and customers can be prevented.

BY CIARAN BYRNE, OPSRAMP



BUSINESS KPIs connected to AIOps include employee productivity, customer satisfaction and web site metrics such as conversion rate or lead generation. Bottom line, AIOps can help companies cut IT operations costs through automation and rapid analysis; and it can support revenue growth by enabling business processes to run smoothly and with excellent user experiences.

These common KPIs can measure the impact of AIOps on business processes.

Mean time to detect (MTTD)

This KPI refers to how quickly it takes for an issue to be identified. AIOps can help companies drive down MTTD through the use of machine learning to detect patterns, block out the noise and identify outages.

Amid an avalanche of alerts, ITops can understand the importance and scope of an issue, which leads to faster identification of an incident, reduced down time and better performance of business processes.

Mean time to acknowledge (MTTA)

Once an issue has been detected, IT teams need to acknowledge the issue and determine who will address it. AIOps can use machine learning to automate that decision making process and quickly make sure that the right teams are working on the problem.

Mean time to restore/resolve (MTTR)

When a key business process or application goes down, speedy restoration of service is key. ITOps plays an important role in using machine learning to understand if the issue has been seen previously and, based on past experiences, to recommend the most effective way to get the service back up and running.

Service availability

Often expressed in terms of percentage of uptime over a period of time or outage minutes per period of time, AIOps can help boost service availability through the application of predictive maintenance.

Percentage of automated versus manual resolution
Increasingly, organizations are leveraging intelligent automation to resolve issues without manual intervention.

Machine learning techniques can be trained to identify patterns, such as previous scripts that had been executed to remedy a problem, and take the place of a human operator.

User reported versus monitoring detected

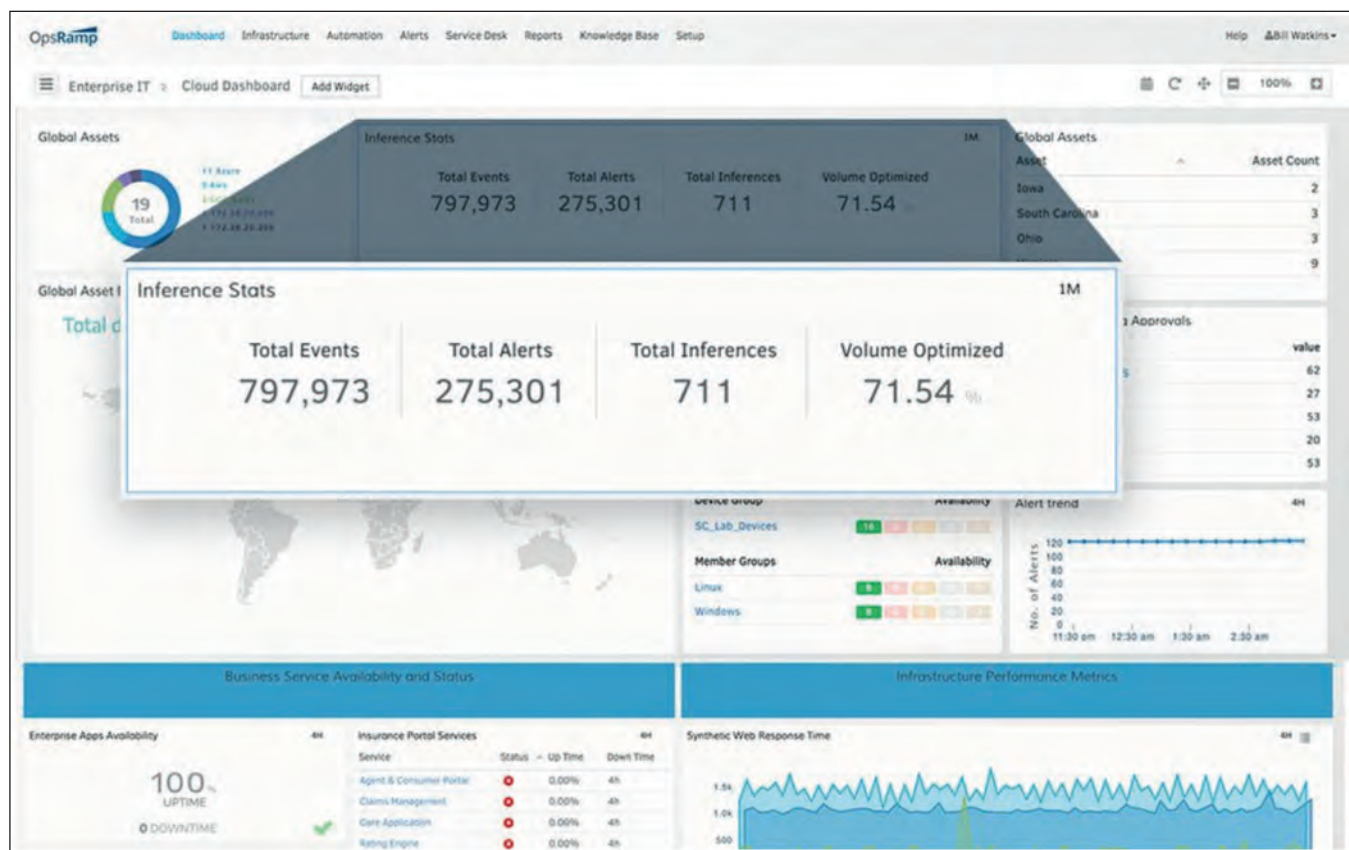
IT operations should be able to detect and remediate a problem before the end user is even aware of it. For example, if application performance or web site performance is slowing down by milliseconds, ITOps wants to get an alert and fix the issue before the slowdown worsens and affects users. AIOps enables the use of dynamic thresholds to ensure that alerts are generated automatically and routed to the correct team for investigation or auto-remediated when policies dictate.


Time savings and associated cost savings

The use of AIOps whether to perform automation or more quickly identify and resolve issues will result in savings both in operator time and business time to value. These have a direct impact on the bottom line.

In summary

These KPIs can be correlated to business KPIs around user experience, application performance, customer satisfaction, improved e-commerce sales, employee productivity, and increased revenue. ITOps teams need the ability to quickly connect the dots between IT infrastructure and business metrics to prioritize spend and effort on real business needs. Hopefully, as machine learning matures, AIOps tools can recommend ways to improve business outcomes or provide insights as to why digital programs succeed or miss the mark.





The importance of visibility in AIOps for digital transformation

The concept of digital transformation is a highly complex one. It means many things to many people, and involves a variety of complex environments and infrastructures from hybrid to multi-cloud, containers and more. To do this at scale, requires automation. To have any hope of managing digital transformation properly, forward-thinking organisations are also looking to machine learning and AIOps.

BY PAUL BARRETT, CHIEF TECHNOLOGY OFFICER, ENTERPRISE, NETSCOUT

AIOps, however, is not a panacea for all of an organisation's digital transformation ills, and what's crucial to making it work as intended is the concept of observability, which I like to think of as the degree to which you can understand the internal state of a system.

Good AIOps and bad AIOps

What separates good AIOps from bad AIOps is the quality of the data that you're feeding into it. If you've got a system that has good visibility, one can generate high quality data about what's happening, which can then be used to underpin machine learning and automation. But if high quality data isn't driving machine learning and automation, one can't expect good results.



The old IT adage of 'garbage in, garbage out' is still as relevant now as it has ever been. Ultimately, automation has been defined by a human being whether they put instructions into a user interface or whether they wrote a configuration file or a script. If we give the wrong instructions, or there's errors in instructions that we give our automation system those errors will get replicated at scale. Without visibility, and the high-quality data that enables, things can go wrong fairly quickly.

Loss of visibility increases the risk of unintended consequences

One example I always like to go back to in my discussions on automation is the stock market crash of October 19th, 1987. One of the primary causes



of which was the collective behavior of automated trading algorithms at scale. For those people who wrote those algorithms at the time, they never expected to crash the stock market, they were trying to make money.

However, no one had foreseen how these individual algorithms might interact with one another when taken out of their comfort zone, i.e. when the markets started to slide.

This is a very real consequence of not having that full visibility, and it's entirely possible to create unintended control loops. If you start linking enough systems together, before you know it you've created a super-system, and one system may have an influence on another system far, far away that you really hadn't intended.

General estimates suggest there are around 30 billion connected IoT devices now, with that number growing rapidly. Once again, that's a call to arms for visibility. Organisations need to make sure that we have independent visibility and clear oversight of how all of these systems are working and interacting with each other. So that, if something does go wrong, and we have created an unexpected feedback loop, that we can stop that and get ahead of it before it becomes a real problem.

Loss of visibility through complexity

Unfortunately, one of the trends we're seeing at the

moment is actually a loss of visibility. This is caused by increasingly complex hybrid multi-cloud environments. As an organisation redistributes its IT infrastructure across a mixture of public cloud, colocation facilities and traditional data centres, an increasing amount of that IT infrastructure is no longer directly under the organisation's control, and that means less visibility.

Organisations therefore have to ask the question of how they can gain insight into what is happening inside these third-party infrastructures, and regain that visibility and control.

Having end-to-end and end-through-end visibility is crucial. For example, the boundaries of different IT domains where there's a change in ownership, or responsibility, is a great place to try and establish visibility.

Although all data is useful for driving AIOps, I believe that packet flows are the most important data source. Packet flows encompass a wealth of information about the status, performance, and security of IT infrastructure. They also provide insight into the experience of end-users and the efficacy of machine-to-machine communications when people and devices interact with each other through apps and services.

With that level of visibility into what's going on in your IT infrastructure, AIOps can be a successful approach that truly revolutionises how businesses operate.

Improving cloud application performance

Masergy has enhanced its Masergy AIOps feature by applying artificial intelligence (AI) and machine learning to optimize Software as a Service (SaaS) applications on global networks. The advancements help companies of all sizes to more quickly and easily solve the problems of application management while also automating IT processes and preventing performance degradation.

Businesses need enterprise cloud applications to be readily available to their employees no matter where they are, and yet it remains difficult for IT teams to ensure a high-performance user experience. Complex IT infrastructures and multi-cloud environments obscure visibility, requiring AI analytics to effectively identify and solve the root causes of performance degradation.

“Unplanned downtime is still largely due to manual processes and human error. AIOps eliminates these challenges, revolutionizing IT operations,” said Zeus Kerravala, Founder and Principal Analyst, ZK Research. “The value of Masergy’s AIOps stems from its ability to evaluate bandwidth usage patterns, identify anomalies, and predict outages all within a fully managed SD-WAN or SASE service. It’s unique because it’s native to the network and security platform, offering prediction and propensity features.”

“Masergy created the industry’s first AI-based network intelligence service that analyzes the network and makes recommendations to enhance reliability. And we remain the only SD-WAN and SASE provider with a fully integrated AI-based network, security, and application optimization solution,” said Terry Traina, CTO, Masergy.

“This is the next innovation and another step forward on Masergy’s path to offering a fully autonomous network,” said Chris MacFarland, CEO, Masergy. “While our clients benefit from the automated analysis and intelligent recommendations of AIOps, Masergy is delivering on the future faster than our competitors. Our AI-powered cloud networking platform is pushing the boundaries of what’s possible.”



“Masergy is a visionary in the arena of SD-WAN and AIOps, and has a clear plan for delivering secure, highly available, self-managed IT and fully autonomous networks,” said Paul Constantine, Executive Vice President of Supplier Services – Digital Distribution at Intelisys.

“As Masergy’s largest Master Agent, we are always pursuing opportunities to bring innovative technologies to our sales partner community, and combined with support from our best-in-class sales engineers, we can together provide industry-leading solutions that address the challenges that business leaders face in the evolving distributed work environment.”

Predictive intelligence and proactive optimization

Masergy AIOps is deeply embedded into the network, security, and application layers and was developed using an unprecedented amount of historic data patterns, leveraging the company’s 20 years of network and security logs. As a result, Masergy is uniquely positioned to provide a more mature AIOps algorithm and therefore the deepest insights for its SD-WAN and SASE solutions.

Because Masergy embedded AIOps into the application layers of its software-defined network, the AI engine now has direct access to more of the data it needs to deliver advanced performance optimization. AI analytics enriched from

traffic flows, applications, and log data offer proactive recommendations for application performance. Here are some of the added features and business benefits:

- Accelerate application troubleshooting and management**
 Masergy AIOps observes application service quality and provides insights that isolate and identify the cause of performance degradation.
- Make data-driven decisions around resource allocation**
 Masergy AIOps observes application propensity for bandwidth consumption, helping IT managers intelligently manage application policies as well as forecast bandwidth and network capacity needs.
- Prevent potential application outages and performance degradation:**
 Masergy AIOps evaluates historic patterns of bandwidth consumption, providing predictions, proactive recommendations, and real-time alerts on application bandwidth utilization.

As a built-in feature, Masergy AIOps is included with all SD-WAN and SASE solutions and delivers insights inside the unified management portal, where clients have real-time visibility and control over bandwidth.

Masergy AIOps was first released in November 2019 as the industry’s first integrated AI-based, digital assistant for network optimization.

Logz.io launches Prometheus-as-a-Service

Logz.io has introduced Prometheus-as-a-Service for their Infrastructure Monitoring solution.

Logz.io's new Prometheus-as-a-Service significantly enhances the company's observability platform and enables engineers to more easily scale cloud and Kubernetes workloads without the overhead or complexity of managing siloed open source tools.

Prometheus, the second project to graduate from the CNCF, is the most popular open source time series database and is praised for its ease of use in cloud-native ecosystems. "451 Research's Voice of the Enterprise: DevOps, Organizational Dynamics 2020 study shows that 58% of organizations are using open source software (OSS) for monitoring," said Liam Rogers, Research Analyst for 451 Research, part of S&P Global Market Intelligence. "Cloud-native application usage in particular has been a driver for open source, fueling the adoption of projects such as Prometheus."

With traditional Prometheus, users often run multiple servers as they scale, have very short data retention and struggle

to investigate incidents. With Logz.io's Prometheus-as-a-Service, engineers can now continue scraping their metrics with Prometheus, while offloading the metrics storage to Logz.io. This has the added benefit of correlating metrics alongside logs and traces all in one unified observability platform.

Prometheus users can now easily send their metrics to Logz.io's Prometheus-as-a-Service by adding just a few lines of code to their configuration files to get started in seconds. In addition, users can also easily import all their dashboards to Logz.io and complete the migration in a matter of minutes.

"We find a lot of value in using Prometheus to provide metrics and insights into our cloud infrastructure, Kubernetes clusters, application performance and beyond, but attempting to run a Prometheus time series database at scale can be incredibly time-consuming," said Daniel Seravalli, Lead Engineer, Holler. "Thankfully, Logz.io helps alleviate this for us by running Prometheus at scale and providing the expertise and support needed to simplify management of the metrics backend."

The launch of Prometheus-as-a-Service reflects Logz.io's strong commitment to empowering the open source monitoring ecosystem, as the product extends Logz.io's unified, open source and SaaS-based observability platform.

Prometheus-as-a-Service accelerates the ability of DevOps and SRE's to unify and correlate logs and metrics for faster troubleshooting, incident response and remediation.

Additional product highlights include:

- Correlate logs, metrics and traces for accurate incident detection and investigation
- Easy migration from Prometheus to Logz.io
- Out-of-the-box advanced dashboards
- Alerts that can combine both logs and metrics to identify and notify of production issues
- Option to add annotations to overlay production events over metrics visualizations
- Sub Accounts to provide the Role-Based Access Control
- Ability to use the best-of-breed open source for logs, metrics, and traces side-by-side on a fully managed platform.

LogicMonitor expands global partner network

LOGICMONITOR has added 11 new partners to the LogicMonitor Partner Network. Of these new additions, eight partners are located in Europe, a region where cloud computing is expected to grow exponentially in the next few years.

According to a study by Global Market Insights, Europe's cloud computing market is anticipated to grow from its current value of \$25 billion to over \$75 billion by 2026. Cloud computing adoption rates continue to grow among European enterprises due to cloud offering a scalable approach that often results in reduced costs, increased operational efficiency and enhanced business agility. European partners are turning to LogicMonitor to give enterprise customers unprecedented visibility into their networks, infrastructure and applications as they begin or continue their cloud migration journeys.

"Enterprises have accelerated their digital transformation initiatives over the past 12 months and are eager to adopt cloud-based monitoring services in order to manage increasingly complex hybrid infrastructures and multi-cloud deployments," said Sanjay Gupta, Global Vice President of Channel and Alliances at LogicMonitor. "Our cloud-based LogicMonitor



platform is indispensable to modern enterprises seeking visibility into and control over every facet of their environments, and we're proud to align ourselves with best-in-class partners around the world to extend its reach."

New LogicMonitor partners in Europe include Amasol in Germany, CDW in the UK, Corporate Finance International (CFI) in Switzerland, Exccon AG in Germany, Kaemi in Germany, Netsecurity AS in Norway, Proact in Europe and the US, and SoftwareOne in The Netherlands. Additional partners joining the LogicMonitor Global Partner Network outside of Europe include Arvensys Technologies in Australia, Total eBiz Solutions in Singapore, and Xylex Technologies in the US.



It's all about scale

Why a delivery app company needed true cloud-native monitoring.

BY CHRONOSPHERE

Constant packet loss

At first, the delivery app company was using a combination of StatsD monitoring for its cloud-native stack and another solution to monitor its virtual machine environment. As the cloud-native environment scaled and developers delivered new features, however, the monitoring system kept breaking down.

"We were experiencing constant packet loss," explained the observability lead. "Developers could break the whole system by making some benign change that has some weird bug and it crashes the whole system." There was an extreme noisy neighbor problem – a developers change could easily impact StatsD's ability to monitor some other, seemingly unrelated application in an entirely unpredictable way. It was so bad, he said, that when he first joined the company no one was actually paying attention to

the metrics generated by StatsD, instead relying on work-arounds like counting log lines. When the team upgraded StatsD, they discovered that they had been losing metrics because of a bug in the system.

"We did the upgrade, and there was an instant change in the pattern," the observability lead said. "That's the bigger problem: The pattern changed. That means we had been flying blind for a long time. We had been seeing a zigzag, but then it became a flat line, which was more or less what we expected." The team had been suspecting that something was off, but hadn't had a way to verify it.

In general, the company thinks of itself as being data driven. Everyone is crazy about numbers, from the CEO down to the newest engineer. Data is used to make better decisions about technology and about the business. If everyone loses observability, it means

the entire company loses that competitive edge. Because software is a core part of the company's product, losing visibility into the application suite was simply not acceptable.

Looking for a solution

As the observability lead started looking for other options, there main criteria were:

Open source. A solution that was based on open source technology was really important because of the team's experience with closed-source solutions in the past. The proprietary format made it difficult for engineers to learn how to use the system and meant that any customization was essentially impossible. Ideally, the solution would have a minimal learning curve and build on technology most engineers were already familiar with.

Scalable. The new solution needed to be able to scale without losing data and without becoming extraordinarily expensive. The company already had a massive StatsD cluster and was experiencing timeouts because it was unable to control the incoming data traffic. It felt like the current solution was at the limits of what it could handle, scale-wise – and yet the company intends to keep growing.

Reliable. Given the problems with data loss, reliability was key. The company was looking for something that developers couldn't break with a seemingly innocent code change and that didn't buckle when asked to scale. Relatedly, eliminating the noisy neighbor problem was important, so that the monitoring system wouldn't experience cascading outages.

Fully distributed. When you're operating at scale, central operations become a bottleneck, explained the observability lead. Dealing with hundreds of millions of datapoints per second could overload the endpoint processes. In his view, the key to being able to scale was having a fully distributed monitoring system. This distributes the load, obviously, and limits the impact of any individual failure. In other words, a distributed system was the only way the observability lead thought that a monitoring system could meet the company's scale and reliability requirements.

Using Chronosphere

"We don't even discuss metric loss," the observability lead said, about the difference between using StatsD and Chronosphere. "Of course, we are always going to monitor our monitoring. But we have a lot more



peace of mind now." In general, engineers are able to set and forget the monitoring tool because they know it is working. Chronosphere's much simpler metrics pipeline has also reduced not just packet loss but all kinds of other operational issues.

Moving to Chronosphere has also underscored just how ad-hoc the company's tagging system had been, and now they are working on a much more unified, consistent naming and tagging practice.

Perhaps most importantly, the company is no longer flying blind. "We have metrics, and we are using them," the observability lead said. The company doesn't need anything fancy when it comes to monitoring, he said – just a solution with a strong foundation that provides scalability and reliability they need.

When you're operating at scale, central operations become a bottleneck, explained the observability lead. Dealing with hundreds of millions of datapoints per second could overload the endpoint processes



The move to SaaS

Gaining visibility over applications to close the security gap

BY VINCENT BERK, CHIEF SECURITY ARCHITECT, RIVERBED TECHNOLOGY



THE RISE in remote working has greatly accelerated digital transformation. Businesses worldwide are migrating to cloud and software-as-a-service (SaaS) applications to maintain efficiency, smooth operations, and optimise team productivity. In fact, Gartner predicts that SaaS revenue will grow to \$140.6 billion by 2022, up from \$102.1 billion in 2019. It's easy to understand this uptick given the cost-effectiveness, low maintenance, and ease of access for users any time, anywhere. However, it's been said that "nobody ever washes a rental car", and enterprises frequently take the same stance toward upkeep of SaaS applications. This approach leaves businesses open to serious security gaps, unexplainable performance issues, and a lack of cost control for a sprawling SaaS fog, all of which are detrimental to business growth.

When operating SaaS applications, businesses are left struggling to detect external threats such as phishing attacks. This is because there can be the perception from enterprises that SaaS applications are inherently secure, and the enterprise is therefore no longer responsible for effectively securing the data they store within them. With a mindset of 'out of sight, out of mind', businesses are failing to put the correct measures in place to ensure they have the right level of visibility over their employees and applications. Without full-fidelity visibility, organisations are increasingly vulnerable to cyberattacks that could result in serious implications if left undetected.

For example, loss of sensitive customer data causing long term reputation damage. As such, it is crucial

– especially during periods of remote working – that organisations use visibility tools to recognise activities unique to the threats to SaaS applications. Only then will they be able to reap the benefits of these applications without the risk. Below, I've set out what companies should keep in mind.

Remote working and the rising threat

As a large portion of the European workforce continues to operate remotely, organisations are faced with the challenge of employees carrying out activity outside of the tightly controlled corporate network. One of the primary drivers for this is that staff are increasingly turning away from business-sanctioned SaaS applications in favour of those that enable them to better maintain and carry out business processes.

For example, an employee may find that sharing files using the fileshare takes too long over the corporate VPN. Consequently, they spin up a WeTransfer account, and begin sharing their files with other employees in this way, all outside of the visibility of the enterprise. This modern form of "shadow IT" is creating a serious security gap as businesses don't have any visibility over unsanctioned applications used by their workforce and can therefore not secure them.

Similarly, with the direct path to SaaS, using personal, or at least unmanaged devices from a corporate perspective, it has become significantly harder to detect when individual users have been compromised. After all, the data in the SaaS application is only as secure as the user accounts that have access to them. Phishing, malware, and nefarious apps and browser extensions may all be vectors through which the users' devices may be compromised, giving a direct path to the data in the SaaS applications that user has access to. Given the fact that it is only a matter of time before any business falls victim to a security compromise this lack of visibility is simply unacceptable.

Narrowing the gap

It is clear that SaaS applications are helping to improve employee productivity and boost the efficiency of business operations during remote working. In fact, collaboration applications such as Zoom and Slack have experienced a drastic increase in usage as employees continue to use them to maintain communication. Zoom alone has grown its customer base nearly five-fold since last year but to fully realise the benefit of these applications, the security risks inherent in SaaS must be addressed.

The first step in overcoming these challenges is carefully considering which SaaS applications the business wants to deploy. When making these decisions companies should remember that the most affordable choice is not always the best. Take some of the most popular video conference applications, for example, which may not possess

end-to-end encryption. Without this, the enterprise is left vulnerable to interception or eavesdropping on confidential business calls and obtaining information that can later be used to perform spear phishing campaigns. However, these kinds of flaws or the resulting exposure are often not factored into the business decision making. To redeem control, and bolster security, enterprises must assess which applications they are deploying and whether they are able to gain visibility over them. In doing so, organisations can minimise disruptions to employees and maintain smooth business operations.

There is a related issue with respect to a cost-first approach to SaaS selection. The ease of deployment and re-deployment of SaaS may lead companies to feel that they can make cost gains through flexibility, choosing one SaaS offering one year and switching to another the next. While there may well be significant cost reductions in spend, every change represents a window of significant security vulnerability. Corporate IT needs time to learn how to monitor new systems and understand what normal access patterns and access needs look like, and users are particularly vulnerable to phishing while they are still learning how to interact with new systems. Continual changes can leave a company constantly vulnerable to attacks that would otherwise be easy to resist.

However, to truly begin closing the security gap, businesses also need to record and collect as much data from across the virtual enterprise as possible. This includes monitoring the traditional network border, the user's laptop, and their application log in. The more information that can be collected from as many angles as possible, the better prepared the enterprise will be to spot compromises, espionage, or sabotage. Besides security solutions, performance issues of SaaS applications are hurting performance and businesses must consider dual-use solutions such as network performance management tools to help them achieve this duality of visibility for performance and security. With these mechanisms in place, the enterprise can gain visibility over SaaS application usage and forensically analyse the data collected to identify any concerning activity. This will enable businesses to quickly resolve any issues and ensure their employees and the applications are working securely and efficiently.

Owning your security and gaining visibility

SaaS applications are the future. It is therefore vital that businesses recognise their responsibility to secure them and make the necessary investments to achieve this. This includes adopting technologies that will provide much needed visibility over applications and help IT teams to close security gaps across the business. In doing so, enterprises will be better positioned to keep business operations running smoothly – safe in the knowledge that they are able to detect and overcome any unwelcome security threats.

Observability vs. Monitoring:

What's the difference?

One of the more delicate debates in the DevOps world is what observability has to do with monitoring.

BY CHRIS TOZZI, SUMO LOGIC.

IS OBSERVABILITY just a trendy buzzword that means the same thing as monitoring? Is observability an improved version of monitoring? Are monitoring and observability different types of processes that solve different problems?

The answer to those questions depends in part on your perspective. Let's take a look at the different ways of thinking about observability and monitoring, and what they have to do with each other.

Monitoring and Observability: A Brief History

Long before anyone was talking about observability, they were talking about monitoring. Modern software monitoring tools, like Nagios and Zabbix, trace their origins to the late 1990s or early 2000s. That's when IT teams first began deploying tools that systematically monitored application environments for problems and sent alerts when something seemed amiss.

Monitoring tools have grown more sophisticated since then. Instead of merely pinging servers to see whether they respond, or tracking CPU and memory usage to see if they are maxed out, modern monitoring tools use

sophisticated analytics to identify problems within complex software

environments. Nonetheless, the core foundation of monitoring remains the same: it's all about finding problems.

Observability is a much newer term, at least within the context of software. Although you can find occasional references to the observability concept from the earlier 2010s, it wasn't until circa 2016 that observability came into vogue among Site Reliability Engineers (SREs) and DevOps teams, as The New Stack explains.

Probably because the term is so new, observability is harder to define in a clean way than monitoring. Different people and companies offer different definitions. Most definitions of observability, however, focus on the idea that observability means collecting actionable data in order to gain a true understanding of problems identified by monitoring tools.

Similarities between Observability and Monitoring

At a high level, then, observability and monitoring share some key characteristics. For one, they both help support the overall reliability and performance of software environments. They help find and (in the case of observability, at least) fix problems.



They are also similar in that they both leverage multiple data sources. Although monitoring tools may not perform the advanced data correlation and analytics of observability tools, and monitoring tools usually rely primarily on metrics, they can leverage other types of data points – including logs and traces – to collect data from a software environment. Observability solutions, too, leverage a variety of data sources to gain the depth of visibility necessary to understand complex problems.

How Is Observability Different from Monitoring?

From that perspective, at least, the difference between monitoring and observability boils down to the end goal. Whereas monitoring focuses on finding problems, observability focuses on understanding and resolving them.

To put this into context, generating an alert when a node fails in your Kubernetes cluster would be an example of monitoring. That's simple enough to do: you'd run an agent on the node itself (or possibly within Kubernetes) that would monitor whether the node is responding, and trigger an alert if it's not. In contrast, observability of the node failure would entail collecting data from multiple sources – the node, the hypervisor hosting it (if it's a virtual machine), the various Kubernetes services that were interacting with the node, and the pods hosted on the node – and then correlating those data sources to determine why the node failed. Maybe it was hosting a pod that suffered a major crash and took the node with it. Maybe it had a kernel panic at the OS level. Maybe it was shut down by an IaaS provider because you didn't pay your bill. Whatever the root cause of the issue, monitoring alone wouldn't expose it, but observability would.

Using Monitoring and Observability Together

In general, DevOps teams seem to agree that observability and monitoring are distinct types of operations that address different problems. Still, monitoring and observability are inseparably linked. Monitoring tools can tell you when something goes wrong, and observability tools can help you investigate the issue after you detect a problem. Pairing monitoring and observability is beneficial because not all problems identified by monitoring tools require sophisticated investigation. Maybe your monitoring tools send an alert telling you that a server went offline, but it was part of a planned shutdown, for instance. In that case, you need not collect and interpret multiple types of data to understand what happened. You can just log the alert and move on.

But when serious problems arise and you need to troubleshoot them quickly, observability data is crucial. Although you could technically collect the same type of data manually that observability tools deliver automatically, the data collection would take

You could argue that observability is a buzzword in the sense that teams were performing observability before the term existed. The idea of analyzing diverse data sources in order to understand complex application performance and reliability issues is hardly new. Indeed, platforms like Sumo Logic, which was founded in 2010, have been enabling this type of work for more than a decade – long before it became trendy to talk of observability

time and delay incident resolution. Observability tools ensure that you always have the data you need on hand to interpret a complex problem. Many solutions also offer recommendations or automated analyses that can help teams sift through complex observability information and identify root-cause problems more efficiently.

Note, too, that a system must be observable – meaning that it must generate data that can be collected and interpreted – in order to be monitored. A closed system that offers no facilities for data collection (such as a proprietary IoT device that generates no metrics) can't be monitored for failure because it lacks observability. And, of course, it can't be observed, either, because there is no data to use as the foundation for observability.

Is Observability Just a Buzzword?

You could argue that observability is a buzzword in the sense that teams were performing observability before the term existed. The idea of analyzing diverse data sources in order to understand complex application performance and reliability issues is hardly new. Indeed, platforms like Sumo Logic, which was founded in 2010, have been enabling this type of work for more than a decade – long before it became trendy to talk of observability.

On the other hand, the buzz surrounding observability does seem to reflect some important changes that have taken place in the DevOps ecosystem in recent years. For one, the emergence of endlessly complex cloud-native environments has made the ability to

collect and correlate complex data sets absolutely critical.

When all of your applications ran as monoliths on virtual machines, simply monitoring them and then investigating issues manually may have been enough to keep everything running. In the cloud-native world of microservices, containers, serverless functions, and scale-out infrastructure, however, you need deeper visibility than you can achieve through monitoring alone.

At the same time, the observability trend perhaps reflects increased awareness that DevOps practices should be tied to business goals. Monitoring helps technical teams identify technical problems, like a server that is down or an application that is not responding. Observability provides finer-tuned insights into application behavior and performance. It can help teams home in on a microservice that is causing problems for a particular subset of users, for example, or a resource allocation that is insufficient for handling application demand at certain times of day.

In these ways, observability helps align technical issues with business issues. It makes it clear which

business priorities (like serving a certain user group or handling a business event that leads to a spike in traffic) are impacted by technical shortcomings. In contrast, monitoring provides only high-level indications that something is wrong, without offering much opportunity for measuring the specific business impact of the problem.

Conclusion: DevOps Teams Need Monitoring and Observability

You could go on and on in debating the precise similarities and differences between monitoring and observability, or arguing about whether observability is actually a useful term or a mere buzzword. But that's not going to improve the outcome of your DevOps processes.

Instead, perhaps the best way to think about monitoring and observability is to approach them both as must-haves for modern DevOps teams. Yes, there may be some nuance surrounding how you implement them or exactly what each term means. But viewed from the perspective of DevOps operations as a whole, you need both monitoring and observability to find and fix application issues and, in turn, ensure an optimal end-user experience.



DIGITAL PACKAGES

CHOOSE one of our bespoke Digitalisation World multi-channel marketing packages to maximise your coverage and exposure to the industry at the most cost effective rate.

3 MONTH PACKAGE

- 3 Sponsored newsletters of your choice (2 banners & sponsor message)
- Run of page section sponsor (3 month banner)
3 x Static/Animated banners Digitalisation World magazine (3 Issues)

Separate Costs: £7,155

Package Cost: £4,995

6 MONTH PACKAGE

- 3 Sponsored newsletters of your choice (2 banners & sponsor message)
- Run of page section sponsor (6 month banner)
3 x static/animated banners Digitalisation World magazine (6 Issues)

Separate Costs: £14,310

Package Cost: £8,995

12 MONTH PACKAGE


- 12 sponsored newsletters of your choice (2 banners & sponsor message)
- Run of page section sponsor (12 month banner)
3 x Static/animated banners Digitalisation World magazine (10 Issues)

Separate Costs: £28,620

Package Cost: £16,995

Other package combinations can be put together on request:

- **Webinar:** Management and/or delegate package
- **Research:** ask questions of our readership... qualified results from targeted demographics
- **Expert blogs:** independent writing by Digitalisation World editorial staff for your company, leverage our expertise



Lift-and-Shift cloud migrations

The good, the bad, and how to avoid the ugly

BY RICARDO NEGRETE, SR. DIRECTOR SERVICES AND TECHNOLOGY, VIRTANA.

THERE ARE MANY PATHS to the cloud, and the one you choose depends on your particular digital transformation requirements and resources. About a decade ago, Gartner cleverly developed an alliterative nomenclature to describe five different migration strategies: the five Rs. That list has evolved over time and there a lot of 5-, 6-, and 7-strategy variations out there. Here's one version:

- **Rebuild** – rewrite the application from scratch to be cloud-native
- **Replace** – swap out the application entirely with a new, cloud-native application
- **Refactor** – rearchitect parts of the application so it better supports the cloud environment
- **Replatform** – move the application with minimal upgrades to take advantage of cloud benefits
- **Rehost** – move the entire application as is to the cloud

Rebuild and Replace both ditch the current on-premises application with the objective of implementing business logic that enables significant strategic

advantage; the difference, compared to the other options, is the time that it takes to implement it, who architects and re-writes the code, who supports it, and who selects the new cloud version.

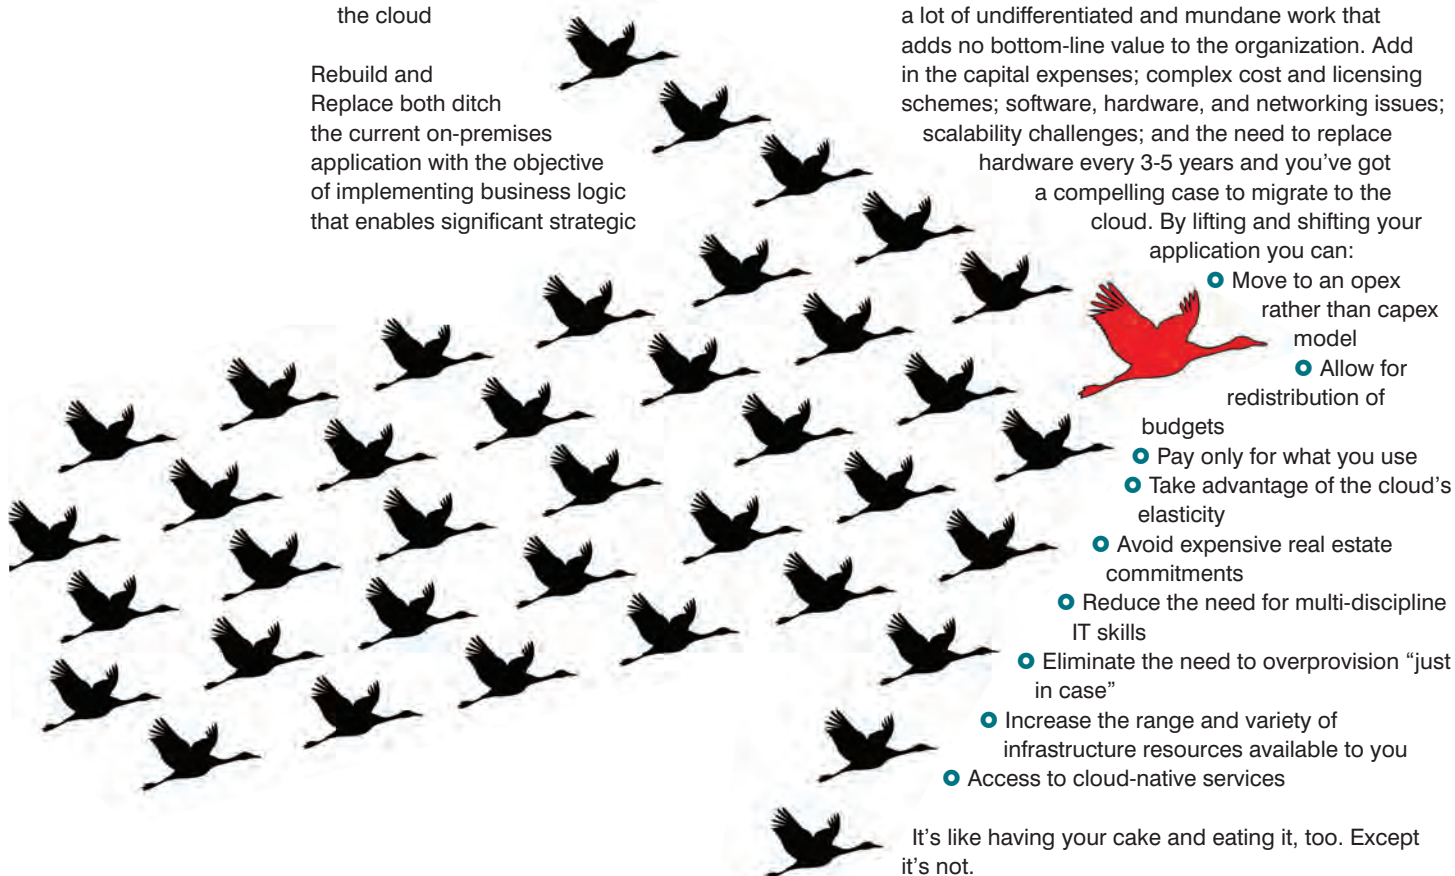
Replacing is a move to SaaS – a top choice for many enterprises. In fact, Gartner projects worldwide SaaS spending to hit \$177.78 billion in 2021. Why? Enterprises view it as the path of least resistance. Their goal is to migrate business functions not infrastructure, and to take advantage of enabling technologies that can deliver a competitive advantage (AKA digital transformation). But Replace is not always a viable option, which is why Rehost, often referred to as “lift and shift,” remains a popular approach.

The upside of lift and shift

Maintaining applications on premises necessitates a lot of undifferentiated and mundane work that adds no bottom-line value to the organization. Add in the capital expenses; complex cost and licensing schemes; software, hardware, and networking issues; scalability challenges; and the need to replace hardware every 3-5 years and you've got a compelling case to migrate to the cloud. By lifting and shifting your application you can:

- Move to an opex rather than capex model
- Allow for redistribution of budgets
- Pay only for what you use
- Take advantage of the cloud's elasticity
- Avoid expensive real estate commitments
- Reduce the need for multi-discipline IT skills
- Eliminate the need to overprovision “just in case”
- Increase the range and variety of infrastructure resources available to you
- Access to cloud-native services

It's like having your cake and eating it, too. Except it's not.



The potential pitfalls

With lift and shift, you essentially just pick everything up on top of the operating system from an environment and move it to another. The problem is that the cloud destination is a very different environment from the on-premises starting point, and there are some dangers. Here's a simple analogy:

Let's say you're moving to a new house. It's pretty easy to lift and shift your display case of hand-cut crystal from one single-family home to another single-family home. The living room may be a different size and shape, but as long as you have room for it, you shouldn't run into any problems. But, if you are instead moving to a houseboat, you can't just plop it down and call it a day – one small wave could tilt the whole thing and reduce your collection to a pile of smashed glass.

When it comes to lifting and shifting to the cloud, there are a couple of potential issues. One is that your workloads may not perform as well in the cloud, sometimes creating latency problems that prevent you from preserving the original on-premises SLAs.

Or, you might have solved a high I/O block storage demand challenge for a workload by running it in a physical server using multiple HBAs, but you may not realize that it will create an unexpected big cost increase in the cloud, or worse it may be unfit to run in the cloud due to unavailability of computes able to handle the aggregate of all the original I/O spread across the on-premises HBAs. And, of course, by choosing rehosting over the other Rs, you're choosing the option that takes the least advantage of cloud-native benefits. In the worst-case scenario, your migration effort fails and you have to repatriate your workloads.

Unfortunately, this is far from rare. In fact, a recent survey we conducted of 350 cloud decision makers found that a stunning 72% of respondents stated they've had to move at least some of the applications back on-premises after migrating them to the public cloud.

The importance of planning

Despite the risks, lift and shift is a viable option – as long as you perform a comprehensive, data-driven assessment of your workloads before embarking on the migration. At a minimum, you need to take the following three key steps:

- Baseline assessment
- Application dependency mapping
- Cloud mapping and cost assessment

Baseline assessment

A baseline assessment uses time-series data captured during peak hours to understand your workloads in the current on-premises environment, such as:

- **Health** – Are there anomalies, such as I/O aborts or chronic overloads, affecting the workloads?

- **Utilization** – Are there any workloads, such as “bully” or Zombie virtual machines, that are unfit for migration? The utilization time-series data is critical for rightsizing the compute and storage services required in the cloud.
- **Performance** – What is the minimum acceptable latency the applications can tolerate?

The baseline assessment creates the reference point for setting utilization demands, performance requirements, and SLA expectations in the cloud.

Application dependency mapping

Before you start your migration, you need to be able to answer the following:

- Which applications can migrate to the cloud?
- What should be moved together?
- In what sequence should workloads and services be moved?
- Which workloads and services should remain on premises?

The application dependency mapping process reveals the landscape of services supporting your workloads, including the workload boundaries, external/internal communication, type of network services, port numbers, bandwidth, egress cloud traffic, etc. This information helps you fully understand your workload interdependencies and consider many “what if” scenarios so you can answer those key questions.

Cloud mapping and cost assessment

Now that you understand your workloads in their on-premises environment, it's time to turn your attention to their final destination. A rightsizing algorithm analyzes inventory data (such as storage size and relevant utilization and temporal characteristics of the targeted workloads) as well as all relevant time-series measurements (CPU MHz, memory, read/write throughput, IOPS, and network utilization) and offers mapping options based on as-is, peak, or 99th/95th percentile values from the time-series measurements.

The mapping selections will identify the best fit for compute and storage resources for the lowest cost from the latest published options for the selected cloud service provider (CSP) regions.

If you're going to lift and shift, you need to know before you go

The bottom line is that lift and shift is a viable cloud migration option that has benefits. But once you've made that decision, you need to set yourself up for success by getting data-driven information before starting implementation.

Virtana can help. With the Virtana Migrate module, which is part of the Virtana unified hybrid cloud optimization platform, you can get your lift-and-shift cloud migration right the first time and every time.

Honeycomb launches refinery

HONEYCOMB has introduced a new solution, Refinery, to help enterprises refine their observability data at scale. Managing large, complex applications means dealing with potentially hundreds of billions of events per month. Refinery collects 100 percent of telemetry data and provides multiple ways to observe only the events that best represent important system changes, resulting in a high-fidelity debugging experience while also controlling costs.

“Refinery has reduced our event volume to 25% of what it used to be, and that reduced traffic is also more consistent on a daily basis,” said Matt Button, Engineering Team Lead, Geckoboard. “Now my team has confidence that we can use Honeycomb even more in our services to better understand what’s happening, without worrying so much about our spend.”

With the complexity of modern software architectures, enterprises need observability to identify and resolve issues in their production apps faster and to build more resilient services. Refinery is the latest addition to Honeycomb’s Enterprise offerings that give developers, site reliability engineers (SREs), and DevOps engineers the tools they need to quickly troubleshoot some of the world’s largest and most complex software applications. In addition to Refinery, Honeycomb’s enterprise offering includes:

Service-Level Objectives (SLOs) that help technical and business stakeholders align on critical customer experience thresholds. Unlike other SLO solutions that only provide alerts and force teams to debug in other systems, Honeycomb’s SLOs are debuggable, so that teams know what is wrong and how to fix issues within a single interface.

Secure Tenancy, Honeycomb’s patented approach that enables organizations to protect sensitive data and comply with their off-premises data security policies

Exclusive training and onboarding support needed at enterprise scale. Customized onboarding packages and top-level support improve an organization’s production service

management skills and accelerate their ability to resolve issues.

“At large enough scale, there’s a tradeoff between providing engineers the data they need to maintain the kind of digital experiences customers expect and the resource costs of sending all that data to Honeycomb,” said Christine Yen, CEO of Honeycomb. “Honeycomb is always looking for ways to democratize the tools and expertise that used to be reserved for only the world’s most elite companies. Previously, solutions like Refinery have been proprietary and closed-source. Now, Honeycomb makes it easy for anyone to only keep the most important data they need to debug their production services and stop paying for the rest.” Refinery delivers significant benefits for the enterprise, including:

Cost predictability and control without sacrificing fidelity. Capture 100% of your observability data and use a variety of out-of-the-box configuration options to determine which events best represent important system changes. Honeycomb uses those samples to rebuild data accurately, providing a high-fidelity observability experience for a fraction of the cost. The control customers need, backed entirely by Honeycomb support. Refinery runs within a customer’s infrastructure, giving them full control over what they choose to send to Honeycomb. Honeycomb supports every aspect of using Refinery, including installation, configuration, usage, troubleshooting, maintenance, and more.

Works with Secure Tenancy and OpenTelemetry. It’s critical that all components of an enterprise stack work together seamlessly. Refinery simply just works with our patented approach to data privacy and with vendor-neutral instrumentation options like the OpenTelemetry project.

“Modern applications and microservices require more data than ever for effective troubleshooting based on system observability, which can become prohibitively expensive,” said James Governor, RedMonk co-founder. “With Refinery, Honeycomb is addressing these cost concerns to encourage wider enterprise adoption of its web scale technologies.”

GAVS Technologies launches the GAVS Healthcare Technology Institute

GAVS Technologies has launched the GAVS Healthcare Technology Institute. The institute was inaugurated by Pramod Bhasin, Founder, Genpact and ex-Chairman, NASSCOM, and some of GAVS’ partners.

GAVS, along with its partners, is focused on enabling Healthcare companies with technology-led business outcomes from Digital Transformation that are driven by the true power of AI and Analytics.

The curriculum of the institute is structured to cater to the growing need for Healthcare Technologists who understand the intersection of technology, business, and human-centered care. Some of the key focus areas will be – AI in Healthcare, Population Health Analytics, MedTech, Epidemiology, Telehealth, etc.

“We at GAVS are committed to evolve as a key player in the Healthcare space. To that end, we have launched GAVS Healthcare Technology Institute to enable technologists with Healthcare domain expertise.” – Sumit Ganguli, CEO, GAVS Technologies
 “This Healthcare Institute by GAVS will play a crucial role in enhancing the Healthcare domain expertise within GAVS, Chennai and in India.” – Pramod Bhasin, Founder, Genpact and ex-Chairman, NASSCOM.

Lightstep introduces 'Change Intelligence'

LIGHTSTEP has launched Change Intelligence. While AIOps is frequently advertised as a way to dramatically improve IT operations, it often fails to deliver, in part because today's software applications are too dynamic: software, infrastructure, and user behavior are all changing faster than ever before.

Lightstep has embraced those changes and put them at the center of analyzing application performance with Change Intelligence, inspired by work done by the Lightstep CEO Ben Sigelman on Google's Monarch project. "Lightstep is the best observability vendor that I have ever worked with," said Trevor Rundell, Director of Engineering at Drift. "We're able to get to the root of what caused a change in our distributed systems with just a few clicks."

The problem with 'AIOps' is that it only works by finding patterns in data, in order to better predict recurring situations. While this may be useful in traditional monitoring dashboards that have alerts set on predefined thresholds and predictable situations, the modern practice of observability acknowledges that it's not possible to know everything

that will go wrong in advance, especially in microservices and distributed environments. "Change Intelligence is the ultimate added value that any engineer can get out of their observability system," said Sam Wierema, Senior Engineering Manager of Infrastructure at MessageBird, one of the fastest-growing software companies in the world.

The X-factor that makes Change Intelligence possible is Lightstep's groundbreaking time-series database, which can process over a trillion events each day and is built by the same engineers that worked on the Monarch project at Google. Monarch is the globally-distributed in-memory time series database system in Google, that is used internally to monitor the availability, correctness, performance, load, and other aspects of billion-users scale applications and systems at Google. By tightly integrating these metrics with Lightstep's existing distributed tracing data, engineers can connect cause and effect faster than what was previously possible.

"We took inspiration from the technology we built at Google, took it to the next level, and made it generally available to

all Lightstep users," said Ben Sigelman, Co-Founder and CEO of Lightstep. "With Change Intelligence, any developer, operator, or SRE can instantly understand changes in their service's health and – most importantly – what caused those changes. In this way, we're able to actually deliver on the promise of AIOps: to automate the process of investigating changes within complex systems." Modern DevOps teams need the ability to answer novel questions about the state of their systems in real time. Until today, understanding what changed and why has largely been an arduous, manual process, with engineers combing through dashboards trying to correlate disparate data to arrive at a theory.

"The truth is that companies are already drowning in data from dashboards, alerts, endless logs – as an industry we're so afraid of missing data that we collect all of it," said Daniel 'Spoons' Spoonhower, Co-Founder and Chief Architect of Lightstep. "But from talking with customers, we've learned that this can create just as many problems as it solves. We're not looking to add more data or noise – we're looking to find root causes and resolve issues faster."

Graylog launches Graylog cloud

GRAYLOG has announced the availability of Graylog Cloud. With the full features and functionality of Graylog Enterprise, Graylog Cloud offers choices to customers who want seamless data collection, rapid search, flexible analysis, and greater affordability without the hassle of maintaining and updating the systems it runs on.

"Unlike other offerings in the market, Graylog Cloud is on par with our Enterprise version day 1 – this isn't a rebuild or subset of capabilities - this is a full-featured log management platform with all of the benefits the cloud has to offer," said Lennart Koopmann, founder, and CTO of Graylog. "Infrastructure costs associated with on-premise deployments are a consideration when customers select a log management platform. For those who want to decrease these TCO costs, Graylog Cloud is a viable choice."

For IT professionals who solve security, operational, and application support issues, Graylog Cloud offers infrastructure cost savings coupled with immediate operational agility so organizations can quickly find meaning in data and take action. Graylog Cloud offers:

- Industry-leading 90 days of live data and 1 year of archived data included, ensuring everything needed for daily work is at the ready
- 99.9% uptime SLA provides confidence that Graylog Cloud is there when needed
- Full control of the customer environment for maximum flexibility and speed in adding new data sources, processing rules, search templates, alerts, visualizations, reports, and unlimited users
- The assurance of a SOC2 Type 2 certified environment
- All the features of Graylog Enterprise v4.0, including Graylog Illuminate for pre-built data visualizations of supported data sets
- Built-in integration with IPinfo for fast, reliable IP address information that automatically enriches log data processed by Graylog Cloud

According to Research Firm MarketsandMarkets, the global log management market is expected to grow at a CAGR of 14.1% from 2020 to 2025, with the cloud delivery portion of the market growing fastest over that period.



The data-driven IT operations organization

Data analytics has always been the foundation to drive actions for IT operations—supporting capacity planning, resource optimization, workload rebalancing, cost projections, and security predictions. Now, there are new demands on IT operations to deliver inclusive data intelligence for managers across IT and the business.

BY BHANU SINGH, SVP ENGINEERING AND DEVOPS, OPSRAMP

ITOps TEAMS have a growing mandate to positively influence business and customer-related outcomes; their ability to collect, analyze and use data to support many types of decision-making is now top-of-mind. Amid digital-first business models, IT organizations have also been leading an accelerated migration to cloud infrastructure. The opacity of managing heterogeneous environments and hybrid infrastructure where numerous containerized or virtualized layers are dynamically reconfiguring has complicated the ITOps job immeasurably.



IT operations teams must build a data-driven view of the organization from the customer perspective: cross-cutting different integrated and siloed business and technology functions. These pain points of traditional

data analysis methods must be resolved to align with enterprise goals:

Lack of customer focus: IT operations can't understand the customers' holistic needs along with short and long-term impacts on customer experience. This means a potentially large disconnect between the business (product or service) and the market (customers).

Wrong strategy: Broken visibility also leads to poor investments. The organization will under or over-invest in key areas such as development, security, infrastructure and monitoring, without an accurate, holistic picture of where the important gaps exist and how everything ties together.

Reactivity: ITOps must carry on with a reactive approach, operating blindly and moving from solving one issue to the next without a big-picture view of priorities and needs.

Competitive drain: We see daily examples of enterprises going out of business because they didn't adapt to changes from digitalization—a stark reality in 2020. As just one example, take ClassPass, a leading fitness and wellness subscription service/app which allows customers to sign up for in-studio or digital live streaming fitness events from boutique studios and spas all over the country. The company has raised \$549 million and booked 100 million reservations since its founding in 2013. Meanwhile, traditional gym chains without viable digital offerings such as 24-Hour Fitness and Gold's Gym, fared poorly in 2020, shuttering locations and filing for bankruptcy. It's easy to find similar examples across restaurants, hospitality/ lodging, retail and banking, among others.

A New Technology Framework for Holistic Data Analytics in IT

The advancement in affordable technology innovation on compute, storage, AI, and ML has made it possible to derive meaningful insights from disparate, discrete, and federated data sources. IT operations teams must build a data-driven view of the organization from the customer perspective: cross-cutting different integrated and siloed business and technology

functions. Gaining this data-driven operating environment, however, requires fundamental changes in technology and process:

- This begins by leveraging the discovery, monitoring, APM, networking and AIOps tools in a cohesive and integrated way.
- The goal is to gain end-to-end contextual visibility in near real-time and apply data analytics to drive the best decisions and actions.
- IT leaders will need to develop an overarching strategy for digital operations tools selection and tailor skill sets and create new organizational roles necessary to get the ROI.
- Domain-specific processes must go away as part of this effort.

Consider an ecommerce business which needs to track metrics such as shopping cart abandon rates and transaction time. It's IT's job to understand how to collect, analyze and correlate the right data from infrastructure resources to those business metrics. The goal is to use analytics to identify website bottlenecks caused by poor-performing systems or where a user process could be better by, for instance, reconfiguring the cloud architecture. There are similar stories across other verticals like healthcare, banking and hospitality. Understanding business KPIs by leveraging customer and transactional data will help the organization create the right data strategy and select the best analytics and automation tools.



DW **DIGITALISATION WORLD**

Modern enterprise IT - from the edge to the core to the cloud

New product and process development is the foundation for the growth of the DW industry.

If you want to highlight the recent important breakthroughs that your company has made, please submit an abstract to philip.alsop@angelbc.com

It is imperative that Digitalisation World Magazine remains a timely resource for this industry, so we are especially interested in highlighting very recent work.



Reducing operational costs and improving productivity in a post-Covid world

How the aftermath of the pandemic has caused a shift in IT operations worldwide.

BY ANN HALL, MARKETING MANAGER, HEAL SOFTWARE INC.

THE CURRENT STATE of IT Operations has shifted to the widespread adoption of multiple AIOps Tools within an enterprise, each with niche capabilities such as observability and proactive anomaly detection. A single tool is frequently insufficient to deal with the increasing complexity of hybrid digital environments and multiple tools are needed to provide visibility across disparate silos. Operations teams are making the shift from monitoring / managing alert storms and disconnected events via dashboards, to focusing on root causes of issues in various dimensions like network, infrastructure, end user transactions and application code via analytics and intelligent event correlation.

2020 forced many teams out of their erstwhile comfort zones. While the pandemic made businesses rethink their overall organizational budgets, one area where expenses had to be pruned was IT Operations. Traditionally necessitating the presence of large teams – mostly comprising contract workers – physically co-located at data centers, teams were now forced to start working remotely from disparate locations due to offices being closed over Covid-19 concerns. On one hand was the challenge of having to operate with skeletal staff, and on the other, the heightened need to keep enterprises running 24x7 with virtually no downtime; most businesses from banking to retail storefronts started operating exclusively online, with

branches and brick-and-mortar stores being shut. IT Operations teams were ill-equipped to straddle both these seemingly paradoxical requirements.

While 2020 was a wake-up call to teams to optimize IT Operations expenditure and look judiciously at making data center costs leaner, 2021 will be the year where these lessons are put into practice. In this whitepaper we look at the consequences of the pandemic on IT Operations and the steps that can be taken in 2021 to ensure that enterprises implement cost-cutting measures effectively.

The Pandemic Effect on IT Operations in 2020

1. Leaner Teams

Most jobs in Network Operations Centers are contract jobs. The first step that most organizations took in 2020 was to downsize the variable pool of talent. Leaner teams were expected to perform a similar quantum of work, which in turn emphasized the need to move to greater automation in ITOM. Challenges in 2021: Attracting the right talent, doing more with less, introduction of automation in business processes to make do with smaller teams.

2. Remote Management

Businesses were forced to give up leased and rented office spaces to cut back on expenses – and this affected IT Operations teams in a big way. Skeletal staff were deployed in the remaining office(s) and the rest of the team was expected to move to a work-from-home mode of operation almost overnight. Managers had to oversee teams remotely and ensure that all employees had the required resources and know-how to use VPN, online meeting tools and collaborative solutions. Gartner predicted that services such as apps, platforms and equipment spent on remote workers would increase by up to 19% in 2020, with this number potentially being even higher by the end of the year and in 2021. One of the largest areas of cloud to experience an increase in demand was desktop-as-a-service (DaaS), with a predicted growth of 95.4% in 2020. DaaS is an inexpensive option for organizations that are looking to support their remote workers by providing secure access to enterprise applications remotely.

Challenges in 2021: Greater security, availability of right talent with the required software and hardware resources, need for collaboration among teams.

3. Automated Deployments

In 2020, organizations had to rethink their deployment strategies and come up with mechanisms for remote and automated deployments due to staff shortage and absence of a centrally located workforce. This created the need for dedicated automation architects who could guide DevOps and SRE teams on the use of Agile practices for breaking down organizational silos between software developers and IT operations personnel. DevOps and Agile are

collaborative practices, focusing on faster deployment, maintenance and upgrades, and automated recovery and failover in case of issues. Adopting continuous deployment allows applications to be quickly brought up in several different environments.

Challenges in 2021: Adoption and continued use of DevOps and Agile practices, increased automation in the application deployment and maintenance process.

4. Revamping Application Infrastructure

Cost cutting measures extended to infrastructure as well, with organizations choosing to review their spending on dedicated hardware and moving to cloud (virtualization, microservice and container-based architectures) as well as assessing the number of software solutions being deployed in the enterprise and seeing if a switch to open-source was possible. Virtualization reduces the number of physical servers required in the enterprise, a container's smaller infrastructure footprint provides a more secure and easier to configure environment and microservices make applications easier to understand, develop and test. Thus, costs of maintaining applications can be significantly reduced.

Challenges in 2021: Deployment of virtualization management systems, evaluating software tool replacement options, co-sourcing environment management to cut costs.

5. Higher level of Automation for Maximum Uptime

Downtime in any business has a direct impact on brand value, customer experience and revenue. This fact was accentuated during the pandemic due to more businesses starting to transact digitally and seeing a marked increase in online traffic due to storefronts being shut. The primary challenge was to provide close to 24x7x365 uptime with reduced IT Operations personnel, something made possible only thanks to automation. Enterprises adopted AIOps



Many enterprises opted for co-sourcing their environment management functions. In this arrangement, the enterprise would still have the systems in its environment, but another firm would run and operate them. These services can be far less expensive than hiring a cadre of full-time IT Operations personnel and provide the added advantage of having the right talent managing the environment with technical know-how and service guarantees in place

solutions providing proactive incident detection and autonomous resolution capabilities coupled with ITSM integrations, so the entire ticketing process was completely automated without the need for human intervention.

Challenges in 2021: Adoption of Preventive Healing Solutions for zero-downtime and complete automation, scaling businesses intelligently for increased online presence and traffic.

Ways to Achieve IT Operations Cost Reduction

In this section, we see how some of the challenges mentioned above can be tackled by enterprises in 2021, to cut back on costs and recover losses sustained due to the pandemic in 2020.

Cutting back on non-discretionary costs
Part of the non-discretionary costs in an enterprise comprise NOC personnel who are mostly contractors. They are tasked with assessing dashboards for alerts and anomalies, escalating notifications up the chain of command, and overseeing the ticketing process. Intelligent tools for notification and escalation, AI/ML solutions with which alert storms can be reduced, and automation of ticketing workflows via ITSM integrations are desirable, so the contract workforce can be reduced. Further, some amount of inhouse work can be shifted to consultants and contractors, so you only contract for the work you need, and the contractor/consultant is no longer on the payroll after the project completes.

Infrastructure as a Service for Intelligent Scaling
One area where costs can be cut for immediate returns is cloud infrastructure; cost savings in cloud services have a real and perceptible cash impact. Many organizations want to move more of their systems to virtualized ones but are still figuring out how to deal with configuration, provisioning, and overall management. If they deploy a virtualization management system to make this move a reality, they will enable faster adoption of cloud platforms.

Many enterprises opted for co-sourcing their environment management functions. In this

arrangement, the enterprise would still have the systems in its environment, but another firm would run and operate them. These services can be far less expensive than hiring a cadre of full-time IT Operations personnel and provide the added advantage of having the right talent managing the environment with technical know-how and service guarantees in place. Maintenance and management cost saving occurs in the following areas when you move in-house IT servers and services to the cloud.

Moving to the cloud reduces capital expenditures for servers and related network equipment, transforming one-time capital costs to monthly operating expenses. Labor costs are handled by the cloud provider on an as-needed basis, instead of using dedicated staff. Cloud providers can provision additional resources like disk space, CPU, memory, and communication lines faster and cheaper than an enterprise would be able to for on-premise servers and infrastructure. Cloud providers can dynamically provision temporary capacity increases as service demand peaks. With on-premise servers, SRE teams need to plan capacity to handle peak demand.

Accurate provisioning of resources in the cloud – whether it is virtual machines, cloud storage or containerized services – is important to provide intelligent scaling for the business. To this end, enterprises can incorporate a workload trend-based capacity forecasting tool that can highlight under-provisioned and overprovisioned resources by running a what-if analysis on projected transaction volumes. This can help avoid a lot of headache in terms of unexpected outages due to resource crunches and bottlenecks, or unnecessary expenditure on cloud services.

Evaluating Software Tool Licenses

In any enterprise today with a hybrid digital infrastructure in place, multiple AIOps and monitoring tools are required to collect telemetry data across silos in different environments. However, the cost of onboarding any new tool includes hidden license expenses as well, and 100% integration with the existing toolset is required to make the incorporation of the new tool as seamless, cost-effective, and

pain-free as possible. Several new vendors are offering many of the same capabilities as enterprise providers. An enterprise may be able to purchase new software and implement it for less than the cost of annual maintenance for its current vendor. In most organizations, there are software tools performing duplicate functions. This could include polling to monitor performance, faults, or other key functions. Eliminating vendors will reduce the annual maintenance bill and reduce staff time to keep the systems up and running.

Moving Toward 100% Automation

Automation replaces labor costs with software and configuration costs. DevOps enables collaboration, integration, and automation in otherwise siloed IT environments. By anticipating and reacting to failures faster and working to deliver feature improvements to end-users on a continuous basis, organizations spend less on deployment hurdles and incident resolution, and devote more resources to run the business. Dedicated automation architects can ensure that DevOps and Agile practices are implemented across the enterprise. Data Center Automation (DCA) reduces the need for manual configuration and processing. Part of this also includes autonomous resolution of incidents and automation of ticketing workflows. Dynamic workload optimization, dynamic resource provisioning and proactive identification of hotspots lead to the mitigation of issues before they even occur, thereby allowing a smaller team to run a more efficient and error-free data center.

Reducing MTTR on Unpreventable Issues

Suppression of alert storms and false positives as well as predictive alerting can be achieved through AI/ML techniques, and preventive healing can help nip issues in the bud before they occur, leading to what is nowadays being referred to as “negative MTTR”. However, some issues may still slip through the cracks due to sudden network or storage outages, hardware glitches or 3rd party dependencies like APIs and payment gateways being unavailable. In such cases, accelerated root cause analysis with event correlations and suggestions on where the error originated can significantly reduce MTTR and the number of personnel required to bring an issue to a satisfactory closure, ultimately leading to cost savings.

Such root cause analysis is aided by supplementing incident details with all time-synchronized contextual and forensic data available at hand. This could include logs, diagnostic data, business error codes, recent configuration changes, query-level statistics from the database as well as code level tracing and instrumentation. In the hands of a skilled IT Operations analyst, this data proves invaluable in establishing the chain of causation and bringing the incident to a logical end as soon as possible.

Integrations for Ticketing, Visualization, Notification and Collaboration



2020 accentuated the need for remote work and collaboration among endpoints spread across geographies. In 2021, there will be increased focus on the sustainability of a such a model, and the ability to secure and manage large populations of remote devices without overtaxing IT resources. Endpoints need to be resilient and capable of self-healing to minimize the need for IT intervention.

Other important integrations are needed to ease the life of the IT Operations personnel, who are themselves tasked with keeping remote systems for the entire enterprise up and running. The “new normal” needs to deliver more agile support for remote workers and this is putting an increasing burden on ITOps teams. Rex McMillan, Principal Product Manager at Ivanti, states that their recent research shows that for 63% of IT professionals, IT workloads have increased 37% since going remote. Hyper automation is the only way for IT to scale and handle the additional challenges. Integrating notification and ticketing platforms can save substantial time and is well worth the investment, as staff will not have to correlate information from multiple systems to determine what the true problem is and then report in another system.

Preventive Healing – A New Paradigm for IT Ops Automation

Preventive healing solutions adopt patented techniques that map an application’s workload to its underlying behavior and learn these workload-behavior correlations over a period so they can flag anomalous transaction patterns or behavioral metrics ahead of time. This helps in true predictive detection of issues before they even occur and allows for remedial steps to be put in place so the issue can be averted. Some modes of preventive healing include dynamically optimizing or shaping the workload so the underlying system behavior remains unaffected, dynamically provisioning additional resources in cloud environments so the system can handle workload surges, or projecting resource requirements based on a what-if analysis of future workload trends, so businesses can perform app-aware scaling.

Survey examines digital workspace deployment and performance monitoring

SeG INNOVATIONS has published the results from its 2021 “Digital Workspace Deployment & Performance Monitoring in the New Normal” survey. The survey report, jointly conducted with xenappblog, is a compilation of responses from 1,050 digital workspace professionals from across the world and includes learnings, analysis, and trends that will be useful for any IT professional working with digital workspace technologies.

A digital workspace is an integrated technology framework that allows employees and remote workers to access their apps and data in real time - on any device, from any location, regardless of whether the information is stored in the cloud or in a data center.

Key digital workspace technologies include Citrix Virtual Apps and Desktops, Citrix Cloud service, Microsoft Remote Desktop, VMware Horizon, VMware Horizon Cloud Service, Microsoft Windows Virtual Desktops (WVD), and AWS WorkSpaces and AppStream.

This is one of the largest surveys ever conducted in the digital workspace domain. It was conducted between December 2020 and February 2021 and specifically focused on changes that organizations had to make due to the COVID-19 pandemic.

Key Findings: Digital Workspace Deployments



- Centralized management is the primary reason why organizations are adopting digital workspaces.
- Citrix remains the dominant player in the digital workspace landscape. VMware has significant market-share.
- Cost and security concerns are two of the major reasons why organizations are not actively moving digital desktops to the cloud.
- More than 50% of digital workspace professionals are spending over a quarter of their work week troubleshooting issues.

Key Findings: How the COVID-19 Pandemic Influenced Digital Workspace Deployments

- 16% of respondents deployed digital workspaces for the first time due to the pandemic.
- 82% of respondents had changes in their digital workspaces due to the pandemic.
- 41% of respondents felt that digital workspace technologies have now become more important for their

organizations.

- Network connectivity and bandwidth limitation at the user end was the topmost challenge of organizations supporting remote workers.

“The last year has been pivotal for digital workspace technologies. As organizations suddenly had to support most employees working from home, digital workspace technologies became critical for the business. This survey sheds light on how organizations have adapted to the changing requirements and highlights process and technology improvements that organizations can make to get better return on investment in digital workspaces,” said Srinivas Ramanathan, CEO of eG Innovations.

“The report provides very valuable information for anybody working in the end user computing (EUC) space. Not surprisingly, adoption of Microsoft’s Windows Virtual Desktop offering is growing fast,” said Eric Haavarstein, founder of xenappblog.

Datadog completes Sqreen acquisition

DATADOG has confirmed the closing of the previously announced acquisition of Sqreen, a SaaS-based security platform that enables enterprises to detect, block, and respond to application-level attacks. The closing of this deal will bolster Datadog’s existing APM functionality and move the company closer to providing customers with a robust, full-stack security monitoring solution for the cloud age.

“The application layer is currently one of the most vulnerable and exploitable attack surfaces,” said Olivier Pomel, CEO, Datadog. “In combining Sqreen with Datadog, we’re closing the gap between application developers and security teams and providing our customers with robust application security, without the disjointed visibility, high implementation costs, or steep learning curve of traditional application security products.”

“This is an exciting new chapter for Sqreen that will enable us to deliver on our mission of creating a more secure future for all organizations,” said Pierre Betouin, CEO, Sqreen. “As part of Datadog, we will accelerate the way we build security tools that automate application protection, threat detection, and security monitoring while delivering more value to customers.”

Poor visibility and silos impact business outcomes

Digital.ai has released the Digital Transformation Progress Report, detailing the findings of a survey of over 600 enterprise IT, Security, and Business decision makers on the state of digital transformation in their business during the COVID-19 pandemic.

Survey respondents overwhelmingly agreed that visibility into planning processes and the alignment of software development and delivery processes to business objectives are critical to driving better business outcomes. However, more than half reported issues providing that visibility and alignment, calling in to question long term success of today's digital transformation initiatives. The research also revealed a wide range of issues relating to digital transformation, including concerns about ROI, security, and the ability to compete in today's digital environment.

The overwhelming majority (99%) of leaders agreed that organizations need visibility into business planning processes to produce better outcomes, but only 62% said they can provide that visibility and 34% admitted they are doing poorly at providing visibility. Additionally, 94% of respondents said organizations need software development and delivery to be better linked with business objectives, but only 54% said their Business, IT, and Security teams are strategically aligned and working toward the same goals and objectives. What's more, 78% of leaders said COVID-19 has further impacted the alignment between business and IT value streams in their organizations.

"Most of today's agile and DevOps tools are designed for the workgroup, making visibility and alignment fairly easy to achieve at the team level. However, scaling to the enterprise is far more complex, as organizations must break down silos and manage teams of teams that embrace different cultures, tools, and systems," said Derek Langone, Head of Strategic Transformation at Digital.ai. "We've found that in large enterprises that succeed at true digital transformation, business and technology



teams work collaboratively to define value and drive outcomes.

These companies typically organize themselves in terms of value streams, employ an open approach to sharing information, and leverage enterprise-grade solutions that offer full visibility and rich, predictive insights across the entire organization."

While 94% of respondents said they are familiar with the concept of value stream management, and 95% believe their organization is definitely or probably practicing value stream management, only about half (53%) of the organizations said their business and software value streams are very aligned, and just 60% say their organizations are very customer centric. Moreover, there was no clear consensus on what value stream management means to the enterprise.

Gartner predicts that "by 2023, 70% of organizations will use value stream management to improve flow in the DevOps pipeline, leading to faster delivery of customer value and the use of value stream delivery platforms to streamline application delivery will grow from 10% to 40%." 98% of survey respondents reported they are likely to use a platform that connects software delivery process to business objectives and provides insights needed to make decisions that increase value delivered to customers, and 63% said they would be very likely to use such a platform.

Alan Brown, Author and Professor in Digital Economy at the University of

Exeter specializing in agile approaches to business transformation, noted "The gap between investing in business transformation versus reaping the benefits has always existed but has been heightened in the last year. At the heart of it, digital or business transformation is not about doing everything faster, but about doing more of the right things efficiently. Automation and Value Stream platforms can go a long way to help organizations overcome key obstacles and enable them to achieve more of the intended results of digital transformation."

Additional Findings

Although most leaders said digital transformation was going well at their organization, the survey also found the majority have serious concerns about the state of their programs:

- 91% said they need to get more out of their digital transformation initiatives,
- 49% are not seeing the results they expected from their digital transformation initiatives,
- 56% are worried about the return on their digital transformation investment,
- 63% are worried about releasing security vulnerable products, and
- 54% believe problems with digital transformation have hurt their company's bottom line.

With regard to the COVID-19 pandemic, the majority of leaders said their organizations have responded well to the situation. However, when asked, more than half (55%) are concerned about their organization's ability to respond to change and their ability to compete in today's digital environment (54%).

Tackling cloud complexity with the right tools

Securing the right observability tools will help enterprises get control over their software and spot problems before they become a more widespread issue.

BY STELA UDOVIVIC, DIRECTOR OF PRODUCT MARKETING, VMWARE TANZU



CLOUD ADOPTION has soared in recent years with 92 percent of enterprises deploying a multi-cloud strategy, and 82 percent a hybrid cloud strategy. Change has largely been driven by businesses seeking to cut costs, deliver and scale operations faster, and offer more competitive products and services. As IT teams have accelerated their use of modern cloud applications, managing and running them has become an increasingly complex task. New research we conducted shows 86 percent of software engineers believe modern cloud applications are significantly more complex than just five years ago. Another 84 percent report using hundreds or even thousands of compute instances across a single organisation, with many containers and microservices in play.

For IT teams to keep on top of their applications, they must be able to monitor and troubleshoot any issues within this complex software landscape. However, only 16 percent of the IT practitioners we questioned are using modern observability tools. The good news is that this proportion is set to rise, with ca third (34 percent) planning to implement observability tools in the next six months.

With observability on the rise, what are the 'red flags' companies must consider for implementation?

Why monitoring tools need to evolve

Cloud complexity means application requests, in the context of distributed tracing, are now going through dozens of technologies including many third-party



APIs, cloud databases, and digital queues. Nearly half of IT practitioners state application requests touch over 25 different technologies, and at the top end, that number can reach over 100. Across this dense software application delivery, troubleshooting application issues has become even more demanding, and the risk of bottlenecks has increased.

More complex environments with distributed applications that are updated frequently, means teams must act fast when issues arise. With so many endpoints to monitor, tools must be able to scale to avoid incomplete coverage. However, many enterprises use too many tools for different elements of the cloud, creating multiple disparate systems. Half of the respondents in our survey stated they use more than five monitoring tools across their apps and infrastructure, from traditional style logging and NPM tools, to siloed container tools. And the larger the enterprise, the bigger this issue becomes. Drawbacks include a lack of or siloed visibility, which prolongs troubleshooting and makes it difficult to maintain existing tooling. More than ever, enterprises need to secure the complete picture, and better observability holds the key to managing this effectively.

Red flags to look out for

Although different business models have different demands, there are several red flags teams can look for when considering observability.

Companies with 'tech DNA' – software companies – have already adopted DevOps processes and automated deployments in their operations, so should find these red flags easy to spot. For more traditional companies moving to the cloud, it's important they build these factors into observability strategies from the get-go.

Firstly, teams must carefully assess the kinds of performance metrics they need to collect, and what they must do to measure these at high granularity. Some observability platforms collect lots of metrics and offer flexibility in fine-tuning collection, so only the details important to the business are shown. Secondly, teams should consider which observability platform will integrate well with existing tooling. Open-source monitoring is very popular and will require integration. Our research found easy integration with open-source and existing tooling is essential for half of respondents. Last but by no means least, teams should consider the cost of un-optimised infrastructures running the open-source observability platforms being maintained. It's vital to factor in such things to ensure companies aren't left with any unexpected bills to pay.

What next?

With observability on the rise, it's important to take stock of what comes next for IT practitioners and

DevOps teams. Once observability is in place, linking it to business outcomes will follow.

For organisations with operating models based around applications, tracking data insights for better decisions provides the ability to spot operating difficulties and reduce the risk of revenue loss through poor cloud service availability. Sub-optimal performance or critical delays on an application can translate into a direct loss of business – if on a Friday evening you can't order a taxi via a ride-sharing app, there's a direct revenue impact for the respective sales team involved.

But with observability, a ride-sharing service can understand its business better by observing cloud service metrics, which make clear when its service is used the most, how many customers wait for rides, and how many drivers are on the streets during peak hours. In addition to observability being important to engineering teams, a company could provide its executives with access to the tools, to better showcase how the business is performing based on both real-time and historical insights.

Another impact of greater observability uptake will be a rise in self-healing. If you instrument your observability and CI/CD properly, by the time SRE teams react, the alert is probably already resolved. Similarly, implementations of automated technologies, such as AI/ML will only increase. For example, 'surfacing' is powered by AI and helps automatically spot problems for DevOps teams to troubleshoot – particularly important when you look at highly distributed environments with millions of traces and metrics.

Finally, as with lots of software technology at the moment, we'll see an intersection between observability and DevSecOps. As developers start adding observability into applications early, they'll need to define which metrics to use, and code must be secure. The challenge for teams is to prevent vulnerabilities from being introduced, both during development and eventual deployment, and to validate that every container deployed to production is compliant and secure.

Final thoughts

It's an interesting time for software businesses looking to scale and secure their application delivery chains in increasingly complex cloud environments. The more sophisticated applications become, the greater the risk that troubleshooting will have an impact on delays for customers – and a bottom-line impact for sales. Securing the right observability tools will help enterprises get control over their software and spot problems before they become a more widespread issue. This will only improve communication across the cloud network and boost user experience at the other end of the chain for customers – a win-win scenario for everyone involved.



Telecom: Enabling automation everywhere through AI and analytics

As operators drive to increase the velocity and scale of Service Assurance processes across their complex, converging enterprise IT and network domains, automation has become a critical priority.

BY IAN RUSS, EMEA CHIEF ARCHITECT WITHIN BMC SOFTWARE'S OFFICE OF THE CTO



IN MY LAST BLOG, I discussed the need for new approaches to key Service Assurance processes. In this blog, I'll delve more deeply into how operators can leverage automation to meet key challenges, including:

- Bringing context and intelligence to trouble ticketing
- Reducing the cost and customer impact of change
- Operator-guided automation

Bringing context and intelligence to trouble ticketing

Converging trends are bringing trouble ticketing to the top of the agenda for telecommunications transformation. Rapidly changing technologies and rising complexity are increasing the potential for problems to emerge throughout the network, often with causes and impacts that are less well understood.

Automated, AI-powered analysis of complex datasets, presented through graphic interfaces, can help operations teams identify patterns in performance, alarm, and ticket data across network, infrastructure, and cloud environments

At the same time, intense competition makes customer satisfaction critical for success and growth. Meanwhile, network staff must find faster, more efficient ways to resolve problems to avoid becoming overwhelmed.

To solve these challenges, operators need to be able to use data more effectively to provide the context that drives intelligence. Automated, AI-powered analysis of complex datasets, presented through graphic interfaces, can help operations teams identify patterns in performance, alarm, and ticket data across network, infrastructure, and cloud environments. Armed with this insight, network staff can intelligently correlate multiple problems into a common root cause or set of customer impacts in order to create, assess, assign, and remediate trouble tickets more quickly and efficiently.

Reducing the cost and customer impact of change

Whether human-initiated or machine-generated, the rate and complexity of change continues to increase. Traditional approaches to change can't accommodate this at the speed and scale necessary to meet ongoing network transformation, changing customer requirements, and performance management.

Operators need to be able to improve the success rate of change, better understand its impacts on cost and quality, understand complex change interactions, and optimize scheduling to minimize customer impact. By automating change assessment and scheduling for agent-initiated change, operators can reduce human involvement while reducing cost and increasing success. As changes are planned, potential customer disruption can be predicted and mitigated more effectively, and the change itself can be automated. AI and automation can also support fully automated change assessment and scheduling for machine-generated changes to eliminate human interaction entirely.

Operator-guided automation

Operator-guided automation is an intermediate but often critical stage of Service Assurance on the journey to fully autonomous networks and self-governing systems. Before operators turn over control entirely to AI and closed-loop automation, they can first choose to take a semi-automated approach. Following an assessment in which the system determines what level of human interaction is needed, AI-driven intelligent automation can guide operators

to the right trouble remediation actions or enable fully automated remediation. In either case, the success of actions is measured to improve accuracy. As manual intervention is reduced and eliminated, the operator can process toward fully closed-loop automation.

Evolving to AIOps

To deliver a competitive customer experience, operators must enable a shift of Service Assurance from fixing problems that have already emerged, to using system data to enable a predictive approach in which problems are averted before they compromise service quality, and AI makes it possible to assess common underlying problems across systems. More broadly, operations teams can use big data, machine learning, and analytics to identify patterns in monitoring, capacity, and automation data across complex technology infrastructure. Based on this insight, they can work more quickly and effectively to improve the speed, quality, and cost-efficiency of service delivery. Ultimately, the goal is to enable fully autonomous delivery in which real-time automation ensures optimal service quality with the speed, scale, and efficiency required for modern operator environments.

In my next blog, I'll discuss using AIOps to deliver an improved customer experience, including improving operator visibility across all customer touchpoints, automating problem resolution, and improving service quality management.



The importance of AIOps within value stream management

IT organisations today are charged with both running the business as well as reinventing it.

BY GAURAV REWARI, CTO AND GM OF AI AND VSM AT DIGITAL.AI



THE FORMER is usually managed by the infrastructure and operations (I&O) teams, whereas the latter is done by the application development and delivery (AD&D) function. As 70 percent of IT budget goes towards keeping the lights on (KTLO), artificial intelligence for IT operations (AIOps) began as a way to drive cost reduction through greater automation of operations and infrastructure activities, shifting the burden of KTLO from humans to machines – the ultimate goal for AIOps.

However, for modern day organisations, applying AIOps for I&O teams alone is not enough. For organisations to reap all the benefits of AIOps, it needs to be part of the application development organisation's strategy to reinvent the business. Enter Value Stream Management, or VSM – an approach which encompasses every step in the software delivery process and which is centred around helping companies move faster with higher quality, whilst

helping them to align technology activities with business outcomes.

Incorporating AIOps into VSM provides all streams of the business with data-driven insights that help organisations to simultaneously manage risk and deliver high quality digital services more efficiently. The development of and key roles of AIOps Recently AIOps has been applied to many use cases. The first wave involved the use of AIOps for event noise suppression by filtering out unnecessary alerts generated by application, infrastructure, and network monitoring tools. AIOps techniques then extended to automating the process of understanding the root cause of issues to enable the swifter remediation of major incidents thus helping to reduce the downtime of critical business services.

AIOps has also found applicability in the area of service management. AIOps techniques have led to the development of tools such as chatbots to help provide an answer to an employee's query by leveraging the company's knowledge base system as well as historical incident and service request patterns. Another compelling AIOps use case in this area is the application of machine learning techniques to predict whether an incident that seems relatively benign right now has the propensity to become a major incident based on patterns detected from historical IT service management (ITSM) and monitoring data. Additionally, AIOps and process mining techniques can help uncover bottlenecks in service delivery such as recurring patterns of ticket reassignments, clusters of incidents that may have similar underlying characteristics even though they are not tagged as such, etc. These insights can help drive process optimisation and automation decisions improving the cost, quality and effectiveness of service delivery.

As AIOps matures as a field, we are seeing more extensive capabilities such as change risk prediction and release schedule risk prediction come to the fore,



highlighting the growing applicability of these techniques to the world of DevOps. With the broadening of the relevance of AIOps beyond the initial shores of the IT Ops organisation, it also has a key role to play in automated governance across the IT estate, facilitating the transition from a realm of key performance indicators to one of machine generated key risk indicators. The benefits of such a transition from lagging to leading indicators are extensive and range from lowering costs of IT operations and providing a higher quality of service, to maximising uptime of services and ultimately lowering risks for the business.

So how does AIOps benefit Value Stream Management?

There is currently a mounting dissatisfaction with the ROI being realised by digital transformation initiatives. In fact, a recent survey of 600 Business, IT and Security decision makers found that the majority of enterprise leaders are concerned about just that. Businesses are learning that practices such as Agile and DevOps are necessary, but not sufficient. To reap the full benefits of digital transformation implementations and optimise business outcomes, organisations need to address the challenge of how best to tie the work of the development team with the needs of the business.

Additionally, as incumbents embark on their digital transformation journeys, they must be able to release new capabilities rapidly without impacting the availability and performance of the production systems that power the business. How to move faster with quality is another key challenge.

It's no small task but, when implemented properly, VSM can help address these two challenges and be the missing piece required for digital transformation success. The required capabilities in this area have been well described in recent research by Gartner and Forrester. Their research shows that one fundamental impediment is that most enterprises use multiple point products in their development and delivery toolchains. This siloed landscape limits the ability of organisations to harness data across development, operations, and the business to support the goal of moving faster with quality. Gartner states that the solution for this is a VSM platform (VSMP) "that enables business and IT to align their goals," and helps "organisations to use leading indicators and predictive analytics to gain a competitive advantage and quickly adjust course to drive customer-centric outcomes."

As one may have already surmised from this description, AIOps techniques will have an increasingly important role to play in VSM. As an example, one of the required capabilities for a VSMP highlighted by Gartner is change risk analytics, a capability also described in their AIOps market guides. A change can be as small as a bug fix or a modification of a configuration, or as large as an entire

release. Change is where the work of development meets the world of operations. It is also the largest cause of outages and business service interruptions. Change risk prediction uses AI techniques to analyse historical change data from across the software development and delivery lifecycle to understand the drivers of failed changes, quantify the impact they have on business services, and score the risk of upcoming changes failing.

This allows an organisation to automate the process of pushing low risk changes into production and also hit the pause button to further scrutinise a proposed change deemed risky by the ML model before it goes into production. The ability to do this is at the very heart of being able to move faster with quality.

Similarly, there are other AIOps use cases such as release schedule and quality risk prediction, as well as automated governance across development and operations that are fast emerging as foundational capabilities in the VSM approaches of companies seeking to optimise and transform their businesses through digital transformation.

AI and humans working in harmony

AI today has great applicability in both decision automation and decision augmentation. As consumers we experience decision automation via new tools like automated appointment schedulers, whereas decision augmentation capabilities are regularly woven into our online shopping experiences such as relevant recommendations for books to read or items to purchase based on shopping history and inferred preferences.

The same applies to IT organisations though, as we saw in the AIOps for VSM examples just discussed, the focus is largely on decision augmentation. This is inherently a coexistence model where the goal of the AI capabilities is providing machine generated insights into monitored patterns and delivering early warning of risks to arm individuals in specific roles with the proactive insights they need to fulfil their responsibilities efficiently and effectively.

Looking to the future

With businesses expanding and systems evolving every day, the challenge of delivering higher value at lower costs gets tougher. AIOps, as a way of shifting some of the growing burden from humans to machines, can help ensure business continuity and drive efficiencies across the different departments. Companies that choose not to consider the AIOps-powered VSM approach are missing a significant opportunity.

VSM is all about smart decision making, being able to predict risk and avoid a drop in operating efficiency. As the world continues to invest in digitisation, businesses that embrace all AIOps has to offer will have a lasting competitive advantage.



The future of service management in the DevOps era



An open, collaborative platform enables teams to scale operations quickly and ensure the organisations' critical services are always on and operating at high velocity.

BY GARY BLOWER, SOLUTIONS ARCHITECT, CLEARVISION

WHETHER YOU VIEW your organisation as having an agile approach or not, in 2020, companies had no choice but to drastically change their way of working as the world rapidly pivoted to remote working.

Organisations that had already embraced agile principles had the advantage of being able to adapt faster to the pandemic and meet the demands of their employees, who were suddenly all working from home. Now, as we start to slowly emerge from multiple lockdowns and restrictions, one interesting side effect of COVID-19 is that it has lowered our collective tolerance for slow, overly bureaucratic processes. We all crave an agile approach, whatever our definition of agile might be.

COVID-19 has accelerated digital transformation. Digital innovation has fundamentally changed how the world operates. COVID-19 demonstrated just how much we rely on technology. And, as modern technology permeates every area of our lives, our expectations around the availability of information and the speed with which we can obtain it are even higher than they were pre-pandemic. Therefore, as lockdowns ease, the world is continuing to change just as rapidly to keep pace with the demands on

businesses, who must accelerate out of recession and aggressively compete to remain relevant. The knock-on impact of this acceleration is that organisations need their IT teams working together as efficiently and effectively as possible. Likewise, their IT service management (ITSM) capabilities must be nimble and efficient to support shifting organisational priorities, capitalise on new opportunities, and satisfy growing end-user demands for immediate and seamless service, wherever users are located.

To meet this increasing demand and requirement for speed, the flow of work between the support, DevOps and operational teams must be unified, and teams need to be empowered to deliver work with agility. IT teams are under huge pressure and are required to become even more adaptable to the challenges they face. This means that practices and workflows need to remain flexible so that teams are better positioned should situations like we just experienced in the past 12 months arise again in the future.

Traditional service management approaches can't keep pace with demand

However, even the smallest request for change is not an easy task for some organisations and must be approved by layers of bureaucracy, which can take weeks or sometimes months. Additionally, this increased demand, combined with the ongoing pressure to lower costs, runs counter to traditional approaches to service management that emphasise risk mitigation and control over efficiency and agility – leaving some IT teams hamstrung and unable to play to their full potential. In our 'always on', digital world, this will disadvantage those companies unable to respond, with end-users and customers no longer willing to accept long wait times. And why should they? The COVID-19 experience showed that, when we really need to, we can completely change our way of working overnight. Therefore, many customers are now unforgiving of those that cannot accommodate their requirements or promptly meet their expectations.

One way that organisations can accelerate their service management initiatives and introduce more efficient methods to serve ever-growing business demands is by implementing Jira Service Management. This is the only ITSM solution built on the Jira software development platform. This means that users don't have to seek the Jira application separately, and they benefit from having everything they need in one platform.

DevOps, IT support, and IT operations must all collaborate

This accessibility is important because IT teams using other service management tools often end up integrating their application with Jira for additional functionality, which can be clunky and not as streamlined. The co-existence of Jira Service

To meet this increasing demand and requirement for speed, the flow of work between the support, DevOps and operational teams must be unified, and teams need to be empowered to deliver work with agility

Management and the Jira software development platform has huge benefits because it means that support and development teams can collaborate on the same platform and fix software issues and incidents faster. Jira Service Management was also designed with both IT and development teams in mind and provides streamlined requests and change management processes. This allows teams to make change requests without complex approvals and link incidents to problems in one click.

With other service management platforms, siloed tools between development and IT operations can result in context switching, lack of visibility, and decelerated work. As a result, integrations between Jira Software and service management tools tend to be weaker and cumbersome to manage. In contrast, tight integrations between Jira Software and Jira Service Management mean seamless and accelerated workflows between development and IT. Teams can link issues across Jira and ingest data from other software development tools, providing IT support and operations teams with richer contextual information to respond rapidly to requests, incidents, and changes.

Jira Service Management also offers customisable templates for ITSM, customer service, and business teams such as HR and finance. Furthermore, an intuitive portal in Jira Service Management makes it effortless for customers to ask for help, while the simple UI makes it easy for teams to use. And, with easily configured automations, IT teams can prioritise and resolve requests quickly.

Service management built for the DevOps era

In today's world of digitised services and support, being able to deliver a rich and collaborative service desk, modern incident management, and change management is critically important. The world is changing fast and, to keep pace, organisations need a service management platform built for the DevOps era. An open, collaborative platform enables teams to scale operations quickly and ensure the organisations' critical services are always on and operating at high velocity. This will ensure they can respond quickly to business change while delivering great customer and employee service experiences.



Weeding your digital garden to provide outstanding software experiences

Providing a digital experience is like tending to a garden – there's always room to grow and improve.

BY DAVE ANDERSON, DIGITAL EVANGELIST AT DYNATRACE



Over the past few months, it's safe to say many of us have spent a lot more time in our gardens. Being face-to-face with rosebushes, lawns, and trees every day has encouraged us to spend more time cultivating and pruning our gardens to keep them looking tidy. Another of our behaviour changes during and post-lockdown is engaging more than ever before with

digital services – from socialising online, click and collect groceries and order-at-table service in pubs.

Just as we want our gardens to look their best when we're spending more time there, we've also begun to expect more from our digital shopping, working and socialising experiences.

However, the difficulty for organisations trying to provide the perfect digital experience for customers and employees is that, like with a garden, there's always going to be 'weeds' that need removing to keep it pristine and avoid performance problems creeping in. Take for example, cloud complexity.

The average mobile or web application transaction now crosses 37 different technology systems or components, and the dynamic nature of cloud-native architectures like microservices and containers also means that things change very quickly.

As a result, the cloud environment resembles a big garden with lots of plants in it, which can fast become overgrown by redundant infrastructure, excess workloads and performance issues if not carefully tended to by organisations. It's almost impossible to completely stop performance problems from arising, so the goal of IT teams is to stop them before they ever affect the user experience.

Identify the 'root' of the problem that needs weeding

The modern cloud environment provides flexibility and scalability benefits that enable organisations to focus on developing new services, creating business value, and keeping customers happy. In fact, 93 percent of companies have said that successful cloud initiatives have led to improved customer experience. However, whilst the cloud brings huge advantages, the modern development models that it enables significantly increases speed and scale within digital ecosystems, which brings a lot of complexity.

With containerisation, for example, new infrastructure is spinning up every few seconds in many cases, making it difficult to keep on top of and prevent the 'garden' being overrun. Database queries, JavaScript errors and service requests can all contribute to millisecond slowdowns that, when scaled, can cause front-end performance slowdowns.

IT teams can end up digging around dashboards and cloud ecosystem data trying to understand the root cause of performance issues. Not only does this drain time and resource, it also affects the user experience, as IT teams are busy weeding out problems from the digital garden, with less time to develop new functionality.

What IT teams need amidst all this complexity is precise answers on issues that need 'pruning'. Organisations can harness AI-assistance to help them here, pinpointing the root cause so teams can continuously optimise services and respond to problems before they affect the user experience. This way, armed with actionable insights, a fix can be deployed instantly, as teams do not have to spend time hunting around for the one 'weed' that is impacting user-experiences and degrading the performance of the entire IT ecosystem.

Offering re-leaf for IT teams

Once IT teams have found the cause of the problem, the next stage is 'pruning' or solving it. However, manual pruning can become a huge time drain – frustrating given the task is often relatively straightforward and repetitive, and there is no shortage of 'weeds', making it harder to prioritise the smaller ones.

As such, IT departments need to start looking at where automation can help speed up this process. Organisations can create a culture of optimisation by automating continuous delivery and operational processes. For example, putting in place automated runbooks for spinning up cloud capacity if certain usage patterns occur or for decommissioning redundant cloud instances. This will free up valuable time for teams to spend on digital transformation projects and the delivery of new functionality for the business.

As close to perfection as possible

Providing a digital experience is like tending to a garden – there's always room to grow and improve. Organisations now face more pressure than ever to deliver new digital services quicker, meaning weeding out problems that slow them down is crucial. By harnessing AI-assistance and automation, IT teams can build a new, more efficient approach to operations based on a culture of continuous optimisation.

This will ensure IT teams don't spend precious time 'keeping the lights on' and can instead focus their energy where it's needed most – delivering innovation to provide a digital garden that employees and customers enjoy spending time in, and that competitors strive to plant for themselves.



It's now or never for more perfect software

Now is the time to celebrate how software is helping us, everyday consumers, to get our daily tasks done and make our lives easier. For instance, from paying bills via an app or streaming a film, most of our day-to-day activities are digital and our reliance on digital infrastructure has never been greater than during COVID-19.

BY GREG OUIILLON, CTO AT NEW RELIC



Prioritising software

WITH OUR DAILY CONSUMPTION being largely digital, companies need to prioritise software, or they will risk losing customers and revenue. There is currently a real struggle to retain talent in the tech industry, which shows no signs of abating, and the global shortage of skilled software developers is at the heart of it. Recent research commissioned by New Relic found that 37% of UK respondents had difficulty recruiting the right development and infrastructure skills, making it a major barrier to software and systems performance within their organisation. Interestingly, the figures showed that Australian businesses (50%) were struggling more than those in Japan (43%), UK (37%), US (30%), Germany (30%) and France (29%).

In addition, existing engineering teams are facing increased difficulties in managing new software and infrastructure, all while working from home. That said, nine out of ten of the best performing businesses surveyed said that developing and deploying more perfect software has helped them to overcome

challenges during the pandemic. The top 25% of respondents are identified as 'More Perfect Software Leaders,' having reported higher success against business metrics like financial performance and growth, pace of innovation, employee engagement, and brand perception. For these leaders, it is a strategic priority for them to:

- **Leverage cloud technology** – By migrating applications and infrastructure onto the cloud, it allows businesses to enhance agility and speed to market.
- **Observe** – A single source of truth on how software and systems are performing in real time provides holistic insights that can shape software design and development, thus allowing teams to rapidly identify and resolve issues.
- **Monitor data intelligence** – Evaluating software performance data and incorporating it into customer experience with business data can enable companies to develop solutions that answer customers' needs.
- **Build resilience** – Businesses should consider employing chaos testing and automated remediation to minimise errors and downtime while still enabling faster innovation.

The research also found that 77% of leaders adopted chaos testing, compared with just 30% of laggards, and 89% of leaders engage in automated remediation in contrast to 5% of laggards. With these systems in place, 83% of leaders experienced only five software and systems outages per month, while 49% of laggards faced 11 to 15.

Meeting customers' needs

Today, the rise and fall of a company is all down to how well they engage with their customers, no matter which industry the business falls into. So, how do companies meet their customers' demands? Most businesses today are using software to service





their customers and manage their supply chain from a competitive standpoint. With the pandemic speeding up digital transformation, digital channels could soon become the only way customers interact with many of their go-to brands. As organisations undergo this change, they will need to measure themselves against speed, service, and scale to stay ahead of the competition.

Key questions businesses ought to ask themselves include:

- How quickly can new products, services, and improved customer experience go to market?
- Is our business able to service customers 24/7, regardless of their geographical location?
- How does our business scale to meet increasing global demand?

Maintaining a focus on software development and ensuring engineering teams are fully equipped to move fast, fail quickly, and scale easily is vital. Businesses would struggle to achieve this agility if they were stuck in data centres with legacy technology, silos of data, and lack of observability - software is what provides businesses with a competitive edge against competitors.

Empowering teams

Developers need to have good access to authority, data, and tools to make faster and more informed decisions to provide seamless user experiences and enable business continuity during the pandemic. Although engineering teams are facing a huge amount of pressure, having a desirable DevOps culture in

place encourages speed and drives efficiency, optimal productivity, and high job satisfaction.

However, many companies are finding it challenging to expand their software development teams to meet demand in the first place due to skills shortages and poor employee retention. This is as a result of numerous issues including a lack of a 'developer-first' culture within the business, low levels of autonomy, and limited access to the necessary tools and resources. The research provided further insight – 99% of leaders say their company culture and observability technology allows their developers to make quick decisions without fear of repercussions, whereas only 33% of laggards can say the same. Therefore, in order to keep and attract developers onboard, firms need to be able to present exciting and challenging projects that allow them to grow their skillsets. At the same time, they need to be given the best tools and technology to successfully complete the tasks on hand to realise greater job satisfaction, increased tenure, and a highly skilled and motivated workforce.

Now, more than ever, it is essential that businesses are able to deliver bespoke and frictionless experiences to their customers as COVID-19 continues to expedite digital adoption. And it is our developers and engineering teams that are keeping digital infrastructure afloat during these times of unprecedented demand for online services. To realise more perfect software, it is time to re-think processes, put employees first, and ensure they have the tools they need to succeed. It is now or never.

CloudFabrix launches AIOps 2.0

CLOUDFABRIX made major Observability & AIOps 2.0 announcements at its 2021 AIOps Virtual Conference. AIOps 2.0 accelerates AIOps adoption by filling IT visibility gaps, simplifying & automating data preparation and integrations, eliminating re-tooling and by driving business/IT outcomes. Additionally, AIOps evaluations and proof of concept (POC) projects can be implemented in a matter of just days & weeks with the company's newly announced POC Express offer.

Traditional AIOps implementations were expensive and bogged down by lengthy implementations, complex data integrations, extensive retooling efforts and manual instrumentations, while solving only limited use cases. With AIOps 2.0, CloudFabrix is advancing AIOps to overcome these challenges so customers can easily embrace AIOps, achieve faster ROI and realize its true potential and benefits.

"As detailed in our DEJ research report (19 Key Areas Shaping IT Performance Markets in 2020), we see almost a 3x increase in the number of top performing organizations that have AIOps as their budget item," says Bojan Simic, founder DEJ, and also a keynote speaker at the conference. "Enterprises have a lot of hope and expectation from AIOps and we are at a critical juncture where AIOps projects need to be made easy to evaluate & implement. CloudFabrix AIOps 2.0 innovations and announcements are a huge leap in this direction."

Key New Product Announcements:

Observability-in-a-Box & Edge AI:

Enterprises are still struggling with IT visibility gaps in their dynamic and ever-expanding multi-cloud hybrid environments. Success of AIOps projects hinges on the quality and completeness of data that is available for AIOps to analyze. CloudFabrix Observability-in-a-Box was squarely designed to fill these gaps, by providing a unified solution to continuously monitor, collect and send metrics, logs & traces, as a unified dataset to its AIOps solution. Additionally, the solution provides Artificial Intelligence (AI) inferencing at the edge, enabling



dynamic alert thresholds, predictive alerting, anomaly detection, baselining and more. The solution has a minimal footprint and can be run in just one VM using a Docker-based microservices environment.

"By 2025, 50% of new cloud-native application monitoring will use open-source instrumentation instead of vendor-specific agents for improved interoperability up from 5% in 2019," says Gartner. Observability-in-a-Box leverages open-source technologies and open telemetry instrumentation that provide great flexibility and investment protection to switch to any analytics software in the future. The solution can be deployed in a branch, edge, data center or in the cloud within minutes.

AIOps Radar View - Bird's eye view of real-time Application or Outcome performance:

This new feature provides two key capabilities

- 1) Accelerates Root Cause Analysis (RCA): by collecting all contextual data and analyzing it for anomalies and correlation to expedite root cause analysis. It then surfaces the root cause, noticeable data points and important changes that point to a root cause to help narrow the scope of investigation.
- 2) Track performance of Outcomes: Business, service and IT operational outcomes can be tracked by continuously collecting service metrics with full-stack data and correlating it with the health and performance of the underlying technology stack.

CFXDX - Modern Data Preparation & Integration Automation:

IT operations data is generated by many structured and unstructured sources and tools, spread across different domains and environments.

This poses a continuous challenge for AIOps implementations to efficiently acquire and ingest data while also integrating with complementary tools within the IT ecosystem. With CFXDX, AIOps implementation times can be cut in half, enabled by rapid development of new data integrations, reusable and modular data flow pipelines and interoperability with the rest of the IT ecosystem (i.e., Data Lake/Ticketing/BI Tools/RPA). CFXDX accelerated data integration implementation in many advanced use cases like e-bonding, and CMDB synchronization and incident NLP classification using OpenAI/GPT-3.

AIOps POC Express:

A key pain point of AIOps has been lengthy proof of concept (POC) and pilot projects. With AIOps POC Express, POC projects can be completed in a matter of days to weeks, instead of several months. Results can be presented based on historical and offline data, without having to perform tool deployments and live integrations with production tools.

Customers can provide historical incident data and alert data (optional) and identify primary sources of alert noise generation, identify blind spots and establish potential ticket volume reduction that can be achieved with AIOps.

Monitoring and optimising Google Cloud infrastructures

CIRCONUS and Cloudbakers have formed an exclusive partnership in which Cloudbakers will serve as the Google Cloud Platform reseller for Circonus, and Circonus will serve as the exclusive provider of GCP monitoring and analytics to Cloudbakers' customers. Cloudbakers' GCP expertise coupled with Circonus' powerful monitoring capabilities empower companies to more effectively optimize the health and performance of their entire GCP infrastructures.

Cloudbakers, a recent winner of the Google Cloud Expansion Partner of the Year Award for North America, has experienced strong organic growth over the past decade while transforming hundreds of organizations with cloud-based data, application, and infrastructure modernization offerings. Cloudbakers' blend of technical skills and people skills helps clients adapt to the ever changing digital landscape, modernize their applications, and do so in a practical, cost effective manner. Cloudbakers "brings the cloud down to earth" for their clients.

Circonus' advanced analytics and monitoring platform enables real-time visibility of the behavior, health, trends, and performance of all applications, infrastructure, and containers in one unified platform. Built on a performant and proven time series database, the platform ingests, analyzes, and visualizes massive amounts of time-series data and provides unlimited scale, retention, storage.

Circonus' GCP monitoring solution includes a lightweight cloud agent to easily collect metrics from GCP, as well as in-application dashboards. Regardless of size, Circonus gives companies visibility of their entire GCP environment in real time. To quickly identify and resolve issues, customers can view anomalies; find disruptions, bottlenecks, and latencies; and recognize performance patterns.

"Organizations trust and rely on Cloudbakers to help them more easily achieve their most cost-effective and best-performing GCP environment.

Successfully realizing this requires a combination of our GCP expertise along with the best technology solutions, and that's why we're excited to partner with Circonus," said Alan Miller, VP Google Cloud Sales at Cloudbakers. "We realized that Circonus' monitoring and analytics platform would give our customers the power, visibility, and flexibility they need to ensure the health and performance of their entire GCP infrastructures."

"We're seeing a significant acceleration of organizations moving their on-premises IT infrastructure to the public cloud," said Jason Bobb, VP Sales at Circonus. "But the cloud is complex, and most organizations can't realize its full benefits without expert help. With a decade of experience and fluency in the Google ecosystem, Cloudbakers is extremely well positioned to help organizations successfully navigate their GCP journeys. It was clear to us that their deep GCP knowledge and focus on customer satisfaction make them a valuable partner to Circonus."

TET Partners with Centreon

TET, a leading value-adding reseller (VAR), and Centreon are joining forces to help digital businesses build the scalable IT strategies that will drive success in the uncertain post-Covid reality. The partners see ITOps monitoring as an integral and strategic element of future-ready IT solutions. A vision shared by a majority (89%) of organisations that say IT monitoring is a top or high priority to their business's IT strategy.

Founded in 1985, TET's focus is on listening to clients' needs and recommending the best IT strategies to help them create value and face change—most recently in the context of a pandemic that accelerated enterprise innovation. "Businesses are aware that COVID-19 recovery will also be driven by employee and customer expectations," explains Paul Knight, Account Director, TET. "But they still don't know where they'll land in this new phase of the digital journey."

The flexibility to deploy just-in-time infrastructure

A recent survey reveals workers want to decide whether they'll be working from home or the office, and they're evenly split on the matter⁽²⁾, while consumers are split between new and traditional shopping habits.⁽³⁾ Marc-Antoine Hostier, Chief Revenue Officer at Centreon says "the lesson businesses have learned from the pandemic is to be ready to deploy

anything, anywhere, fast. This entails relying on hybrid cloud, and generally speaking, growing a more diverse, spread-out IT estate. In that context, ITOps monitoring becomes a strategic part of TET's services and projects."

The VAR recommends an IT monitoring platform at the outset of any innovation project. "This explains our decision to partner with Centreon—they can monitor anything," says Paul. "We've started using Centreon for large public sector and enterprise clients and one of the key advantages of such a scalable, business aware monitoring platform is that it gives us and our clients more control from initial project design to continuous feedback on the performance of the solutions that are implemented."

Supporting a mix of rapidly changing digital strategies "Through an interconnected, cloud-to-edge approach, Centreon can easily support the transformative journey of businesses," says Juan Lyall, Centreon's Country Manager in the UK and Ireland. This year, the average percentage of IT infrastructure in the public cloud is expected to rise to 26%, private cloud reach 29% and the edge/IoT it will rise to 21%. (1) "One thing is certain—businesses, no matter their size are looking to keep their options open as they build their post-COVID strategy, and we'll be there to support them." concludes Paul.

Why intelligent automation is an absolute must



At a time when competitive pressure has never been more intense, good customer service has never been more important.

BY KAY KNOCHE, DACH LEAD FOR DECISIONING & AI SOLUTIONS AT PEGASYSTEMS

AS THE DEMANDS of the digital customer continue to increase, a poor service experience with insufficient advice or long waiting times will no longer be accepted by savvy customers who realise that they can always go elsewhere. So what's the solution to achieving this customer service success? For many, intelligent customer service is the answer.

A look at history shows that customer service has continually changed over the years. In the past, the 'intelligent' part of customer service depended on the intelligence of the call centre agent. The best agents learned the quickest ways to solve customer problems, understood how to navigate systems, memorised the steps required to resolve common requests, and established best practices.

The quality of customer service thus depended on the competence of the respective call centre employee. In the next phase of intelligent customer service, thinking evolved into trying to understand customers and their needs. As a result, many companies tried to get a 360-degree view of their customers. And although this well-jointed effort was based on a solid idea, building a centralised database ultimately led to a flood of information. As a result, many organisations found that giving agents too much information can be as detrimental as giving them too little.

Around a decade ago, another challenge was added when customer service was extended beyond the traditional means of the telephone to new channels. With the introduction of these channels such as mobile apps, chatbots and social media, it became more difficult to ensure consistent interactions.

The task for companies is to give customers the impression that they are talking to one single, identical 'contact person', even when they switch from the mobile app to the phone or from an agent with email communication to another agent in chat.

As time has progressed, these challenges become even more complex, as communication channels developed further and tools like self-service portals became prevalent. The only way to successfully meet these challenges is to enrich your customer service with intelligence. This means being able to offer a contextual, consistent experience for customers. For this, two points are particularly important: the acquisition of customer insights and the consistent quality across all channels.

The customer insights aspect means that a company really knows and understands its customer, including their contact history with the company, the context of their request and current mood. If a customer is already upset for some reason, organisations need to go to lengths to ensure that their frustration is not exacerbated, which means applying empathy. If, on the other hand, a customer is in a good or neutral mood, it may be a good time to make a cross-selling or up-selling offer.

The quality across all channels refers to the ability to always provide the right answer, regardless of the communication channel used. There must be no differences in the quality of the advice, the speed, or the solution.

The quick resolution of customer challenges and challenges always requires a number of decisions to

The task for companies is to give customers the impression that they are talking to one single, identical 'contact person', even when they switch from the mobile app to the phone or from an agent with email communication to another agent in chat

achieve the right result. Companies need to be able to rely on data quality and availability. Given the amount of customer data generated across different channels and interactions, an artificial intelligence (AI) solution is needed that analyses the data in real time and recommends concrete measures that can inform the next best available action.

Implementing an integrated solution for intelligent customer service brings decisive benefits to a company, as the example of UPC, one of the largest cable network operators in Switzerland, shows. According to Urs Reinhard, Chief Digital and Customer Officer, UPC, the company's agents find all relevant customer information on the first screen of a modern customer service application.



Without any clicks, they need only five seconds to understand the customer. As a result, UPC has reduced the average processing time of each customer by more than 30 seconds.

The customer, for their part, gets a faster and better service.

An intelligent service is the key to meeting or even exceeding customer expectations. It helps

companies create smooth customer experiences and find relevant solutions in near real-time, regardless of the communication channel, be it a self-service portal, a chatbot, or an agent-based channel.

Last but not least, an optimised, intelligent customer service also leads to higher productivity of call centre agents.



Accelerate your automation journey with a Centre of Excellence

An automation CoE is the essential foundation needed to successfully launch automation initiatives, and, in doing so, becoming a business that thrives within the ever-changing and challenging business environment.


BY VIJAY KURKAL, CEO, RESOLVE



AUTOMATION has been at the heart of technological progress as far back as the first industrial revolution. In 1733, John Kay's revolutionary invention of the flying shuttle increased weaving speeds and changed manufacturing forever. The fourth industrial revolution, also known as Industry 4.0, began in 2010 and ushered in an era focused on digitisation and virtualisation fuelled by IT automation. In fact, digital transformation and IT automation are inextricably linked, and success with both requires organisations to embrace a digital-first mindset.

Businesses have turned to digital transformation over the last few decades to cope with a variety of challenges. The last twelve months in particular have seen an incredible acceleration in these initiatives in response to the challenges posed by the pandemic. IT teams faced a new set of demands as workforces became remote and digital channels became the lifeblood of our personal and professional lives. Digital transformation has also been driven by astronomical growth in data volumes. According to a recent IDC report, data volumes increased from 2 zettabytes in 2010 to 59 zettabytes in 2020, with claims that this figure will reach 175 zettabytes by 2025. More data

means more insights, better predictive capabilities and further business optimisation, yet these volumes of data are overwhelming and require new solutions, like automation and artificial intelligence (AI), to make sense of all that data. Businesses that invest in automation and AI will reap the rewards with rich analytics that drive agility and innovation, enabling them to not only endure constant global changes, but to thrive in our current environment. Not surprisingly, strategic leaders are looking for ways to quickly advance their organisations' automation maturity. It is no longer enough to simply automate repetitive tasks. Companies must be innovative in what and how they choose to automate. As organisations move from automating simple processes and achieving quick wins, to hyperautomation scenarios that can handle more complex use cases and advanced integrations, part of the challenge is properly assessing the full scope of automation's capacity. With so many promising



opportunities for optimisation, it is difficult to know where to begin. Introducing an automation Centre of Excellence (CoE) is not only an essential part of the journey, but the very foundation from which to start building.

What is an automation Centre of Excellence?

A Centre of Excellence brings together a cross-functional group of experts who are dedicated to

deploying, scaling, and leveraging a specific technology (or group of technologies). Their sole focus is to improve business outcomes by leveraging that technology to its maximum potential, and ensuring that it is being utilised in an innovative and successful way, which also complies with corporate governance and compliance standards. To form an automation Centre of Excellence is to guarantee this technology is tested, analysed and brought into operation effectively and efficiently.



In practice, a CoE establishes the overarching automation strategy and framework, develops the operating model, and ensures the appropriate skills are in place to support the strategy. The very existence of the CoE also signifies an organisation's dedication to automation and can foster the culture required for automation to succeed.

Scaling and optimising automation with metrics

Measuring, sharing, and communicating the value of automation is critical for the long-term success of your automation initiatives and is a core responsibility of the Centre of Excellence. Illustrating automation ROI secures the cross-functional support required to scale, from the C-suite members who control the purse strings to the practitioners who are in the automation trenches.

Adopting a data-driven approach also enables a CoE to identify where automation can be leveraged to provide the best business benefits, including areas for quick wins. Tracking and reporting on key metrics, such as hours saved, cost savings, and improvements in service delivery, can be done at a granular level for each automated process. These metrics offer visibility into which automation candidates should be prioritised next for the greatest gains. And, of course, aggregating these metrics across the entire automation ecosystem measures the overall impact of automation on business outcomes.

There are many additional KPIs that can be tracked as indicators of automation success, going well beyond costs and hours. Examples include mean time to repair (MTTR) problems, call volumes, the number of

incidents or requests handled solely by automation, or the time required to complete service requests. Customer satisfaction and employee happiness are also important metrics for many organisations. An automation Centre of Excellence should develop a committee to determine which KPIs are most important based on business goals and dynamics, and then develop a baseline for each one in order to show demonstrable improvements post-automation.

Cultivating collaboration and culture

Fostering a culture of automation is critical for automation initiatives to succeed, and every CoE member should be prepared to take an active role in driving the cultural shift. The CoE bears responsibility for creating an environment where new ideas can flourish and that prioritises innovation, cultivates collaboration, inspires participation, and offers employees new opportunities.

Within a CoE, collaboration and unity are key tenets. Siloed thinking and competitive departmental rivalries must be avoided at all costs; working together towards a positive result should be the central focus. From here, it is possible to achieve a company-wide sense of collaborative culture and ensure there are enthusiastic automation ambassadors within every department.

An automation CoE is the essential foundation needed to successfully launch automation initiatives, and, in doing so, becoming a business that thrives within the ever-changing and challenging business environment. Gartner predicts that 69% of routine work currently completed by managers will be fully automated by 2024, making it clear that automation is here to stay and will continue to play an integral role in digital transformation. Do not get left behind.



WHAT IF WE looked at legacy infrastructure in IT through the same lens? Too often, tech teams rush to discard mainframes and legacy infrastructure, immediately aiming to build something shiny and new. However, just as the Saqqara was revolutionary at the time of the built, and acted as the origin point of Egypt's 96 other pyramids, mainframes that we now classify as 'legacy' were the epitome of innovation when first set up.

Like the thousands of blocks that built the pyramids in Saqqara, mainframe applications have thousands, even millions of lines of code. Each serves a very specific purpose: a purpose that has probably outlasted the developer who first wrote the code. Only once this purpose is examined and understood, can enterprises move onto successful app modernisation.

Learning from a Pharaoh's thought process

Organisations have been rushing to modernise their IT infrastructure in the last decade, and the race to digitally transform has picked up even more since the beginning of the 2020 pandemic. Caution is important here, as decisions made in a rush today may end up hindering progress for decades to come.

What if we moved away from terms such as 'app modernisation' and 'digital transformation'? Instead of associating our approach to making applications timeless with 'transform' and 'overhaul', it may be beneficial to just focus on extending software to generate valuable business outcomes. Map the goals of a tech project, and then work backwards to understand what you need to achieve it.

Understanding legacy infrastructure by looking at the Egyptian pyramids

When you look at the Saqqara Necropolis, the oldest of Egypt's pyramids, do you immediately discard it as an irrelevant relic? Not at all. You admire it for the architectural skill involved in its construction, and the story behind its eternal beauty.

BY MARC ZOTNER, FIELD CTO, VMWARE TANZU



Imagine how your most business-critical applications would look like in a decade. Would they be simply structured, easy to understand, highly-functional, low-effort to evolve?" App modernisation shouldn't be approached 'tool-first', 'technology-first', or 'pattern-first'. Don't start mingling with tools and technologies without first clearly understanding what you are trying to build, and why you are building it. For example, an event-driven or microservice architecture does not make sense for all projects. Taking a step back to understand the business goals first is essential. Then a proper notional architecture can be derived before the right tools and frameworks should be selected to achieve these.

Most modernisation projects are focused around fixing a certain problem IT teams have identified by adopting a not necessarily related technology. In reality, modernisation for our customers is about making new requirements possible, saving cost, scaling capacity, increasing stability and security, at a higher speed. It digitally provides them a competitive edge while boosting business growth.

The key to success lies in first understanding the truly desired outcomes of the modernization work and keeping them prioritized. Defining early qualitative objectives and measurable key results is a spot-on practice to guide transformation journeys. Without such a compass, it is extremely easy to get lost in the depths of our majestic monolithic construction or swamped in a never-ending analysis-paralysis.

This mindset empowers companies to successfully approach brownfield transformations, as well as greenfield cloud initiatives. While the root causes of the core concerns are often obvious, proper solutions are not. Monolithic problems cannot be solved with the same thinking used to create them.

Ensuring a timeless approach to app modernisation

There are four elements of a successful application modernisation initiative. The first is to begin having the end in mind, aligning all stakeholders on goals and non-goals. It is key to start small: even if your portfolio contains thousands of apps. Start with a single business unit and a handful of applications that matter, where a huge impact could be made. Test-driven development, continuous integration and continuous deployment can then be used to reduce manual processing time and increase determinism through automation. It is also important to just plan enough to start, learn on the job, informing strategy and building new skills through hands-on effort, and loops of rapid feedback and result measurement. Lastly, break things down – iterate quickly and continually on thin slices of complex systems focussing on high-level architecture before technology. Following these tenants enables tech teams to quickly focus on the right things and deliver impactful, iterative results.

Recently, VMware Tanzu worked with a large European public sector customer to support them with the modernisation of a mainframe application being over 30 years old and requiring several millions of euro yearly as maintenance and running cost. This involved reworking software and processes impacting millions of citizens and tens of thousands of officials. The challenge was to find a way to start rewriting the monolith to a cloud-native area, with a modern programming language, without impacting the implementation changes required by the latest legislation. By using the approach outlined above, in only a few weeks, the customer was able to get the first components of the mainframe rewritten in Java and running in production. The path from code change to production has been reduced from four weeks to 30 minutes. Several engineers and developers have been empowered to further modernize the application, as if they would be building a new one, delivering value in shorter development cycles on a weekly basis – with no downtime.

Now, back to Egypt

In the years following the Saqqara, many Pharaohs based their monuments on the same concept, look and approach. Over time, other innovations were incorporated - including etching hieroglyphs on the stones - and plans modified.

Just as historians and archaeologists approach the Egyptian pyramids, empathy and context are crucial tools for app modernisation. The end game for archaeologists is not to put beautiful relics in museums, but to really understand the story, culture and the motivations of the people who built them. Similarly, CIOs must first understand the mindset of developers and technology teams at the time of the initial concept, to truly be able to embody a fresh mindset for app modernisation.



Catchpoint releases enhanced version of the WebPageTest API

CATCHPOINT announces the general public availability of the WebPageTest API (WPT API), previously limited to a small number of users, along with several upgrades including deeper performance metrics, immediate test results, dedicated support, developer resources, and CI/CD integrations.

The enhanced WPT API provides direct and programmatic access to WebPageTest data and test infrastructure, which is the industry's de-facto web performance testing solution. WPT offers in-depth front-end performance metrics and side-by-side video comparisons of user experience from the latest browsers around the world.

"We're thrilled to make the WebPageTest API available to the entire development community," says Mehdi Daoudi, CEO, Catchpoint. "It will be exciting to see the ways developers leverage the API's flexibility to optimize site performance and integrate page speed metrics into their development workflow."

Continuous Delivery of Faster Web Pages

The API's portability enables developers to run performance tests from simulated user environments without leaving their CI/CD tool to continuously deliver fast web pages.

Catchpoint launched their official GitHub Action for WebPageTest, making it possible for developers to test their code changes against performance budgets and get detailed performance information right in their existing workflow.

"Having WebPageTest data integrated with your CI/CD process brings performance into the development conversation, ensuring sites meet speed and usability standards before they're released and experienced by users in production," says Patrick Meenan, Founder of WebPageTest and Engineering Fellow at Catchpoint. "Developers will benefit from shortened feedback loops, accelerated delivery, and deeper performance visibility at all stages of development."



Deeper Performance Investigations, Automated

In addition to development integration, the WebPageTest API also provides tremendous value to performance engineers and technical SEO consultants, who can now automate deeper performance investigations and audits to surface new opportunities to improve site performance.

"WebPageTest is the industry standard for page load performance analysis. Being able to isolate individual factors with high confidence and repeatability makes it easy to pinpoint and communicate how incremental changes can improve things for the better," says Alex Russell, Software Engineer, Google Chrome. "The best professional web performance investigators I know use WPT as a critical part of their workflows, and it's the center of mine."

Supporting a long-term Core Web Vitals strategy

The need to push page load times even faster is underscored by Google's plans to incorporate performance indicators, Core Web Vitals, into their search ranking algorithm starting in May 2021. With lucrative search traffic at stake, Core Web Vitals metrics have become a universal barometer of performance and top-down organizational imperative.

"WebPageTest highlights the factors impacting Core Web Vitals so developers can optimize the performance and subsequently, search ranking, of their top pages," Says Tim Kadlec, Performance Engineering Fellow at Catchpoint. "Enforcing Core Web Vitals performance budgets within your release process ensures all future pages are high-performing in the long-term."

3rd Party Platform Integrations

Additionally, performance-focused technology providers leverage the WebPageTest API to build complementary product and solution integrations to automate image, delivery, and vulnerability analysis, designed to improve customer page load speed and security.

"Today, vulnerability management is no longer a nice-to-have but rather a critical requirement for anyone running web applications or interactive and static websites," says Liran Tal, Director of Developer Advocacy at Snyk. "These kinds of public-facing assets are common attack vectors for malicious actors seeking unauthorized access to systems and data. Snyk uses WebPageTest to perform 3rd party library detection and vulnerability scanning that reveals potential security vulnerabilities in your website pages."

Prioritise, predict and act with BMC Helix

BMC has introduced the new AI-driven IT Operations (AIOps) and AI-driven Service Management (AISM) capabilities for the BMC Helix portfolio that will enable IT service and operations teams to predict issues better, resolve them faster, and provide always-on service. These capabilities are powered by the new BMC Helix Platform, which delivers open, cross-domain engagement, observability, and actionability.

“A digital enterprise requires some way to orchestrate the ecosystem of observations, people, data, actions, and assets which deliver value to both customers and employees,” said Shannon Kalvar, Research Manager at IDC. “IT Service Management tools, which have traditionally done that coordination on a small scale, have grown and evolved in the last few years to meet the need more broadly, allowing for both semi-autonomous IT operations and more integrated enterprise service management.”

Service and Operational Excellence for the Autonomous Digital Enterprise

With industry-leading capabilities, the BMC Helix portfolio seamlessly provides comprehensive, unified IT service and operations management solutions that enable organizations to:

- Work smarter thanks to AI-derived insights that help prioritize and resolve problems quickly and proactively;
- Be more productive and efficient through reduced rework and automation of tedious manual processes; and
- Continuously deliver high-performing digital experiences and uninterrupted service through improved collaboration and visibility across IT service and operations teams.

Discover and Connect Assets and Relationships

The new BMC Helix Platform helps enterprise teams, including DevOps, meet their quality, velocity, and compelling experience goals through observability and automation capabilities. It joins IT service and operations teams by providing a common user interface with rich visualization that expands visibility and context, supports cross-launching from one solution to another,

and offers configurable automated workflows to free up resources and quickly take corrective actions. These workflows can be used to automate tasks, such as deploying software requested by a user or collecting additional details about an incident for enrichment.

Finally, the platform supports rapidly changing environments and keeps them up to date with dynamic service models that ingest metrics, events, and topologies from BMC and third-party solutions.

Streamline Proactive Problem Management

By applying AI-natural language processing to detect clusters of recurring incidents automatically, the BMC Helix solution reduces investigation time significantly by figuring out root cause and other causal impacts and enables a seamless transfer from problem identification to investigation to closure. This capability frees up employees to focus on high-value problem areas, reduce incident management load, and prevent adverse impacts to service performance and availability.

Improve Service Quality with Event Management, Noise Reduction, and Probable Cause Analysis

The BMC Helix solution reduces the amount of time it takes to investigate and resolve groups of incidents by automatically identifying incidents related to the same situation. Users can collect, correlate, and apply intelligence to analyze events based on policies to reduce event noise and better understand the issue. Probable cause analysis can then be used to identify the root cause, allowing teams to resolve issues faster, more accurately, and with less effort.

Increase Service Performance and Availability with Service-Centric Probable Cause Analysis

BMC Helix shows the configuration items and related events that are most likely the root cause of a potential performance or service impact. Users can analyze data and identify a single or group of metrics that are behaving abnormally to trigger events and notifications accordingly.

Gain Real-Time Visibility with Service-Centric Monitoring

BMC Helix gives users the ability to use new, user-friendly visualizations, such as heatmaps and tile views, that show a business service's current state. This real-time view helps teams avoid any potential issues based on trends and increases service performance, availability, and quality.

Implement Advanced Analytics for Service Desk and Change Management Processes

Additionally, BMC Helix identifies areas and opportunities to optimize and drive efficiencies in service desk speed, quality, and resource allocations. It also helps users gain a more comprehensive understanding of change interdependencies and the best ways to implement new services.

Improve Service Assurance and Optimization

Users can dynamically align and continuously optimize infrastructure resources to meet the unique needs of applications and changes in resource demands across applications and services. The BMC Helix portfolio supports Kubernetes, microservices, containers, and pods, as well as private, public, hybrid, or multi-cloud environments.

Deliver Enhanced, Personalized Experiences

The BMC Helix solution set now comes with new personalized dashboards that visualize data about incidents, changes, service requests, chatbot conversations, metrics, events, capacity, and more across IT service and operations management. It includes a flexible, persona-based user interface and user experience that optimizes and maximizes productivity and efficiency.

“We continue to embed AI, analytics, and automation into new and existing offerings to help customers evolve their businesses with better intelligence and decision making so they can prioritize, predict, and act with confidence,” said Margaret Lee, General Manager, Digital Service and Operations Management at BMC.



The key to unlocking optimum Digital Employee Experience: Proactivity

In the age of remote working, IT is now vital to ensure the success and running of organisations. Even with work almost entirely digital, there is still a considerable gap between the optimised and seamless services IT teams aim to deliver, and the actual quality of the digital experiences of employees.

BY CHRIS TERNDROP, BUSINESS TRANSFORMATION ARCHITECT, NEXTHINK



IN A RECENT STUDY, 96% of technology executives agree that digital employee experience (DEX) is an essential part of what IT teams do, yet more than a third (34%) rely on manual methods to collect experience information and nearly half (46%) don't measure their employees' digital experience at all. If you're not measuring your employee's experiences, how can you ensure that this essential element is maintained?

Too often IT teams are one step behind, playing catch up with their employees' IT problems. Systems look good, all signs point to green, but bluescreens, application failures and slow interfaces cause dissatisfaction. Over time as a result there has been a culture shift; many employees will fail to report technical issues, losing trust in IT's ability to help, or perceiving the process as lengthy and frustrating. A

survey released in 2020 revealed that workers report just over half of all the technical issues they face.

The rapid transition to working from home in 2020 only exasperated these IT cracks further, introducing new problems like messaging application failure, slow internet, or video calls dropping which can cut off employee communication entirely. In a recent study, 81% of employees say they either don't want to go back to the office or would prefer a hybrid schedule from now on.

Therefore, with remote or hybrid working here to stay, it will be important for IT teams to prioritise gaining visibility and insight into the digital experiences of employees, while taking a more proactive and forward-thinking approach to technical solutions.

IT remote actions and automation - Employee 'Zero touch'

The smartest and most effective way to prevent disruption and avoid IT downtime, whilst improving the employee experience from 'anywhere', is to identify and automatically resolve problems before they become service impacting issues. It is now possible to remove employees from any involvement in many regular and persistent disruptions impacting productivity. Using employees either as a monitoring system for IT or in the fix scenario should be considered a thing of the past.

So, employee 'zero touch' - is this a dream or reality? It will be a reality as there will be far less need to contact the IT team. Involvement of employees in resolving IT issues is definitely reducing, whether by lack of trust or otherwise, and is going down the path to continue to do so. Real time analytics as a data source is key in enabling intelligent focus on the issues that occur the most. IT teams can then apply fixes with remote actions for employees automatically, before employees realise there is even an issue at all.

There are a wide variety of issues that impact many businesses, including those related to high log-on times, performance of collaboration services, software compliance and even incorrect resolution of customer interfaces. Any of these can escalate to kill productivity company wide and quickly bump up IT ticket count. The following two examples showcase how those disruptions were remedied by a more proactive IT approach.

Remote remediation of software glitches

Take the example of a law firm, where IT and employees now benefit from this proactive approach. One key IT issue was that critical software was opening incorrectly and minimised, always causing a support ticket and a poor employee experience. A remote action is now able to detect this issue by checking the registry, and can change the value to automatically remediate the issue, thus reducing productivity impact and the need for raising a support ticket.

Automating screen resolution

At this large national retailer, in store screens for customer use were displaying at the wrong resolution, which made a higher number of them unusable. This directly impacted trading and their bottom line. In order to solve this, a remote action has been put in place which runs to detect the resolution. If it is incorrect, it will resolve this to the correct setting automatically. This results in fewer tickets and improved satisfaction for employees.

Smart device management

Outdated or non-compliant devices are responsible for a myriad of potential IT issues, from slow performance and application instability to security weaknesses. Despite their impact, they remain one of the areas



which IT teams notoriously struggle to identify and fix. This problem is even harder when teams are working across multiple network and device versions. In the majority of cases, IT teams manually resolve outdated devices or applications on an individual basis, as their visibility of issues is limited to when employees submit tickets. Acting proactively to prevent outdated issues would minimise disruption and improve resource efficiency. The answer lies again in real-time monitoring and automated self-help campaigns.

For example, a proactive IT team could remotely identify devices that required updates, and deploy targeted restart alerts. A simple click by the employee could trigger a forced update, stopping the issue in its tracks before it could get started and save the IT team hours of productivity. By keeping every device in their landscape compliant and updated, IT teams can avoid future catastrophe and easily maintain a high standard of Digital Employee Experience.

Stay proactive

As we move forward into a future where WFA (work from anywhere) and hybrid home-office working is the norm, it's imperative that IT teams are one step ahead of technical issues as organisations cannot afford to be playing catch-up in this area. Employees should be able to rely on a seamless digital experience to collaborate with colleagues and remain productive.

By adopting a proactive approach based on real-time analytics to inform targeted alerts, problems can be pre-emptively resolved by employees themselves or by IT detecting and automating the solution completely, removing the need for any employee involvement. Through engaging the workforce in this way, IT can become a source of positive employee sentiment, rather than a topic synonymous with disruption or confusion.



Three tips for creating outstanding digital customer experiences

If you look at businesses that have thrived during the pandemic, the likes of Dominos and Amazon, you will see that they have something in common; excellent customer experiences that are driven by agile and innovative internal business cultures.

GREG OUIILON, CHIEF TECHNOLOGY OFFICER EMEA AT NEW RELIC



It has never been easier than it is today for customers to shop around for alternative products and services - agility and business innovation are the backbone of customer experience. Now more than ever, business leaders need to ensure they are agile enough to deliver a frictionless journey to maintain a competitive edge on an even playing field. Possessing the correct tools, agility, and internal processes to spot and solve

issues long before users do is critical as it creates an enhanced developer and customer experience; it can be the difference between businesses that thrive and those that fall behind.

There are a few top tips businesses can follow to ensure happier customers, fewer bugs, and faster page loads.

Legacy “page load time” metrics insufficiently measure what matters most to customers. Modern websites are dynamic in comparison, they look to user-centric performance metrics to measure success that leverage effective benchmarking tools and improve overall page performance. Combined with technology that provides user-focused performance metrics and crash diagnostics, this approach will ensure that technology rollouts have the customer at their core

Agile developers are productive workers

Agile developer teams tend to be more creative and productive due to their ability to swiftly adapt to new environments. Key to this is a user-centric approach, Digital Experience Monitoring (DEM) helps teams focus on customer journeys and web performance as the primary signals for software health.

Companies proven to have successfully transitioned during the pandemic are also adopters of approaches that focus on the user to engage with customers. Operating under a strong, cohesive unit helps to foster a team culture of customer success.

Customer-first businesses

Legacy “page load time” metrics insufficiently measure what matters most to customers. Modern websites are dynamic in comparison, they look to user-centric performance metrics to measure success that leverage effective benchmarking tools and improve overall page performance. Combined with technology that provides user-focused performance metrics and crash diagnostics, this approach will ensure that technology rollouts have the customer at their core.

For example, Australian fintech Nimble went from not knowing when or why a customer abandoned their cart to effectively observing customer behaviour and conversions in real time. The developer teams can now proactively address potential challenges faced by customers and make informed decisions.

Answering market needs

To be better placed to meet the needs of the markets, businesses are migrating to the cloud and continually optimising the resulting user experience for their customer base. It is vital to plan the move to the cloud carefully and establish baselines beforehand so that issues can be tracked and traced. In turn, improvement and success can be demonstrated as soon as the migration is complete.

The top obstacle that DevOps teams face is an increasingly complex web of systems and infrastructure along with the intricate, distributed web

of agents monitoring all components of the stack. Devs require a single source of truth to instantly observe what is going on in real-time so that they can immediately fix issues well before they escalate. The right tools and data assets are needed to enable swift action to be taken.

Businesses cannot afford to deliver poor customer experiences as downtime has a hard price to pay, Gartner has found that the average cost of downtime is approximately £4,100 per minute. That said, there are also additional fees that stack up including productivity falling, teams unable to carry out other tasks, as well as growing stress and fatigue taking a toll on employees.

The secret to maintaining seamless experiences for the business lies in the dev team - if they can access full observability across the entire stack, maintaining uptime, availability, and performance will come naturally.

It is important that they are able to collect, correlate, and contextualise data easily - the key to success is a cross-functional team with a unified customer experience management strategy, a shared understanding of the user journey, and a full-stack view of softwares’ performance.



Companies consider work from anywhere

APPNETA has revealed the results of AppNeta's 2021 State of Work from Anywhere Outlook Report, which identified end-user expectations for the future of work, highlighting the rise of "work from anywhere" and how this shift in the workforce model will impact IT teams spanning industries in the near future and post-COVID.

As many companies across the nation continue remote work, the rise in vaccine availability has challenged business leaders to identify a logical next step for their employees, keeping safety and talent retention top of mind. The study, of 1,000 people across the United States, examined the views of Americans that rely on the internet to do their jobs and identified critical areas to address in order for a business to successfully take on a work from anywhere model.

The Workforce has Moved on - and they are Unwilling to Return

According to AppNeta's 2021 State of Work from Anywhere Outlook Report, nearly 80% state that their preferred work environment long-term would include an element of remote work. Employees are looking for flexibility and a hybrid model, either hoping to work remotely permanently or have the option to come into a physical office only when needed.

Further, nearly three-quarters would be interested in their employer adopting a "work from anywhere" approach, allowing them the freedom to pick and choose where they set up work each day. Surprisingly, since March of 2020, 21% of respondents relocated from their original address, with the American workforce now nearly evenly dispersed among urban areas, suburban areas and rural areas, identifying the need for companies to be able to provide the same level of user-experience across vastly different regions.

This broader trend of urban decentralization creates new challenges for already strapped IT teams, as delivering optimal internet connectivity to residential and rural communities can be a challenging task, requiring employers and IT teams to set user expectations around the quality of their network performance.



Managing Employee Expectations with Technology

The events of the past year have challenged companies regardless of industry to move entire workforces to operate productively at home, essentially creating IT support needs for individual offices, where each employee resided amid the pandemic. Of the technology-related issues causing frustration in the past year for remote employees, the biggest gripe was internet connectivity, with almost half (44%) of respondents expressing their frustration. Another top stressor was issues with video calls, with 40% of respondents identifying freezing screens and challenges with popular tools such as Zoom. Unsurprisingly, more than one-third of end-users shared that they were frustrated by technology challenges and experiences with their employer's IT team since the pandemic, even though 21% of respondents acknowledged that the IT team may be doing their best in the circumstances.

"Now a year into the pandemic we are seeing organizations define more concrete plans for their long-term work strategy. Whether the approach is fully remote, or a hybrid setting, IT should be involved in these conversations early on. Anything is possible with the right resources, but Enterprise IT, and the Solution Partners they depend on, absolutely need to play a strategic role in how the business transforms to better support the new style of working in the year ahead. That may include adding additional solutions to their arsenal," said John Tewfik, Director of Global Alliances, AppNeta.

As business leaders strategize what the "new normal" will look like for their

organizations, those adopting "work from anywhere" will want to set clear expectations for IT support with remote employees. In a "work from anywhere" environment, employee expectations are high, with nearly half of those surveyed wanting support from the IT team with critical applications they use, 37% expecting support for internet connectivity issues, 35% expecting support with hardware and one-third expecting IT to support them in learning any new tools provided by the company.

"The pandemic completely blew up most IT departments' support model for their remote users," stated Matt Stevens, CEO of AppNeta. "Remote office visibility for the end user experience of business critical applications was already a challenge for many, but in the new normal, the 'user-to-problem ratio' is now often 1-to-1 vs. hundreds-to-1 that used to be the norm with large groups of users consistently working on a regular schedule from a given office location. The IT prioritization process of understanding truly critical applications and their associated users and achieving alignment with the lines of business has never been more critical to the success of the overall business."

Employees want IT's attention, and they want it quickly

In 2021, the IT team at many organizations has been in high demand. According to survey results, nearly 30% of employees anticipate a response from IT within the hour of them reaching out with a tech-related issue. An additional 33% think that a response within the business day is acceptable, while unsurprisingly only 11% feel that a response within the week is acceptable. In a "work from anywhere" model, employee expectations need to be tempered in order to maintain clarity around IT's responsibilities and protect them from burnout.

Where possible, business leaders should add clarity around responsibilities and response time and seek out technology solutions that are able to scale to the dispersed workforce and assist in network performance monitoring to ease the burden of internet connectivity and support with critical apps.

BigPanda launches Automatic Incident Triage

BIGPANDA has launched Automatic Incident Triage, a new platform component that significantly reduces the manual toil associated with the triage phase of incident management. Automatic Incident Triage reduces the “mean-time-to-resolve” (MTTR) for applications and services by enabling IT Ops and NOC teams to quickly triage incidents by reducing the steps required to fully understand the business context of an incident and assign it to the right response team within their desired collaboration platforms.

“Streamlining processes is a critical component of technology operations,” said Rob Scarmuzzi, Executive Director of Operations Technology at E*TRADE Financial. “Automating tasks, like consolidating events, equips us with the tools to manage our workflow efficiently and ultimately freeing up time to deploy manpower to areas that require attention. BigPanda’s Automatic Incident Triage enhancements put additional firepower behind these automated capabilities.”

Enterprises with complex hybrid IT infrastructures and organization structures face a growing number of challenges, including centralized and decentralized Ops teams, and hybrid environments with on-prem and cloud-based applications and tool sprawl, making it difficult to rapidly understand, investigate, remediate and resolve incidents. According to Gartner,



“organizations are struggling to reduce incident response times because of delays around manual incident routing and cross-team collaboration challenges with incident response.” Gartner goes on to state, “Depending on the organization, gathering the context of the incident often takes 15 to 30 minutes, which significantly impacts mean time to resolve (MTTR).”*

An inability to quickly gather business context in the incident triage phase delays incident response times, which negatively impacts service availability and reliability, creates user satisfaction issues, and drives up operational costs. BigPanda’s Automatic Incident Triage helps IT Ops and NOC teams solve this pain point, improve NOC productivity and reclaim high-value L3 and DevOps FTE hours.

“Time is one of the biggest enemies of IT Ops and NOC teams. Incident responders know all too well how long it takes to answer the ‘What next?’ question once they’re presented with an incident,”

said Elik Eizenberg, Co-Founder and CTO at BigPanda. “Automatic Incident Triage turns what used to be a technical incident into a business incident automatically, helping incident responders rapidly triage and handle more incidents than before and quickly route critical incidents to the right teams for follow-up and resolution.”

With Automatic Incident Triage, BigPanda customers can:

- Automatically calculate and incorporate detailed business context into incidents, such as validated incident severity, impacted services, business priority and routing information using easy-to-create custom incident tags.
- Quickly and easily sort, filter, visualize and respond to the incidents, prioritizing those with either the most pressing validated incident severity or the number of impacted services.
- Bi-directionally sync custom incident tags with collaboration tools such as ServiceNow or Jira to deliver easier mapping of fields and trigger easier workflows within those tools.

Automatic Incident Triage allows Ops teams to handle higher volumes of actionable incidents themselves, without having to escalate as frequently. And when they do escalate, the additional business context makes it easy to prioritize and route incidents to the right teams for faster resolution.

Anodot and Coralogix partner

ANODOT and Coralogix have partnered to provide an integrated solution that monitors business operations and applications faster and more accurately than ever before.

The integration between Anodot’s business monitoring and Coralogix’s data platform enables customers to dramatically improve how they leverage data to find and remediate issues they would have otherwise missed while solving log maintenance autonomously. One of those customers is Payoneer, Inc., the global payment and commerce-enabling platform which powers growth for millions of small businesses, marketplaces and enterprises, including Amazon, Google and Walmart – for nearly 200 countries and territories.

Payoneer delivers a suite of services that include cross-border payments, working capital, tax solutions, risk management and payment orchestration for merchants.

With more than five million customers worldwide, Payoneer monitors millions of business and technical metrics to keep their payment gateway running smoothly. Previously, they used traditional monitoring and log analysis solutions that took at least 24 hours to resolve and resulted in high false positive rates. To hasten the digital payment process and improve customer experience, Payoneer integrated Anodot and Coralogix.

“This integrated solution is a game-changer for us,” said Yuval Molnar, senior director of production services, Payoneer, Inc. “We have over 1,000 services and must understand their behavior to protect our clients. In doing so, we monitor all logs and autonomously detect errors and false positives. Anodot provides a zoom into the exact Coralogix logs that show the root cause. This has improved our time to resolution by 90% which is a true game changer.”

Why is APM important?

Breaking down the benefits

New to application performance management (APM)? Read a quick primer on how it works and why it's become a business necessity.

BY BRENDAN COOPER, PRODUCT MARKETING MANAGER FOCUSED ON APPDYNAMICS SAP MONITORING SOLUTION



LET'S FACE IT: Our lives revolve around applications. We use them to engage and transact with our customers, to drive our productivity at work, in our downtime to stream videos, browse the internet, and stay connected with friends and family. And we expect them to perform flawlessly every time. Think about the last time you experienced poor app performance. Did you try to use the application again, or did you — like most consumers surveyed — simply move on?

The reality is that the new competitive battlefield is digital, and the winners and losers are defined by the experience they provide. But this unforgiving demand from users around the world has put technologists under massive pressure to keep their applications continually available and operating at peak performance, at the risk of losing customers and revenue and, ultimately, damaging their brand's reputation. This is where application performance management (APM) comes into play. And with the above in mind, it's easy to argue that APM is not only important, but an absolute business necessity.

What is APM and why is it important?

In short, APM is the practice of proactively monitoring the many facets of an application environment in order to identify and mitigate issues before they

become major problems. Across the typically complex and distributed ecosystems of today's applications, we can think of APM as our guide in finding the needle (or needles) in a digital haystack that spans multiple locations, across various types of technology.

But why do we need APM, specifically?

Let's look at a few key APM benefits and the role they play in solving performance problems.

Benefit #1: APM breaks down operational silos

APM provides a unified view across your entire application stack, including every component, connection point, dependency, and user interaction. This benefits the different teams supporting your application by equipping them with comprehensive visibility, which allows them to collaborate in a way that would be virtually impossible without APM. This is particularly important in the highly distributed, multi-cloud environments that support so many modern applications today.

Benefit #2: APM allows you to meet — and exceed — customer expectations

As we discussed earlier, when an application experiences performance issues, or is unavailable, you run the risk of losing customers. APM provides

Prioritizing the right way with APM

At any given time, multiple issues could be impacting your application's performance. APM provides the insights you need to prioritize which issues matter most.

Priority One: Business Impact

Is the performance issue impacting a critical business metric? *Example: Revenue impact due to an issue with the check-out functionality of the application*

Priority Two: User Impact

Is the issue impacting the user experience? *Example: An issue impacting load-times, thus increasing user drop rates*

Priority Three: Everything Else

Is the issue impacting overall application performance? *Example: An impact that's degrading performance, but not necessarily impacting business KPIs or user experience*

Defining the application stack

Front-end Monitoring	The performance of your applications through the lens of an end-user
Back-end Monitoring	The performance across the various services built into your applications to support the functionality required by your end-users
Infrastructure and Cloud Monitoring	The performance of the infrastructure supporting your applications including servers, databases, the network, and the various cloud services you may be using

real-time performance insights that allow you to react fast when issues arise, including contextual data that helps you reduce the mean time to resolution (MTTR) and restore your application to normal performance. By proactively resolving issues, you're better able to provide that flawless experience your customers expect from your applications.

Benefit #3: APM protects your company's bottom line

More sophisticated APM solutions provide business intelligence analytics, which can help you visualize and understand how application performance issues impact your mission-critical business metrics — revenue or sales conversions, for example. This not only helps better align IT with the business, but it also helps technologists prioritize issue remediations by focusing on resolving the problems that directly impact key business outcomes.

What's important in the world of APM going forward?

APM solutions are continually evolving to meet the demands of the rapidly changing technologies we use to build applications. At the time of writing, we're on the cusp of a major shift that will bring APM to the next stage of its evolution. The two driving forces behind this shift are observability and OpenTelemetry, which, at a high-level, can be considered to go hand in hand. Observability has a pretty fluid definition, but in general, you can think of it as APM on steroids. Driven by the advanced needs of DevOps and SRE teams, observability provides the raw, granular data

necessary to gain an in-depth understanding of complex and highly distributed systems — typically defined as M.E.L.T. (metrics, events, logs and traces). This fine-grained understanding of how applications and systems should work will help you further reduce MTTR when issues arise. Sounds great, right? But it's extremely challenging to get the correlated M.E.L.T data needed to make this a reality. Enter OpenTelemetry. OpenTelemetry is a vendor-neutral standard for collecting telemetry data for applications, their supporting infrastructures, and services, providing the consistent collection mechanism and format needed to understand and validate performance across the most complex of distributed applications.

While observability and OpenTelemetry are still in their infancy, they're both a testament to the industry's ongoing efforts to simplify complexity and ensure applications are always driving better digital experiences and business outcomes. Be on the lookout for developments in this space.

Further reading

As we've learned, APM is a critical component needed to keep your applications available and performing at optimal levels, and will continually evolve at the pace of the technology around us. This is important because your applications are the storefront to your brand, and customer expectations are only increasing. But we're only scratching the surface here. For a complete overview of important APM capabilities and what the future of APM looks like, check out our definitive guide, [An Introduction to APM](#).

How APM helps reduce MTTR

Baseline Performance	Baselines application performance metrics, helping you understand what normal looks like
Automated Alerting	Generates alerts when anomalies deviate performance from normal levels
Root Cause Analysis (RCA)	Provides root cause analysis insights to help identify the core issue(s) and expedite resolution



How to thrive in the enterprise AI era

Enterprise AI is the ability to embed AI methodology — which combines the human capacities for learning, perception, and interaction all at a level of complexity that ultimately supersedes our own abilities — into the very core of an organisation's data strategy.

BY ALEXIS FOURNIER, DIRECTOR OF AI STRATEGY, DATAIKU

IT ISN'T JUST a trend or something to be leveraged for one-off projects and use cases. The ability to become a true AI enterprise by successfully scaling and employing robust data projects and processes at all levels of the company is an organisational asset pivotal to the success of businesses of the future, regardless of the industry.



Many businesses struggle to get started on their Enterprise AI journey, but there are also some who are thriving. The companies that succeed are the ones that go beyond leveraging Enterprise AI for one

particular project or use case and instead focus on scaling it out to a level that will sustain the business in the future. Scaled Enterprise AI companies manage to build a foundation for data science, machine learning, and AI at an organisational level. Here's a quick look at how some of them are fully embedding AI in the enterprise.

Leverage Horizontal (Team-Wide) and Vertical (Cross-Team) Collaboration

When we think about Enterprise AI, a key focus of companies that are achieving success is the ability

and proven capacity to fully embed AI as a regular capability inside new processes and ways of working. This means bringing all people together around AI efforts, from business people to analysts and data scientists. Traditionally, delivering this has been challenging for companies, because it demands a full transformation and a full value chain reorganisation as well as a mass effort to upskill people.

A key part of getting everyone on the same page is creating spaces where data is accessible. These companies are opening sandboxes where people can learn, experiment, and put governance in place that creates a positive incentive to testing. They are also putting the right checks in place before machine learning-based models are industrialised or put into production.

It all goes back to the virtues of collaboration, and the virtues of sharing and democratising initial access, so that even business people, who may not understand everything in terms of data science, are in a better capacity to actually iterate in the right manner with data scientists. In the same environment, data scientists get closer to the business bottom line so that they really understand the needs of the business, and so on.

We may come to know it as 'Inclusive AI,' or simply a concept that encompasses the idea that the more people are involved in AI processes, the better the outcome (both internally and externally) because of a diversification of skills, points of view, or use cases.

Iterate Rapidly

Being able to iterate rapidly on a spectrum of data applications — whether that means building out a self-serve analytics platform or fully operationalised AI integrated with business processes — is key to fully embedding Enterprise AI within an organisation.

The reality around most data projects is that they don't bring real value to the business until they're in a production environment. Therefore, if this process isn't happening quickly enough — both in terms of total overall start-to-finish time-to-insights as well as the ability to rapidly iterate once something is in production — efforts will fall flat. Proper tools that allow for quick, painless incorporation of machine learning models in production are the key to a scalable process.

Speed is also of the essence in that feedback from models in production should be delivering timely results to those who need it. For example, if the data team is working with the marketing team to operationalise churn prediction and prevention emails, the marketing team should have immediate insight into whether the churn prevention emails sent to predicted churners are actually working, or if they should re-evaluate the message or the targeted audience.



Prioritise Data Governance

Data governance is certainly not a new concept — as long as data has been collected, companies have needed some level of policy and oversight for its management. Yet until recently, it has largely remained in the background, as businesses weren't using data at a scale that required data governance to be top of mind. In the last few years, and certainly in the face of 2020's tumultuous turn of events, data governance has shot to the forefront of discussions both in the media and in the boardroom as businesses take their first steps towards Enterprise AI.

Recent increased government involvement in data privacy (e.g. GDPR and CCPA) has no doubt played a part, as have magnified focuses on AI risks and model maintenance in the face of the rapid development of machine learning. Companies are starting to realise that data governance has never really been established in a way to handle the massive shift toward democratized machine learning required in the age of AI. And that with AI comes new governance requirements.

Data governance needs to be a collaboration between IT and business stakeholders. A traditional data governance program oversees a range of activities, including data security, reference and master data management, data quality, data architecture, and metadata management. Those responsible for data governance will have expertise in data architecture, privacy, integration, and modeling. However, those on the information governance side should be business experts, understanding what the data is, where it comes from, how and why the data is valuable to the business, how the data can be used in different contexts, and how data should ultimately be used for optimised business benefit.

Embedding AI in the enterprise successfully is not an insignificant task. However, if businesses can work to understand data at a more detailed level, including its constraints, they are well on their way to understanding what they can do with that data to develop additional insights to transform the way they are working.



How to train your AI algorithm

Successful AI algorithms are built on a foundation of training data, but sourcing data that fits your needs and meets volume requirements is harder than you might think. Particularly when it comes to developing AI-driven applications and smart voice assistants.

RICHARD DOWNS, DIRECTOR NORTHERN EUROPE, APPLAUSE



BUSINESSES FACE several challenges when it comes to training their algorithms to respond to real-world scenarios. Sourcing data at scale is extremely challenging. Businesses need to be able to leverage large and diverse samples, or crowds, of people representative of their target market. It takes a dedicated resource to deliver projects of this scale. In effect, a crowdtesting (or distributed testing) solution, which provides businesses with access to a global community of skilled testers who work remotely. This model provides an embedded infrastructure that can be scaled up or down to meet requirements.

Enterprises and consumer brands have been using crowdtesting services for over a decade. Crowdtesting has become a well-established model that operates in tandem with in-house teams to complement integrated QA testing. Traditionally used to test apps, websites and other digital properties, crowdtesting has become integral to sourcing the data needed to train AI algorithms. It provides businesses with the scope

and scale they need to bring new AI applications to market.

Despite the advantages this model offers there are still a number of challenges businesses need to address. Here we explore three of the key challenges businesses face when sourcing training data.

1. Quantity of data sources

Enormous amounts of data are required to develop an effective algorithm. In the case of training a smart voice assistant developed for the UK market, the algorithm required over 100,000 voice utterances. This eventually required utterances from 972 unique people who were sourced from almost every corner of the UK.

In another example, a business needed to train its AI algorithm to read handwritten documents. The brief was to deliver thousands of unique handwriting samples. The quantity of individuals was a critical

factor, because the algorithm needed unique samples from a broad spectrum of people. More than 1,000 individuals had to be sourced to provide handwritten documents that met the requirements. The size of the crowd was critical to the success of both projects. The majority of businesses don't have access to the large number of participants needed to contribute data. A business could ask its employees to get samples from friends and family, but that would be ineffective and extremely difficult to project manage.

2. Quality of data

So, how do you produce quality training data? Let's return to the handwriting samples. In that instance, the artifacts had to be legible, easily accessible and meet a host of other requirements based on the individual project goals. More specifically, there couldn't be any defects on the page or even a single folded margin in the middle of the page. When users scanned the documents, they needed good light conditions or the ability to use flash in dark settings. During any project there are always specific requirements that need to be tracked and monitored very carefully.

Every individual artifact needs to be tested for quality to assure the algorithm will work as intended. Again, this process takes up a considerable amount of time and resources. While businesses could do this internally, it would prove to be costly and inefficient. In the case of the handwriting samples example, if the business had taken responsibility for analysing every single document to confirm its quality, it would've taken them months and created a logistical nightmare. Instead, the process was completed in a matter of weeks.

3. Diversity of data

Besides producing reams of quality data, your team must also have a diverse range of artifacts to develop an accurate algorithm. Without diversity in the training data, the algorithm won't be able to recognise a broad range of possibilities, which will make the algorithm ineffective.

When building an AI algorithm, you shouldn't rely on a single person to provide the artifacts used to train the algorithm. To train an algorithm properly, you need different types of data and inputs, including geographical data, demographic information, types of documents and so on. Otherwise, the process will not lead to a strong output that will service the needs of a diverse customer base.

A crowdsourced community provides businesses with access to a global pool of participants. This model enables businesses to select hyper-specific demographics, including gender, race, native language, location, skill set, geography and

many other filters.

Evolve with Your Project

Unfortunately, no project ends up exactly how it started. Needs shift over time and you have to change footing, get new data points, and source new testers or resources to input the information as the project evolves. When embarking on a project like this, always consider how you're going to manage that data input and data quality process.

By continually ingesting new data, algorithms identify new trends and patterns and automatically adjust their predictions and outputs to better reflect the current landscape. However, AI algorithms are only as good as the quality of data they receive. Crowdstesting provides the diversity, connectivity and scale required to meet the demands of training AI algorithms and testing AI applications.



What is an AI strategy and how do you build one?

We now live in an age of AI implementation. AI has completely revolutionised the way we work, interact and go about our lives.

BY MICHAEL CHALMERS, MD EMEA AT CONTINO



MODERN DATA SCIENCE is no longer about researching breakthrough AI, but focusing on solving business problems with existing tools and proven algorithms. From manufacturing to marketing, AI is increasing productivity and streamlining business operations.

In an increasingly crowded marketplace, business leaders are looking to AI-driven solutions for an edge over their competition. However, many businesses adopting AI lack the key ingredients of strategy, operating model and execution framework necessary for achieving business-wide AI adoption.

An AI strategy outlines how the technology will achieve set business objectives and, more importantly, identify AI assets that grant the business a competitive advantage and are difficult for competitors to replicate.

Now that we know what an AI strategy is, here are some tips to create the right one for your business align your organisation and effectively execute.

1. Align your corporate strategy with your objectives

During AI transformation projects, companies often make the mistake of separating the vision from the execution, resulting in disjointed and complicated AI programs that can take years to consolidate. This can be easily avoided by choosing AI solutions based on concrete business objectives that have been established at the project's outset.

It's important to align your corporate strategy with measurable goals and objectives to guide your AI deployment. Once complete, the strategy can be easily escalated down into divisional- or even product-level strategies.

2. Establish a multi-disciplinary AI team

Form a multidisciplinary team to assess how the AI strategy can best serve their individual needs. Having members from different departments in your AI team, for example, web design, R&D and engineering, will



ensure your strategy will meet objectives for key internal stakeholders.

You may not deploy the right strategy in the first instance, so iteration is crucial. By fostering a culture of experimentation you team will locate the right AI assets to form your unique competitive edge.



3. Choose the right problems

This might seem like common sense, but the problems you're looking to overcome have a large impact on your success. Some problems are not AI problems at all, and for the ones that are, the business should advocate the delivery through small lighthouse projects that act as a beacon for their capabilities.

In identifying 'lighthouse' projects, your business will need to assess the overall goal and importance of the project, its size, likely duration and data quality.

Lighthouse projects tend to be able to be delivered in under eight weeks, instead of eight months, and will provide an immediate and tangible benefit for the business and your customers. These small wins are then multiplied to sow the seeds of transformation that swell from the ground up, empowering small teams that grow in competency and display increased autonomy and relatedness.

4. Execute backwards from the value chain

Customer-centricity has become one of the most popular topics among today's business leaders. Traditionally, a lot more businesses were product-centric than customer-centric. Products were built and then customers were found.

When creating your AI strategy, create customer-centric KPIs that align with the overall corporate objectives. It is important to constantly measure product execution back to these customer-centric KPIs. By taking a customer-centric approach, you will find a lot of the technology decisions are now decided by business drivers.

5. Scaling out your AI community of practice

The journey to business-wide AI adoption will be iterative and continuous. Upon successful completion of a product, the team should evolve into what's known as an 'AI community of practice', which will foster AI innovation and upskill future AI teams. In the world of rapid AI product iterations, best practices and automation still apply and are in fact more relevant than ever. Data science is about repeatable experimentation and measured results. If your AI processes are non-repeatable and everybody is changing production by hand, then it is no longer data science but data hobby.

6. The AI strategy is a continuous journey

As with any successful project, the formula for enterprise-wide AI adoption is: nurture the idea, plan, prove, improve and then scale. Lighthouse projects will need to be proven to work. Teams will need to be upskilled. Processes will need to be streamlined. There will be mistakes made and lessons learnt. And all of this is okay. Businesses need to focus on a culture of learning and continuous improvement with people at the centre, through shorter cycles, to drive true transformation.

By focusing on experimental culture and continuous improvement, through shorter cycles, to drive true AI transformation, business can drive true AI transformation. An AI strategy acts as a constantly evolving roadmap across the different business functions (people, processes and technology etc.) to ensure your chosen solutions are working towards your business objectives. In short, let your business goals guide your AI transformation, not the other way around.

U-Haul speeds its digital transformation with Dynatrace

U-HAUL, the world's largest do-it-yourself moving and self-storage company, is using the Dynatrace® platform to speed its digital transformation. In response to the pandemic, U-Haul teams needed to react quickly and develop new digital services, including mobile and web-based applications to enable contactless reservations and pick-up of moving equipment.

With the Dynatrace® Software Intelligence platform, including its Applications and Microservices, Digital Experience, and Business Analytics Modules, the teams gained precise, real-time insights about critical KPIs, such as feature adoption, conversion rates, and user abandonment. This has enabled them to improve user journeys, shorten development cycles, and accelerate the impact of their efforts on the business.

"Moving is stressful enough, and our digital services need to make it as easy as possible for our customers to manage their moves," said Brian Rutherford, Director of Software Development at U-Haul. "Dynatrace is one of the best solutions we've ever deployed because it gives us the precise, data-backed information we need to understand how our customers interact with our apps, and how these interactions impact the

business. Insights from Dynatrace help our team discover new services and features we need to develop to meet our customers' needs. They also ensure we stay focused on activities that deliver maximum value for our company."

With the Dynatrace® platform's Digital Experience Module, teams at U-Haul now understand precisely how their customers are using new features, which helps them proactively optimize user journeys for better conversion rates and higher revenue. Dynatrace automatically connects this user data with application, infrastructure, and log data, enabling the teams to understand all dependencies, anomalies, and degradations across U-Haul's entire, dynamic cloud environment.

This advanced level of observability, with everything connected and in context – from a tap, click, or swipe, through to back-end, hybrid services – allows the teams to improve code quality, accelerate new feature adoption, and shorten development cycles quickly and easily.

With AIOps capabilities built into the core of the platform through Dynatrace's AI engine Davis®, anomalies in application performance or user behavior are automatically flagged and prioritized based on business impact,



which eliminates the guesswork often associated with other less precise AIOps approaches.

"Dynatrace has allowed us to move away from a 'war room mentality, where we were on-call and meeting at any hour of the day to troubleshoot problems," continued Rutherford. "The Dynatrace platform's Session Replay capability is like having an instant focus group with all our customers, across 22,000 locations.

The proactive assistance to improve our customers' experience, and the insights we gain from this, are tremendously helpful, not just for IT and development, but for our executive management as well. I can confidently say U-Haul is coming out of the pandemic stronger than ever, in part because of the way Dynatrace has fundamentally changed how we work, build our applications, and support our customers."

Micro Focus recognizes Gazprombank

MICRO FOCUS says that Gazprombank, one of the largest multi-faceted financial institutions in Russia, has received the Digital Transformation Innovation Award: EMEA Region in recognition of its exceptional creativity in leveraging technology to win in today's digital economy. The award was announced at Micro Focus Universe 2021, the company's flagship customer event.

Gazprombank earned the coveted Digital Transformation Innovation Award, by leveraging a variety of Micro Focus products to create a new strategic IT platform and completely transform the IT department.

The effort leads to increased agility, openness to innovative ideas, and increased engagement with the business. "Gazprombank is well known in the FinTech industry and being at the leading edge of applying technology to help deliver even greater value to their customers, and we are excited to play a part in that," said Gonzalo Usandizaga, President of Sales EMEA and LATAM at Micro Focus. "With a comprehensive portfolio, underpinned by a robust analytics ecosystem, Micro Focus delivers pragmatic, customer-centric solutions that enable customers to both run and transform at the same time. We often talk about taking a High Tech, Low

Drama approach to solving real-world IT challenges in today's ever-evolving market, and we believe Gazprombank is a shining example of how this is manifest."

Gazprombank understands that in today's financial services industry, customers expect to conduct banking transactions digitally as well as in person. Gazprombank's development strategy is in effect through 2022 and involves transformation in many areas. This involves technology, focus on the rollout of the retail business, brand management, the restructuring of internal processes and even corporate culture.

Datadog launches Windows server monitoring of live traffic

DATADOG has extended Network Performance Monitoring (NPM) to Windows. Datadog NPM now monitors the performance of network communications between applications running on Windows Server and Linux, providing seamless network visibility across cloud environments, on-premises data centers, and operating systems.

Datadog Network Performance Monitoring translates distributed traffic of complex network architectures into meaningful application dependencies, so that customers can spot latencies or inefficiencies that negatively contribute to application performance, infrastructure load, and network-related costs. With this enhanced functionality, organizations can monitor their entire network across varying operating systems, providing complete visibility.

"At Datadog, we are pushing the boundaries of what it means to holistically monitor Windows Server

workloads by analyzing every aspect of their health, from infrastructure, application, network through to security," said Ilan Rabinovitch, Vice President, Product and Community, Datadog. "With this latest development, we're excited to create new opportunities for all Windows Server customers to isolate the root cause of their app issues, whether they be upstream code errors, heavy network traffic, or regional outages."

"Assessing the performance of crucial application traffic in our Windows environment used to be very difficult," said Alex Kanevsky, Lead Architect at Generali Global Assistance. "With Datadog Network Performance Monitoring, we can quickly determine if our network is at fault for slow traffic or low connectivity before our applications are affected, so that insuring travel is a seamless experience for our customers."

According to Gartner, Inc., "a well-performing network is critically important

to digital business projects; however, the level of network agility and visibility often fails to meet the requirements of these initiatives."

Datadog NPM enables monitoring of distributed traffic across on-premises and cloud environments, so organizations are able to:

- **Spot cost and performance**

bottlenecks: identify unexpected or costly communication between services and cloud regions to quickly detect where network connectivity and latency issues are concentrated.

- **Isolate the root cause:** determine when application and infrastructure issues are the root cause of faulty dependencies, misconfigured connection pooling, or cloud provider outages.

- **Visibility for every engineer:** visualize connection data at the application layer, so it can be analyzed and understood by network, application, and Reliability Engineers alike.

Park 'N Fly speeds up development cycles by 29% with Dynatrace

PARK 'N FLY, the leading off-airport parking provider, is using Dynatrace® to accelerate development cycles and deliver exceptional digital experiences across mobile, web, and airport kiosks.

Travel in the United States fell by more than 60% last year. To succeed in this challenging and competitive industry, Park 'N Fly increased its investments in digital innovation, which included migrating to a cloud-native architecture based on Kubernetes and Microsoft Azure.

Automatic and intelligent observability delivered by Dynatrace has enabled the company's digital team to tame the complexity of this constantly changing IT environment and innovate faster.

"Working without Dynatrace is like going cave diving without a light – we're just not able to see what's around us, and can't map our technologies," said Ken Schirrmacher, Senior Director of Information Technology at Park 'N Fly. "With Dynatrace, we've cut through the darkness to make sense of our complex environment and gain rich, real-time insights into how our efforts to enhance digital performance are impacting business outcomes."

Park 'N Fly's team has been able to develop new digital services faster and with more confidence due to Dynatrace's precise AI-powered insights and continuous automation. By automatically identifying anomalies that have true business impact, and alerting when problems occur, Dynatrace has enabled the team to resolve issues before customers are ever impacted.

In addition, Dynatrace's automation dramatically reduces manual-intensive tasks, through continuous discovery and automatic configuration and instrumentation. As a result, Park 'N Fly has been able to accelerate development cycles by 29% on average, including a 43% improvement during its busiest two-week sprint.

"Dynatrace allows us to innovate much faster, which has been instrumental in the launch of major new customer-facing features for our mobile app, particularly our Come Get Me and Shuttle Tracker features," continued Schirrmacher. "Dynatrace's AI capabilities have given our digital teams the confidence to try new things, knowing they can quickly assess the business impact of their changes and resolve any issues, instead of spending days or weeks trying to understand the outcomes of their efforts."

The ROI of data-driven development: improving how teams work

Every year, companies spend millions on digital solutions and platforms, hoping to accelerate their transformations. Development teams are at the heart of delivering that value and innovation, and the DevOps practices software teams have begun to adopt are critical to helping them optimize their software delivery value streams.

BY ANDREW DAVIS, SNR DIRECTOR, RESEARCH & INNOVATION, COPADO



VALUE STREAMS illustrate the way teams work together to deliver value, but without visibility into the delivery pipeline, companies are leaving millions in ROI for their transformations on the table.

But where does the ROI of digital transformation come from? How does it get realized?

Put simply, ROI comes from improving the way teams work. While the work development teams produce

is clearly important, how that work is accomplished typically holds the key for opportunities to drive business value faster.

When teams improve the way they work, they gain small, yet impactful, efficiencies continuously. Improvements to the software delivery value stream help reach revenue goals faster, allowing the business to realize ROI sooner and increase its ROI potential. In a recent study of 12 Salesforce customers, analyst



Value streams represent the series of steps an organization uses to provide a continuous flow of value to a customer. The steps of that process are typically some version of planning, building, testing, deploying, and monitoring. All these steps are needed to maximize value delivery in a systematic way

firm IDC found that two of the four primary value drivers of the Salesforce platform were increasing the productivity of the application development team and optimizing the use of IT staff time and IT infrastructure. So, how can teams strategically improve the ways in which they work together?

The Value of Visibility

To unlock innovation at scale in an orchestrated fashion, development teams need to gain visibility into the value stream itself. Teams cannot improve without clarity. Gartner's research has revealed that companies' inability to measure and optimize the value of application development is directly tied to their lack of visibility into the flow of work. While they may "know" how to work together, teams are often siloed, without visibility into processes, flow, and what is working versus what is breaking down.

Value streams represent the series of steps an organization uses to provide a continuous flow of value to a customer. The steps of that process are typically some version of planning, building, testing, deploying, and monitoring. All these steps are needed to maximize value delivery in a systematic way. But, if the steps are not explicitly defined, visualized, and measured, it is difficult to know what a team is doing or how it could improve at any stage.

Most teams who have adopted DevOps practices are unsure of how to take the next step to begin to measure and improve their organization's processes. While their teams' actions throughout the development lifecycle produce all the data they need to do so, data-driven development can only happen when development data is surfaced in a clear way.

Few companies give their teams the tools to have access or visibility into their data and even fewer can see the data organized in such a way that leads to insightful decision-making. This is the critical gap organizations need to fill to maximize potential ROI from digital transformation based in DevOps practices.

Performance data visibility enables and expedites data-driven decisions about future investments in the product, brings to light opportunities for business process re-engineering, and improves delivery velocity and quality while mitigating risk factors.

Insightful decision-making based on data

A single-source-of-DevOps truth can help give teams the insights they need to measure and improve DevOps performance and throughput. Many people talk about monitoring the development lifecycle as the last step in DevOps maturity, but the reality is that it is one of the linchpins to digital transformation.

Monitoring should be applied to every step of the DevOps lifecycle to gain visibility into processes, gather data on what is working and what is not, and adjust accordingly. It is integral to every step of the DevOps process.

Historically, a value stream map has been an essential lean tool for an organization wanting to plan, implement, and improve while on its lean journey. Value stream mapping helps users create a solid implementation plan that maximizes their available resources and ensures that materials and time are used efficiently.

Applied to DevOps, this maps the value stream of a Salesforce delivery process to help unlock opportunities to drive more value. Businesses need to map out each phase of the delivery process: plan, build, test, deliver, and monitor. Each metric represents an opportunity for incremental improvements in time to value and ROI.

DevOps Analytics

Metrics such as change failure rate, mean time to restore, lead time, and deployment frequency - that are tracked on every stage of the process serve as red flags. Changes in these numbers over time help indicate stages of the development process to further dig into. This information is crucial for not only having a holistic view of where components of the roadmap stand in the development lifecycle, but for reporting on value creation of the development team and digging in to optimize the system.

Realistically, a business cannot make improvements across all indicators at once. Just as agile is rooted in simplifying work to drive value faster, DevOps value stream optimization benefits from starting with smaller pieces. Increases or decreases in these metrics over time are just symptoms of underlying causes. Target



one at a time for improvement, and often, multiple areas will benefit.

In this way, Value Stream Maps illustrate process breakdowns such as bottlenecks that prevent the team from delivering value in a reasonable time frame; excess waiting time; or when, where, and how far back in the process work is being sent back. These are all opportunities for business process improvement. For example.

If lead time increases, a business can assess how the work is planned, if the work scoped is too big and should be broken down into smaller, independent pieces?

Alternatively, leaders may need to analyse the build phase. Are developers context-switching throughout the day, as they work on multiple user stories at once? There are then assessments of work quality - do tests or QA checks routinely fail, causing work to be sent back to development? Do deployments routinely fail? Is testing / QA being bypassed?

If deployments become less frequent, leaders should ask if complexity of the planned work increased, and then consider re-scoping projects to release in smaller batches. If teams are still manually testing, could it be automated? Then review the tests to make sure they are written in a way that provides more coverage.

Achieving understanding

These examples show that visibility is not just about tracking metrics to see incremental improvement in them over time. It is about understanding where to look to uncover which processes are working and which are not in a development lifecycle. This, in turn, helps understand where to apply or refine DevOps practices to optimize a team's work output and the way they work, to deliver value to customers faster. This business process management is a virtuous feedback loop that resembles a DevOps mindset: watch the data to understand which process to optimize, adjust the process, watch how the data changes, rinse, and repeat.

As teams practice this process of end-to-end optimization over time, they become more efficient, building, and delivering with higher quality at faster speeds. This is the core of digital transformation: businesses cannot transform (or continue transforming) without refining processes.

In 2018, nearly \$1.3 trillion was spent on digital transformation globally. Of that, more than \$900 billion is estimated to have gone to waste. The need for DevOps, agile development, and value stream visibility is not a nice-to-have, it is a business necessity for digital transformation winners.

Further reading

1. Carvalho Larry, Marden Matthew, Arora Ustav. The ROI of Building Apps on Salesforce. IDC. 2016.
2. Tabrizi, Behnam Lam, Ed Girard, Kirk Irvin, Vernon. Digital Transformation Is Not About Technology. Harvard Business Review. 2019.

Using the greater data ecosystem to drive great decision making

As the fallout from the COVID-19 pandemic continues to disrupt the majority of industries, its impact on supply chains has been nothing short of seismic. As teams continue to face increasing pressure to make the right decisions at the right time - squeezing every last drop of insight and information out of vast lakes of data is now more important than ever.

BY WILL DUTTON, DIRECTOR OF MANUFACTURING, PEAK

THE PHRASE 'data is the new oil' has framed a large amount of discourse in the twenty-first century. The statement, although contestable, does beg the question that should always follow: exactly what data are we talking about? These tricky times call for a new approach to data-driven decision making. There's now a real need for supply chains to focus on the greater data ecosystem, accessing wider sources of data and utilising it to its fullest capacity. While making effective decisions based on data from current systems, or by joining up a few previously-siloed sources across the organisation is becoming easier than ever – there's potential to go even further than this. The more data there is to play with, the more informed supply chain decisions will be. Here are four data sources that can help accelerate smart supply chain decisions:


1. Linking the supply chain with customer systems

The more systems that talk to each other, the better. Linking data from supply chain systems with customer systems, including behaviour data points, can help understand the pain points that arise. For instance, this could be the customer's ERP system or even the logistics systems between the business and the customer. Taking consumer-packaged goods businesses and manufacturers as an example, with a better handle on Electronic Point of Sale (EPOS) and any other sell-out data from customers' systems, the business can better predict what demand is going to be like, and better understand their stock levels in order to help anticipate their own. Factoring into account things like receipts data, what baskets are shoppers generally buying together, and how can this help better anticipate the groups of products that are going to sell together. This closer relationship with customers' systems allows the business to better serve them, increasing efficiency and anticipating demand fluctuations. Inherently it's all about creating

more competitive supply chains which are more cost-effective, with better service levels and a more accurate view of demand.

2. Supplier data for efficiency

By leveraging data points from suppliers' systems, businesses can plan ahead in the most efficient way and execute an effective just-in-time (JIT) inventory management strategy, holding minimal assets to save cash and space whilst still fulfilling customer demand. By employing this methodology, businesses are able to understand when a supplier is going to deliver, to what location, and anticipate the arrival of goods and raw materials whilst also better understanding the working capital implications.





3. Using environmental data

Don't underestimate the power hidden away in external, non-industry related data sources and the impact it can have on supply chain decision making. Think about the ways a business can utilise, let's say, macroeconomic data to understand what could be driving issues connected to supply and demand. Yes, we immediately think of things like GDP, or maybe even exchange rates, but there is now a plethora of data out there, that may be more industry and company-specific, that helps predict demand or implications for business performance. In recent months, appropriate data feeds impacting the supply chain could be an increase in Covid-19 cases near a supplier, hampering their ability to operate as normal. Connecting these data points up to technology such as Artificial Intelligence (AI) could help understand the impact of these incidents with supply performance, and create accurate forecasts on the trends.

4. Sharing data across the network with co-opetition

For many, the rule of thumb is not giving the game away to competitors, so this may seem a little pie-in-the-sky for many businesses at first. However, the benefits of sharing data with the industry and accessing competitor data sources can be enormous. The data of those providing similar products is at first harmless – but using it in the right way, to make intelligent decisions, will allow the business to gain a unique view of what is happening across the rest of the market. This ultimately leads to a better understanding of wider trends and the ability to make smarter decisions. With a mutually beneficial relationship with the wider network, a business can understand supply issues, and work with competitors

or neutral parties to deliver better products and services to customers creating a form of 'co-opetition.'

Accessing the ecosystem requires digital transformation

Tapping into the greater data ecosystem and utilising it in decision making will be essential for supply chain teams to run smooth operations in disruptive climates. However, to truly unlock the potential this offers, a central AI system is needed.

In the same way that business functions have their own systems of record, the ability to power decision making based on a wide range of data sources hinges on the introduction of a new, centralised enterprise AI system. Using AI gives teams the ability to leverage unlimited data points at scale and speed. Utilising AI in this manner, to make decisions that are both smarter and faster to supercharge teams. At Peak, we call this Decision Intelligence (DI).

Decision Intelligence results in being able to connect the dots between data points with AI, to prescribe recommendations and actions to make more informed commercial decisions across the entire supply chain.

By feeding external data from the points above into both demand and supply planning systems, leveraging it with AI, enterprises can optimise that connection between these two core areas. Not only does it allow a better sense of demand with a higher degree of accuracy, but also enables a better understanding of how supplier and operations constraints are affecting supply – automatically making micro-adjustments to optimise the way demand is being fulfilled.