DW DIGITALISATION WORLD

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

ISSUE VI 2021

digitalisationworld.com

# Kubernetes and containers:

## Powering tomorrow's applications


AIOps | Apps + DevOps | Artificial Intelligence | Big Data + Analytics
Cloud + MS | DC Facilities + Colo Digital Business | IT Management + Service
Networks + Telecoms | Open Source | Security + Compliance | Storage + Servers

SEAGATE

DCS DATACENTRE SOLUTIONS

# ROUNDTABLE

## WITH ERIK SALO AND ANDY PALMER

---

## 10:00 – 11:00 AM (+1 GMT)

## REGISTER NOW

**digitalisationworld.com/s/seagate**

In this exclusive roundtable, tech celeb Erik Salo and Seagate enterprise solutions specialist Andy Palmer discuss the latest intelligent storage solutions to optimize data centers.

They go into depth on Seagate Corvault, Salo's invention, and reveal the key trigger that led to it. Afterwards, they will be addressing questions.

**Erik Salo**
**Sr. Director**
Worldwide Systems Distribution

**Andy Palmer**
**UK Channel Lead**
Enterprise Data Solutions

## A New Category of Intelligent Storage

Exos CORVAULT is a high-performance, self-healing block storage system that delivers multi-petabyte capacity, five-nines availability, and hyperscale efficiencies for data center and macro edge environments.

# Editor's View

By Phil Alsop

## What's the future for SustainabilIT?

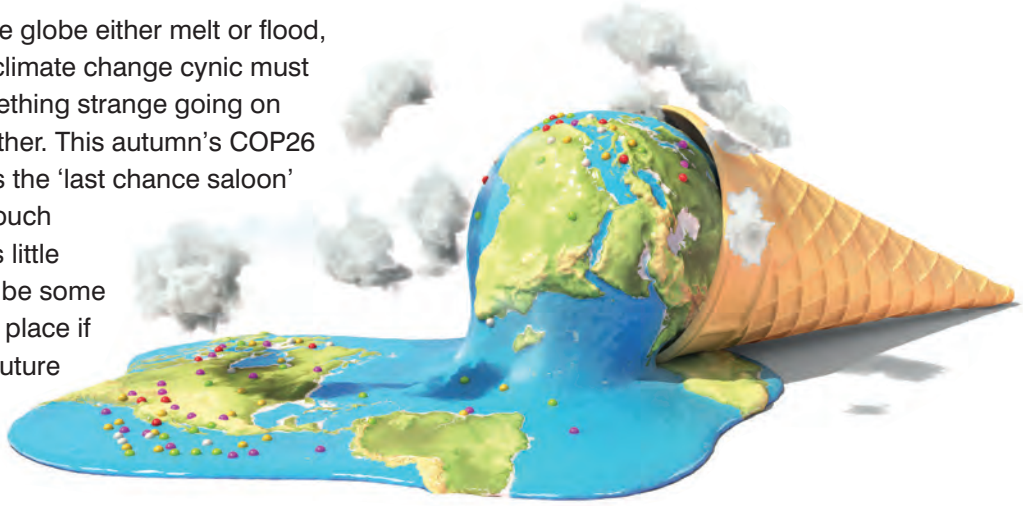

AS VARIOUS PARTS of the globe either melt or flood, even the most hardened climate change cynic must concede that there's something strange going on when it comes to the weather. This autumn's COP26 summit has been billed as the 'last chance saloon' and, even allowing for a touch of hyperbole, there seems little doubt that there needs to be some radical, urgent plan put in place if we are not to hand on to future generations a dying planet.

The IT industry is uniquely placed, in that it is both a major cause of the problem, thanks to the vast amount of energy it consumes, and a major part of any solution, as digital business is already eliminating many energy-intensive and polluting activities (ie video calls as opposed to flying/driving to a physical meeting).

It is inevitable that, at some stage in the near future, the sector will come under intense attack for its energy consumption levels. However, properly prepared, IT companies will be able to demonstrate that they are actually making a major, positive contribution to the various Net Zero targets out there. Tedious though it might be, this preparation requires a major tranche of work, which compares analogue and digital activities, demonstrating that IT solutions are much kinder to the planet than the 'old-fashioned' activities they have replaced.

Such scientific lifecycle analysis will stand up to the closest scrutiny of consumers who are more than capable of seeing through the shallow, rather meaningless green claims that are currently thrown around with abandon.

The good news is that more and more businesses are beginning to understand that environmental policies are no longer just a tick box exercise, but do actually need to demonstrate real change, and a roadmap to further, continuous improvements.

The less good news is that there is still much work to be done to elevate the science of green business to a level where it needs to be. After all, it's not difficult to imagine a future where consumers judge potential suppliers on their carbon footprint just as readily as they currently evaluate them for the digital experience on offer.

It would be a great shame if companies who have survived or thrived in the digital revolution found themselves victims of the coming green revolution.

# Contents

ISSUE VI 2021





## LOW / NO CODE



## CLOUD



## IT MONITORING



## STORAGE





The data centre trade association

## DCA News



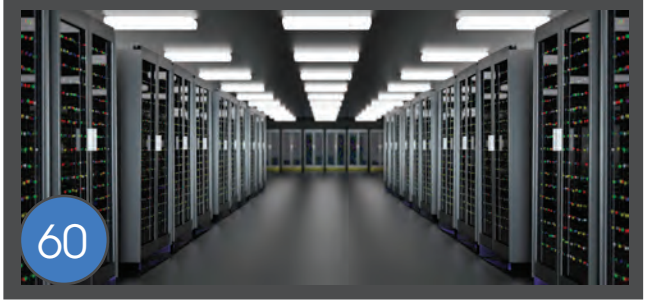



DIGITALISATION WORLD

**Editor**
Philip Alsop +44 (0)7786 084559 philip.alsop@angelbc.com

**Sales Manager**
Peter Davies +44 (0)2476 718970 peter.davies@angelbc.com

**Account 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

Angel BUSINESS COMMUNICATIONS

recycle When you have finished with this magazine please recycle it.

Digitalisation World is published 10 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.



The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. ISSN 2396-9016 (Online)

# Pandemic exposes significant disconnect between IT and the wider business

A LONGSTANDING disconnect between IT teams and the wider business is preventing organisations from adopting new technologies and jeopardising their long-term response to the pandemic, research from Insight has revealed.

Despite the importance of IT to delivering on organisations' strategic objectives, almost three quarters (72%) treat IT as a utility rather than a business enabler, with just 22% giving IT a seat on the board. This has direct consequences for enterprises: 55% of organisations are failing to take advantage of new technologies because they aren't listening to IT.

The pandemic has shone a light on this disconnect. 83% of senior IT decision makers believe ways of working have been permanently transformed. Yet across the wider business, at least 61% of organisations are reluctant to invest in projects that could improve the employee experience or optimise the business because they believe things will eventually return to a pre-COVID-19 "normal".

Without addressing this, there is a real risk that enterprises will invest in projects without believing in their goals; fail to understand the impact of new ways of working on employees; or base strategies on incorrect assumptions. As a result, they will almost certainly see investments wasted, projects failing and competitors taking advantage.

"The pandemic has brought about permanent changes to the way many of us live and work. We are not going to see a return to the status quo, and it's absolutely imperative that organisations adapt," said Emma de Sousa, President, EMEA at Insight.

"There's already a huge risk associated with making investments in the wrong place. But an incorrect investment at this moment in time could prove more damaging than ever before, leaving the enterprise unequipped for new ways of working and doing business. The gap between IT teams and the wider business must be closed as an urgent priority: businesses have to engage with IT on a more strategic basis, and measure it against businesses objectives."


Other findings from the research included:

- **IT teams must measure business impact:** 81% of IT departments have freedom to invest in the skills they need, and 82% are engaged to support business projects. Yet 59% aren't measured against business KPIs
- **Skills gaps must be overcome for new ways of working to succeed:** 57% of organisations say they need to invest more in the skills and technology needed to support a remote workforce, and 60% need to invest more in the skills and technology needed to optimise the business
- **Disconnect between IT and the business puts projects at risk:** 67% of organisations are working on projects designed to improve the employee experience, and 55% on projects to optimise the business. However, the belief of the wider business that things will return to "normal" means that many of these projects do not have the full support of the business, and so are more likely to fail
- **The costs of not engaging IT:** The failure to engage with and listen to IT, coupled with the clear disconnect between IT and the wider business, has almost certainly contributed to enterprises' £3.81m / €4.19m spend from 2018-2020 on projects that either did not provide the full benefits, or failed

"The way IT is perceived and used within businesses has to fundamentally change," de Sousa continued. "Having IT at arm's length from the board is simply not good enough: it must be given a seat at the top table. Without this, businesses risk falling behind at a time when digital technology is driving change across all sectors. IT must be put front and centre, driving organisational change and being made directly accountable for doing so.

If organisations give IT a voice on the board to drive strategy; let IT use that voice to support innovation; consult IT on what approaches will best meet the business's objectives; and trust IT to perform against business KPIs, they will be positioned to face the challenges of 2021 and beyond."

# Meeting the demands of a digital-first future

TO HELP mitigate the rise of emissions from digital devices, the new Digital Economy and Climate Impact report recommends continuing efforts to achieve efficiencies on IT and energy sides, at component and system levels.

Schneider Electric has released a research report to foster an understanding of how digitised and smart applications will be powered in the future. The report titled Digital Economy and Climate Impact predicts IT-sector related electricity demand is expected to increase by nearly 50% by 2030.

Yet, as the electricity system decarbonizes, emissions would not increase by more than 26% by that time. To help mitigate this rise in emissions, the Schneider Electric™ Sustainability Research Institute recommends continued efforts in achieving efficiencies on the IT and energy sides at both the component and system levels.

Released at an exclusive media event presented virtually from Schneider Electric's Boston Hub, the report highlights how the rise of edge computing requires a specific focus as these systems are expected to be less efficient than hyperscale data centres from a PUE standpoint.

"When the world locked down it also logged on and internet traffic soared," said Pankaj Sharma, EVP, Secure Power, Schneider Electric. "It's misleading to assume that digital activity will inevitably result in a deeply problematic increase in CO2 emissions. The analysis from the Schneider Electric Sustainability Institute puts to rest many of the worst-case scenario claims predicting IT-related electricity use will double every five years. That said, as an industry we must remain vigilant in finding new sources of sustainability gains while ensuring resiliency as digital keeps life moving forward."

In addition to releasing the research report, Schneider Electric also announced updates to its EcoStruxure™ IT data centre infrastructure management software and Galaxy™ VL 3-phase uninterruptible power supply (UPS). All introductions are designed to advance the industry forward in meeting sustainability goals while increasing resiliency of IT and data centre infrastructure.

EcoStruxure IT software updates reduce complexity in managing hybrid data centre and edge IT environments Increasing demands on digital consumption, which are explored in the new research report, create a more complex hybrid environment inclusive of enterprise, cloud and edge data centres.

To address the unique management challenges of a hybrid IT environment, Schneider Electric has announced updates to its EcoStruxure IT software to increase efficiency and resiliency, including:

- **Increased remote management capabilities:**
  New granular remote device configuration features enable users to change configurations on one or more devices – including the new Galaxy VL UPS unit – from one centralised platform with EcoStruxure IT Expert. This update, combined with previously released software insights on device security health, enables the user to identify faulty devices or configurations and address them in a matter of clicks, keeping their hybrid IT environment secure.
- **Improved environmental monitoring:**
  Environmental monitoring systems ensure users have eyes and ears on data centre and IT deployments from anywhere, anytime. With this update, users can push mass configurations remotely for NetBotz cameras 750 and 755 quickly and efficiently increasing security across the critical infrastructure.
- **Enhanced remote capacity modeling and planning:**
  With EcoStruxure IT Advisor's new capabilities, users can remotely compare an unlimited number of racks and easily identify available capacity, view what assets are deployed and their dependencies.

Redesigned Galaxy Lithium-ion battery solution enables greater space savings, faster recharge and installation and enhanced safety.

The newly released Galaxy VL, the most compact of its class, modular and scalable 3-phase UPS in the 200 - 500 kW range with efficiency levels up to 99 per cent, now features redesigned Galaxy Lithium-ion battery cabinets, providing a sustainable, high-density and innovative energy storage solution for data centres, industrial processes, and critical infrastructure. The exclusive cabinets are compatible across the full Galaxy V Series.

A Green Premium™ offer, this UL9540A-compliant battery solution reduces battery footprint and weight by up to 70 per cent, allowing more effective use of space.

The new cabinets enable two to three times faster recharge than VRLA solutions as well as faster installation and enhanced system availability with patented redundant self-powered internal power supplies.

Lithium-ion batteries reduce total cost of ownership by doubling battery life, lowering installation and maintenance costs, plus reducing cooling needs, as they operate at higher temperatures than VRLA. The included real-time battery management system improves battery system visibility, predictability, and manageability. The modular, touch-safe design simplifies maintenance and increases operator safety.

# Putting faith in a flexible future

75 PER CENT OF WORKERS believe flexible work should be an essential part of how people work going forward.

Fuze, the leading cloud-based communications provider for the modern global enterprise, has published the results of a new global study of 8,800 workers, which reveals that as flexible work is increasingly viewed as an expectation for the working world, attitudes continue to vary widely across industries, roles and geographies.

The Fuze report, "FLEX Study: Global Findings on the Future of Flexible Work", reinforces that organisations and business leaders should not consider flexible work as a 'one-size-fits-all' benefit, but a personalised experience that should be tailored to an individual's unique work preferences, role within their team, industry practices, aligned to goals and expected outcomes, and established through a foundation of trust.

The FLEX Study was unveiled today at the 2021 Flex Summit Week, a week-long virtual event that brings together industry experts and visionaries to discuss how enterprises are embracing flexible work. Results from the study were collected from more than 8,800 frontline and office workers in the United States, United Kingdom, France and Australia, and across a number of core industries, including manufacturing, retail, professional services, financial services, and software and technology.

Among the key findings from the study, 75 per cent reported that flexible work should be an essential part of how people work. In most regions, a full-time return to the office or worksite mandate could result in destabilising job churn. In the UK, US, and Australia, approximately two-thirds of employees (67 per cent) would consider finding a new job for greater flexibility in when and how they work. In France, more than half of employees (57 per cent) would consider finding a new job for greater flexibility.

However, trust in flexible work varies widely across roles. Across all regions, about half (54 per cent) of workers think management is more trusting of remote work, while 70 per cent of senior


leaders believe management is more trusting, highlighting a gap between management's new-found trust in remote work and employee perception of that trust. The gap in trust toward remote work is also reflected by role, with office workers (73 per cent) more likely to say they are more trusting of remote work compared to frontline workers (50 per cent). This may be credited to the lack of flexible work options traditionally offered to frontline workers versus office workers. Other key insights from the report include:

- Frontline workers perceive flexible work differently than office workers. 42 per cent of frontline workers believe management is trustworthy of remote work, compared to 62 per cent of office workers. 78 per cent of frontline workers also report that their organisation requires them to be in their current working location versus only 60 per cent of office workers. However, six in 10 (63 per cent) frontline workers say they would be willing to change jobs for more flexibility.
- Remote workers feel more productive. For those working remotely full-time, 60 per cent say they are more productive at home than they were in the office. However, this feeling of productivity varies across regions with 70 per cent of Australian respondents saying that they felt more productive versus only 52 per cent of UK respondents.
- Companies must reshape their culture of video meetings to drive engagement. In every industry, every job category, and every region, fewer than 10 per cent of workers consider seeing someone's face the most critical part of an effective meeting. For workers who report being happier at home, if their day is packed with meetings, they become less engaged and productive. In fact, 59 per cent of respondents say they would like to spend less than two hours a day in video meetings.
- Companies must establish a structure for the work day to prevent burnout. Only 66 per cent of respondents say they are ensuring they take a break every day. In addition, a quarter of respondents report that they find themselves working longer hours since they started working remotely.

"Over the last year, the pandemic forced many organizations to digitally transform and embrace flexible work," said Brian Day, CEO, Fuze. "While office workers have been at the forefront of the flexible work conversation for years, other segments of the workforce are playing a critical role in the movement's evolution. This study provides organizations with the foundational data required to inform their approach to flexible work and empower employees to be more engaged and productive."

# McAfee research highlights threat trends

McAfee sees attackers shift from mass-spread campaigns to fewer, more lucrative targets.

McAfee has released its McAfee Threats Report: June 2021, examining cybercriminal activity related to malware and the evolution of cyber threats in the first quarter of 2021. The quarter saw cyber adversaries shift from low-return, mass-spread ransomware campaigns toward fewer, customised Ransomware-as-a-Service (RaaS) campaigns targeting larger, more lucrative organisations.

A proliferation in 64-bit CoinMiner applications drove the growth of cryptocurrency-generating coin mining malware by 117%. Additionally, a surge in the growth of new Mirai-based malware variants drove increases in malware targeting Internet of Things (55%) and Linux (38%) systems.

"Criminals will always evolve their techniques to combine whatever tools enable them to best maximise their monetary gains with the minimum of complication and risk," said Raj Samani, McAfee fellow and chief scientist. "We first saw them use ransomware to extract small payments from millions of individual victims. Today, we see Ransomware as a Service supporting many players in these illicit schemes holding organisations hostage and extorting massive sums for the criminals."

Each quarter, McAfee assesses the state of the cyber threat landscape based on in-depth research, investigative analysis, and threat data gathered by the McAfee Global Threat Intelligence cloud from over a billion sensors across multiple threat vectors around the world.

## Ransomware

Ransomware declined by 50% in Q1 due in part to a shift by attackers from broad campaigns attacking many targets with the same samples to campaigns attacking fewer, larger targets with unique samples. Campaigns using one type of ransomware to infect and extort payments from many victims are notoriously "noisy" in that hundreds of thousands of systems will, in time, begin to recognise and block these attacks. By allowing attackers to launch unique attacks, RaaS affiliate networks are allowing adversaries to minimise the risk

of detection by large organisations' cyber defenses and then paralyse and extort them for large ransomware payments. This shift is reflected by the decline in prominent ransomware family types from 19 in January 2021 to 9 in March 2021. Despite the high profile attacks from the DarkSide RaaS group exposed in Q2 2021, REvil was the most detected in Q1, followed by the RansomeXX, Ryuk, NetWalker, Thanos, MountLocker, WastedLocker, Conti, Maze and Babuk strains.

## Coin Miner Malware

While prominent ransomware attacks have focused attention on how criminals use ransomware to monetise their crimes with payments in cryptocurrency, a first quarter 117% surge in the spread of cryptocurrency-generating coin mining malware can be attributed to a sharp spike in 64-bit CoinMiner applications.

Rather than locking up victims' systems and holding them hostage until cryptocurrency payments are made, Coin Miner malware infects compromised systems and silently produces cryptocurrency using those systems' computing capacity for the criminals that designed and launched such campaigns. The advantage to cybercriminals is that there is zero interaction required of both the perpetrator and the victim. While the victim's computers may operate slower than usual due the coin miner's workload, victims may never become aware that their system is creating monetary value for criminals.

"The takeaway from the ransomware and coin miner trends shouldn't be that we need to restrict or even outlaw the use of cryptocurrencies," Samani continued. "If we have learned anything from the history of cybercrime, criminals counter defenders' efforts by simply improving their tools and techniques, sidestepping government restrictions, and always being steps ahead of defenders in doing so. If there are efforts to restrict cryptocurrencies, perpetrators will develop new methods to monetise their crimes, and they only need to be a couple steps ahead of governments to continue to profit."

## Threats & Victims

### Overall Malware Threats

The first quarter of 2021 saw the volume of new malware threats average 688 threats per minute, an increase of 40 threats per minute over Q4 2020.

### IoT & Linux Devices

A variety of new Mirai malware variants drove increases on the Internet of Things (IoT) and Linux malware categories in Q1. The Moobot family (a Mirai variant) was observed to be mass-spread and accounted for multiple Mirai variants. These variants all exploit vulnerabilities in IoT devices like DVRs, webcams and internet routers. Once exploited, the malware is hidden on the system, downloads later stages of the malware and connects with the command-and-control server (C2). When the compromised IoT devices are connected to their botnet, they can be commandeered to participate in DDoS attacks.

### Industry Sectors

McAfee tracked a 54% increase in publicly reported cyber incidents targeting the technology sector during the first quarter of 2021. The Education and Financial/Insurance sectors followed with 46% and 41% increases respectively, whereas reported incidents in Wholesale/Retail and Public Sector declined by 76% and 39% respectively.

### Regions

These incidents surged in 54% in Asia and 43% in Europe, but declined 13% in North America. While reported incidents actually declined 14% in the United States, these incidents grew 84% in France and 19% in the United Kingdom.

# Citrix research reveals employee experience is about to break

IN RESPONSE to the global pandemic, companies have given their workforce the tools they need to work remote. And research shows it has increased their engagement and productivity. But these gains are on the brink of being wiped out. According to a new study from Citrix Systems, employees feel they've been given too many tools and not enough efficient ways to execute. And it's hindering their ability to get things done.

"People are working the same or more hours, but they're accomplishing less because technology is getting in their way," said Tim Minahan, Executive Vice President of Business Strategy, Citrix. "As companies organise around new, hybrid work models, they need to rethink the role of technology and how they apply it across their organisations so that employees, rather than being frustrated, are empowered to succeed." To help them do this, Citrix undertook Work Your Way, a survey of 1,000 IT decision makers and 2,000 workers across the United States conducted by OnePoll, that revealed a few significant trends:

**App sprawl is out of control**

The number of tools employees are required to use to do their jobs has significantly increased, as has the complexity they are creating in the workplace. As uncovered by Work Your Way:

- 64 percent of workers are using more communication and collaboration tools than they were prior to the pandemic, and
- 71 percent say they have made work more complex

"Employees are frustrated, and to keep them engaged and performing at their best, companies need to eliminate the friction and noise from work and deliver technology that adapts to their workstyles rather than forcing them to learn new ways of doing things," Minahan said.

**A new digital divide is emerging**

But workstyles have fundamentally changed. "People are not going back to working the way they did," Minahan said. Work Your Way confirms this. Nearly 90 percent of respondents to the survey say they want the flexibility to continue to work at home and in the office post pandemic. "Regardless of their physical location, employees need to be empowered with tools that provide a consistent, secure and reliable experience and allow them to work the way they work best," Minahan said.

**Digital workspaces are the future of work**

Savvy organisations recognise this and see digital workspaces as a way to deliver it.

With digital workspaces, companies can:

- **Unify work:** Whether at home, on plane or in an office, employees have consistent and reliable access to all the resources they need to be productive across any work channel, device or location
- **Secure work:** Contextual access and app security ensure applications and information remain secure – no matter where work happens.
- **Simplify work:** Intelligence capabilities like machine learning, virtual assistants and simplified workflows personalise, guide, and automate the work experience so employees can work free from noise and perform at their best.

Almost 90 percent of participants in Work Your Way say their companies use digital workspace software platforms to facilitate hybrid/distributed working. And they are delivering results.

- 72 percent of employees say they have improved productivity, and
- 77 percent indicate they have aided collaboration "In creating a layer between employees and the technology that frustrates them, companies can empower them to efficiently engage with the apps they need to execute work and achieve their goals," Minahan said.

# Current data architectures in danger of crumbling

NEW MARKET RESEARCH identifies demand for real-time model training and inferencing; Highlights major challenges with accuracy, latency, and reliability in current architectures.

As companies look to expand their use of artificial intelligence (AI) and machine learning (ML) to keep up with the demands of their customers, they are facing hurdles getting these projects to production and ultimately deliver the desired results to their bottom line.

In fact, 88% of AI/ML decision-makers expect the use cases that require these technologies to increase in the next one to two years, according to a commissioned study conducted by Forrester Consulting on behalf of Redis Labs. The research looked at the challenges keeping decision-makers from their desired transformation when deploying ML to create AI applications. The study revealed that companies are developing increasingly more models based on real-time data. Still, more than 40% of respondents believe their current data architectures won't meet their future model inferencing requirements.

Most decision-makers (64%) say their firms are developing between 20% to 39% of their models on real-time data from data streams and connected devices. As teams develop more models on real-time data, the need for accuracy and scalability is becoming increasingly critical. Significantly, thirty-eight percent of leaders are developing roughly a third of models on the real-time spectrum. Other key findings include:

- Nearly half of decision-makers cite reliability (48%) and performance (44%) as their top challenges for getting models deployed with their current databases. Equally concerning was the revelation that 41% of respondents believe their databases cannot meet the necessary data security and compliance requirements.
- To achieve the benefits that AI/ML promise, survey respondents said that locating models in an in-memory database would solve key hurdles

currently standing in their way. According to the survey, the benefits architecturally would allow firms to prepare data more efficiently (49%), improve analytics efficiency (46%), and keep data safer (46%).

As Forrester Consulting concludes, "AI powered by ML models mustn't slow down applications by necessitating a network hop to a service and/or microservice for an application to use an ML model and/or get reference data. Most applications, especially transactional applications, can't afford those precious milliseconds while meeting service-level agreements (SLAs)."

"Companies are embracing AI/ML to deliver more value for their mission-critical applications, yet need a modern AI/ML infrastructure to support real-time serving and continuous training. There are still gaps that impede companies from making existing applications smarter and delivering new applications," said Taimur Rashid, Chief Business Development Officer at Redis Labs.

"Customers realize this, and the simplicity and versatility of Redis as an in-memory database is enabling them to implement Redis as an online feature store and inferencing engine for low-latency and real-time serving."

"Fabric was established to help brands migrate from legacy to modern, digital commerce systems," said Umer Sadiq, CTO of Fabric. "In order to offer businesses the best-in-class technologies that enhance and improve customer experiences, we have crafted and continue to deliver applications that rely on Redis Labs' real-time data platform hosted on AWS to ensure real-time feature serving to customers, thus maintaining exceptional user satisfaction.

Additionally, by combining the power of Amazon SageMaker and Redis Enterprise to bolster the efficiency of our market-leading recommender systems, we guarantee low-latency and high reliability for each individual customer interaction."

"The Room's mission is to connect top talent from around the world to meaningful opportunities, and at the core of the technology challenge is a mathematically difficult entity-matching problem," said Peter Swaniker, CTO of The Room. "To address this complexity, we have architected a joint solution using Scribble Data's Enrich Feature Store and Redis Labs' real-time data platform to provide the overall framework for The Room's Intelligence Platform, which is responsible for entity matching. Using Redis' high-performance key retrieval based on nearest neighbor vector lookup, the team was able to achieve a 15x+ improvement in the core similarity computation loop without any memory overhead."

"The use of machine learning (ML) algorithms in simulations continues to grow to improve scientific research with efficiency and accuracy," said Benjamin Robbins, Director AI & Advanced Productivity, Hewlett Packard Enterprise. "By leveraging Redis and RedisAI in SmartSim, our new open source AI framework which advances simulations that run on supercomputers, users can exchange data between existing simulations and an in-memory database, while the simulation is running. The ease of data exchange helps unlock new machine learning opportunities, such as online inference, online learning, online analysis, reinforcement learning, computational steering, and interactive visualization that can further improve accuracy in simulations and accelerate scientific discovery."

# Organisations to spend almost $656 billion on Future of Work technologies

All aspects of how people and organizations work is evolving, enabled by 3rd Platform technologies like cloud and mobile computing and accelerated by the COVID-19 pandemic

FUTURE OF WORK (FoW) is a fundamental shift in the work model to one that fosters human-machine collaboration, enables new skills and worker experiences, and supports a work environment un-bounded by time or physical space. A new forecast from the International Data Corporation (IDC) Worldwide Future of Work Spending Guide estimates FoW spending will be nearly $656 billion this year, an increase of 17.4% over 2020.

"Traditional work models do not provide the agility, scalability, and resilience required by the future enterprise. This was, of course, highlighted by the ongoing health crisis. To drive growth and competitive differentiation, organizations will invest in technologies and services that power automation, human-machine collaboration, new organizational structures and leadership styles, dynamic learning opportunities, a reimagined workplace, and a digital work environment that is not bounded by time or physical place," said Holly Muscolino, research vice president, Content Strategies and the Future of Work.

To facilitate the transition to the new workplace and an evolving workforce, organizations are investing in a wide range of technologies and services. The largest area of investment in 2021 will be hardware, where companies are expected to purchase $228 billion in endpoint devices, enterprise hardware, infrastructure as a service (IaaS), and robotics and drones. Services, including business, IT, and connectivity services, will be the second-largest area of spending at more than $123 billion. Software will see the fastest spending growth with a compound annual growth rate (CAGR) of 21.3% over the 2020-2024 forecast period.

This includes investments in enterprise applications, content and collaboration, analytics and artificial intelligence, human resources applications, security, and software development and deployment.

"Emerging technologies like artificial intelligence, the Internet of Things, and augmented/virtual reality are changing how work is getting done across all industries and across the world. Seeking automated decision support and virtual collaborative approaches, discrete and process manufacturing, the two largest spenders on Future of Work technology over the forecast period, are investing in key use cases like collaborative robotics, operational performance management, and 3D and digital product design and review for improved cost control and higher process efficiency," said Eileen Smith, program vice president, Customer Insights and Analysis.

Together, discrete and process manufacturing will account for just over one third of all Future of Work spending this year. Professional services, retail, and banking will be the next three industries in terms of FoW spend in 2021. The construction industry will see the fastest growth in FoW spending over the forecast period with a five-year CAGR of 23.7%. Media and retail will follow closely with CAGRs of 19.5% and 19.3% respectively.

## Share the image

The FoW use cases that will benefit from the most spending in 2021 include collaborative robotics, operational performance management, and automated customer management. The use cases that are expected to see the fastest spending growth over the 2020-2024 forecast period are adaptive skill development, interconnected collaborative workspaces, and advanced project management. "IDC forecasts investment in technologies supporting Future of Work initiatives to exceed $1 trillion worldwide by 2024 with a robust 17% CAGR over the five-year forecast period.

All aspects of how people and organizations work is evolving, enabled by 3rd Platform technologies and accelerated by the pandemic. Indeed 3rd Platform hardware, such as IoT devices, robots and drones, and IaaS, are more than one-third of the total spend, demonstrating the growing importance of the technologies enabling the reimagined workplace," said Karen Massey, research manager, Customer Insights & Analysis.

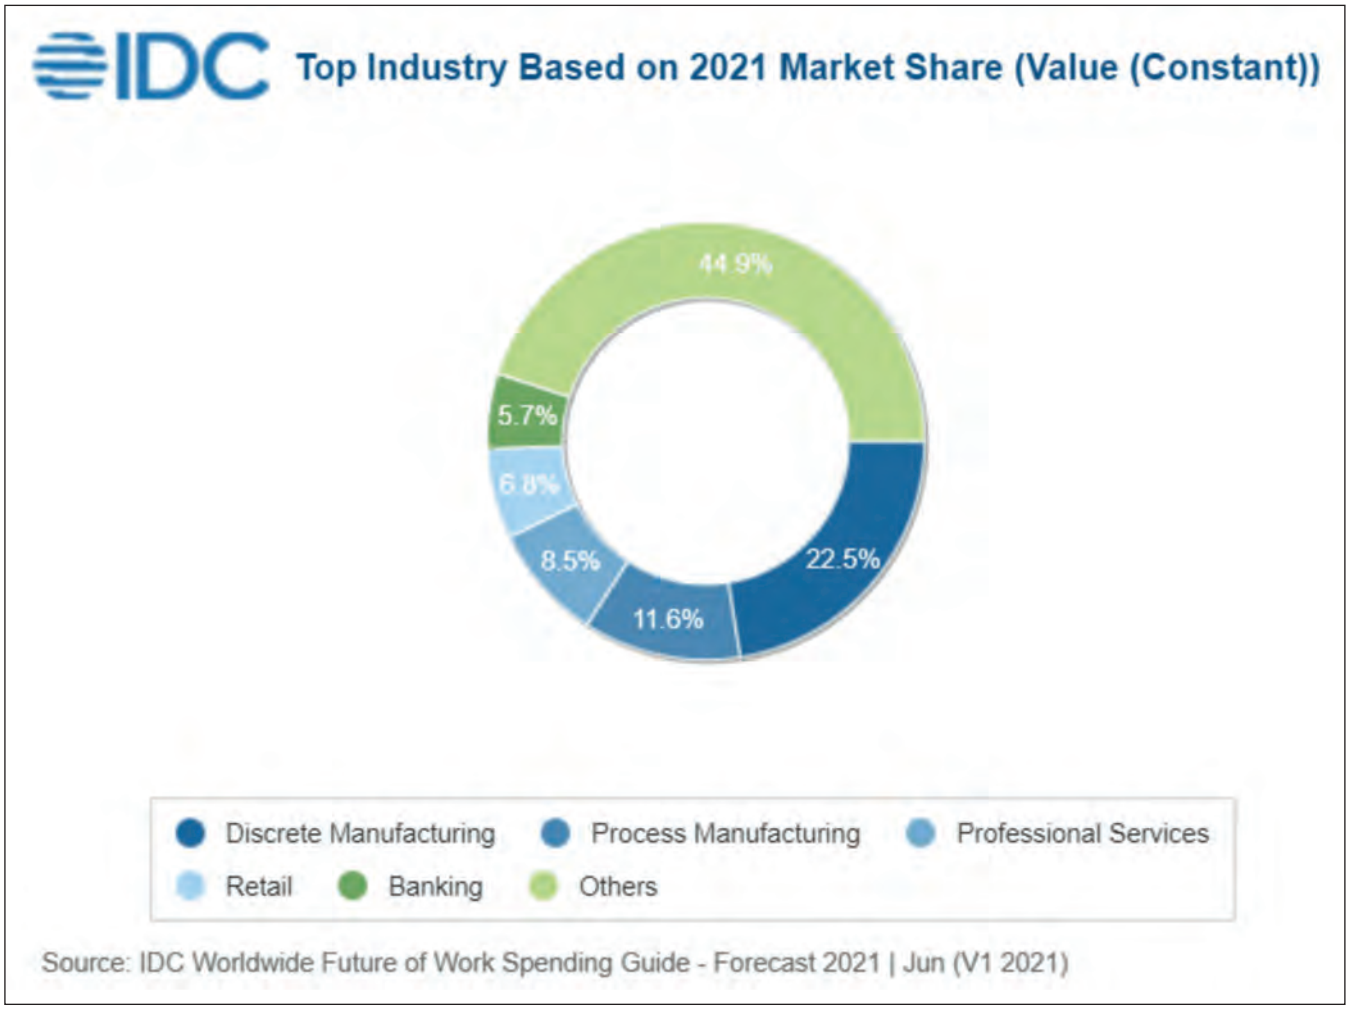


## Double-digit growth for Enterprise Edge spending

International Data Corporation's (IDC) Worldwide Edge Spending Guide estimates that spending on edge computing will reach $24 billion in 2021 in Europe. It expects spending to continue to experience solid growth through 2025, driven by its role in bringing computing resources closer to where the data is created, dramatically reducing time to value, and enabling business processes, decisions, and intelligence outside of the core IT environment. Edge is increasingly being recognized by enterprises aiming to deliver fast, secure, scalable, and reliable products and services to their customers.

In 2021, the largest share of European enterprise edge spending is expected to remain within the services category, driven by connectivity and professional services for IoT, robotics, drones, AI, and AR/VR deployments. Spending on hardware, driven by heavy edge platforms used to perform heavier computing tasks adapted for the edge location or deployment, is almost the same as the spending on software technologies, especially security software used to ensure integrity of data, end points, and infrastructure. In the next few years, the bulk of European enterprise edge spending will continue to be linked to the IoT domain, even though the contribution of other domains, such as AI and AR/VR, will increase further. For example, automated threat intelligence and prevention use cases will leverage AI systems to connect the dots between different pieces of information and to identify possible threats to databases and systems in financial, government, or utilities industries, while AR/VR can help organizations from the manufacturing and utilities industries to improve maintenance, repair, and operations processes. Besides the IoT, AI, AR/VR, robotics, and drone domains, enterprise edge will remain strongly connected to other technology markets such as 3D printing, blockchain, autonomous, wearables, and implantables.

From a vertical perspective, the manufacturing industry will account for the largest share of European enterprise edge spending in 2021, with the majority of use cases related to the IoT and robotics domains. IoT technology can help factories to function autonomically on their own based on a multitude of data collected through sensors and analyzed through various cognitive computing systems as well as using IoT sensors to monitor production systems and

equipment to automatically check final product quality and compliance or detect and anticipate possible failures in production. In this scenario, robots can help factories to achieve full automation by assisting in assembling, inspection, painting, or welding processes.

Retail is the second-largest and fastest-growing industry in the European enterprise edge market, with IoT supporting evolving multichannel retail strategies to provide an excellent consumer experience through any shopping channel and enabling interactive shopping by capturing continuous, real-time streams of data from various devices. Improving customer experience will be increasingly important for retailers, so they will leverage AI to understand customer needs and make shopping recommendations. Other use cases will focus on helping retailers to improve their in-shop losses and optimize restocking using real-time video analytics.

"The pandemic has highlighted the importance of 'remote everything' and of instant access to data and information to enable processes and workloads in locations that are different from the usual core datacenters," said Alexandra Rotaru, research analyst at IDC Europe, Customer Insights & Analysis. "These trends will probably shape the future even after the pandemic ends. Edge computing is a tremendous opportunity to address these new demands, by going beyond traditional IT approaches and extending and innovating on the capabilities found in core datacenters."

## European Augmented Humanity spending to exceed $100 billion by 2025

According to International Data Corporation (IDC), the European augmented humanity (AH) market, including technologies such as augmented/virtual reality (AR/VR), biometrics, exoskeletons, affective computing, ingestibles, injectables, and implantables, brain computing interfaces, wearables, and smart devices, will reach over $50 billion by the end of 2021 and more than $100 billion by 2025.

IDC's The Future of Augmented Humanity in Europe: 2020–2025 Forecast estimates that investments in AH technologies will surge as both people and corporations invest in tech to improve their quality of life and achieve enhanced ways of working.
"Augmented humanity is the advocate of cultural change across the commercial and consumer segments," said Andrea Minonne, senior research analyst and co-lead of the Augmented Humanity Launchpad at IDC UK. "Promoting an AH-oriented culture and complementing human skills with technology will help humans elevate their skills, automate business processes or domestic chores, unlock new capabilities, bring disruption, promote workforce transformation, and enable humanized customer experiences." Consumer needs change fast but streamlining and automating domestic and routine tasks remains a priority. In fact, investments in smart devices and wearables will remain high and will drive a significant share of the overall AH market. At the same time, businesses are searching for tech-based ways to drive innovation and achieve relevant KPIs such as lowered costs, increased productivity, and improved employee safety. Integration is key, and in the future we will see strong investments in integrated technologies such as AI-enabled capsule cameras and AR-infused wearables.

Augmented humanity removes accessibility barriers that limit humans from performing their daily tasks and aims to empower humans and employees with data-driven tech-based innovative solutions to enable them to perform at levels they could not achieve without the support of technology.

## IDC survey illustrates the growing importance of developers

In its annual survey of software developers, DevOps professionals, IT decision makers, and line of business executives from around the world, International Data Corporation (IDC) published a wealth of findings about contemporary developers and the state of software development today. For example, IDC observed that the roles and responsibilities of developers continue to expand. The survey results show that developers are not only the architects and visionaries of digital transformation but also have end-to-end line of sight into the processes that govern the manufacture of digital solutions.

This finding underscores the importance of developers to both technology suppliers and technology buyers because developers have insights into how the operational efficiency of development processes can be improved.
"Given their role in designing, building, and executing digital strategies, developers have become indispensable to the modern enterprise," said Arnal Dayaratna, research vice president, Software Development at IDC. "Developers are product designers, product managers, business analysts, builders, strategists, and sales professionals all rolled into one. This expansion of the responsibilities of developers means that developers are critical to the success and growth of enterprises and organizations as measured by their ability to innovate, execute on strategic and operational plans, and pivot business operations in response to a rapidly changing business landscape."

Key findings from IDC's PaaSView 2021 survey include the following:

- **Full-stack developers are the most common developer role**
  Full-stack developers have proficiency in both the development of business logic and the management of data, as well as the development of rich front-end experiences that are consumed by end users.

**Jacky**
*VP Power Solutions*
*EMEA / France*

**Kathy**
*Sr. Operations Manager*
*China*

**Jitendra**
*DGM Sales*
*India*

*Meet your*

# GLOBAL POWER PARTNER

***Generators, power solutions and 24/7 support.***
***Dedicated to data centers.***

**Kohler-sdmo.com/**EN/application-data-centers

KOHLER®
IN POWER. SINCE 1920.

- **Developer responsibilities have expanded to include deployment, the implementation of automation, performance management, user experience and security**
  Developers are increasingly responsible for the full lifecycle of application development, including operational responsibilities such as the implementation of DevOps and development-related automation, and the implementation of UX and security.
- **Developers feel they have the freedom and autonomy to select developer tools and infrastructures**
  Given this independence, technology suppliers need to ensure that developers are familiar with – and have easy access to – their full portfolio of developer tools products and services.

- **Developers should be considered technology buyers because they have a strong influence over purchasing decisions**
  Between 70% and 79% of developers feel they have either significant or complete influence over purchasing and procurement decisions, including decisions related to the modernization of legacy apps, cloud adoption, and cloud vendor selection.
- **Java is the most popular programming language**
  Java retains its relevance across a multitude of use cases such as enterprise applications, web development, data science, AI/ML, AR/VR, and IoT while modernization efforts have made it more compatible with cloud native infrastructures such as containers.
- **Developers are deploying production-grade applications to the cloud**
  The cloud is no longer principally used for development and test purposes, but is increasingly used for production deployments either in the form of IaaS, PaaS, SaaS or a hosted private cloud.
- **DevOps is now a mainstream attribute of contemporary development**
  This means that automation is increasingly part of development-related operations, and as a result, developers can deploy code and update applications more frequently than ever before. This transition suggests growth opportunities for tools vendors that support DevOps practices.
- **Organizations are investing in legacy app modernization in earnest**
  In 2021, 86% of respondents noted that their organization had modernized more than 50% of their legacy applications, a notable increase from 65% in 2020. This finding serves as a barometer for the pace of digital transformation initiatives and the importance of application development tools and services that facilitate modernization-related work.

"The last year has made crystal clear the value that developers bring to their organizations through software-driven competitive differentiation, and agile development response to changing needs that were unanticipated for many organizations," said Al Gillen, group vice president, Software Development and Open Source at IDC. "IDC's PaaSView and the Developer 2021 brings deep insight into the behaviors, tools preference and technology interests of today's professional developers."

## Cloud explosion?!

Digital transformation (DX) is one of the trends shaping the enterprise world and driving the investments organizations are making to modernize their IT infrastructures and processes. Migration to service-based IT is one of the significant elements of DX initiatives as it enables organizations to utilize IT more efficiently for achieving business goals, pursuing new opportunities, and enhancing customer experience. Until recently, service-based IT was largely associated with public cloud services. However, in the past 18 months, a number of system vendors and cloud service providers have introduced a new class of offerings that are designed to bring cloud experience to enterprise premises.

Dedicated (Local) Cloud Infrastructure-as-a-Service (DCIaaS) solutions deliver compute and/or storage resources dedicated to an individual customer that are deployed on customer premises and consumed as a service. This model is essentially a dedicated version of a publicly available cloud offering, modified to run on premises or in a specially certified colocation environment, including outside of a traditional datacenter environment (edge). The cloud service provider retains full ownership of all underlying infrastructure hardware and software and is completely responsible for delivery, maintenance, updating, and ultimate disposal of the asset when the subscription is terminated.

In its recently published report, Dedicated Cloud Infrastructure as a Service, 2019–2025: Market Trends and Outlook (IDC #US48005321), IDC estimates that the worldwide annual recurring revenues (ARR) from dedicated (local) cloud infrastructure as a service offerings for compute and storage will increase from $138 million in 2020 to $14 billion in 2025 with a compound annual growth rate (CAGR) of 151.8%.

The DCIaaS solutions will be consumed by both enterprise customers and by hosted service providers, which use infrastructure for delivering cloud service to their customers.

"Increasing demand for service-based consumption of IT resources triggered a broad move within the system vendor community to introduce a variety of offerings to fulfill this demand," said Natalya Yezhkova, research vice president, IT Infrastructure practice. "With dedicated cloud as-a-service solutions, enterprises have an opportunity to bring all of what they like about public cloud to their own premises while mitigating restrictions and concerns they have with moving infrastructure off-premises."

HYGOOD

# Stop fires.
# Not hard drives.

## Acoustic Nozzles

For suppressing fires in hazard areas where sound levels from standard nozzles may affect sensitive electronic equipment.

For more information visit hygood.com

The power behind **your mission**

Johnson Controls

# Kubernetes & containers

## Powering tomorrow's applications

Containers and Kubernetes are the driving force behind how the industry is reinventing the way we build and run applications, fueling enterprise IT efficiency.

BY JAMES PETTER, VP INTERNATIONAL AT PURE STORAGE

CONTAINERS are a standard unit of software that packages up code and all its dependencies so that an application runs quickly and reliably from one computing environment to another. Containers make it easier to roll out cloud-based applications because they contain all the information needed to run them in manageable packages. In September 2020 we announced the acquisition of Portworx®, the industry's leading Kubernetes data services platform for approximately $370 million, so it's safe to say we recognise the significance of the technology. Let's take a look at how we got here.

### Importance of data centricity

Data is at the heart of tomorrow's businesses. Leading digital organisations are using a new "cloud native" technology stack to process this data into value and insight. Cloud-native applications are specifically designed to operate in a cloud-like manner, whether in the public cloud or on-prem, from day one. They can be deployed and fixed faster, and can be moved across different environments easily.

Cloud-native applications are typically made up of microservices (more on these later) and are packaged in containers. This new cloud-native stack includes a new set of applications - apps that analyze streaming data in real-time, apps that index massive quantities of data for search, and apps that train machine learning algorithms on increasingly large data sets - undoubtedly this cloud native revolution is being powered by a combination of containers and Kubernetes.

Containers make it efficient to run disaggregated applications at high degrees of scale and fluidity with minimal overhead, and Kubernetes creates the machine-driven orchestration that can juggle all these application fragments and assemble them into a composite application as necessary.

## Container adoption speaks for itself

Adoption rates of this new cloud native stack have been staggering. According to 451 research, 95% of new apps are developed in containers. Enterprises are evolving their cloud strategies to be multi-cloud, and containers are also key to this. Gartner reports that 81% of enterprises are already multi-cloud, working with more than two cloud providers. Gartner also predicts that 85% of all global businesses will use containers in production by 2025 - a huge rise from just 35% in 2019.

It's still an early market with huge growth potential so it's inherently hard to forecast, but IDC predicts that the commercial market for container infrastructure software alone will top $1.5B by 2022, and enterprises are paying attention.

## Microservices and containers - a perfect match

Put simply, microservices are the individual functions within an application, and form the basis of a new architectural approach to building applications. Microservices enable IT teams to more easily build and run the applications their users want and need to stay ahead of competitors. Many of the largest consumer and enterprise applications today run in microservices, proving that it's not just a trend for small organizations but also for the largest and most complex. Indeed, the larger the organization is, the more benefits there are to gain from adopting microservices because teams are often spread out with limited direct communication.

When was the last time you got a maintenance notification from your favourite streaming service to let you know you won't be able to access services? It doesn't happen. There's never a good time to update these services because someone is always binge-watching a new show. The principle of microservices states that you should break an application into smaller pieces that communicate via APIs, where each part can be updated independently from other parts. As a result, if a streaming service needs to update its password-reset functionality, it doesn't need to kick millions of users offline. This feature is a different microservice that can be updated independently. This results in happy developers and happy users.

Microservices are here to stay and will underpin the applications of tomorrow. In what kind of environment should you run them? Containers are the perfect building block for microservices. They present a lightweight, consistent environment for microservices, that can follow the application from the developers desktop, to testing, to final deployment. In addition, containers can run on physical or virtual machines, and they start up in seconds or even milliseconds, which is faster than VMs.

Packaging applications with their dependencies
Traditionally, software packages have included all the code needed to run the application on a particular operating system, like Windows or Linux. However, you need more than just application code to run an application, you also need other applications. For instance, an application for looking up stock prices might use a library to convert company names to ticker symbols and vice versa. This functionality is generic and not value-added, but it's still important to allow a user to type "Apple" and get the stock "AAPL." The library is an example of a dependency. Without IT knowing it, any application might have hundreds of these types of dependencies.

One of the main reasons that containers became so popular is that they provided a mechanism and format to package application code – with its dependencies – in a way that made it easy to run an application in different environments. This solved a big problem for developers who were constantly fighting environment-compatibility issues between their development laptops, testing environments, and production. By using containers to package their applications, they could "code once and run anywhere," dramatically speeding up application delivery.

Not all container services are created equal
In terms of challenges, the first generation of cloud native applications were designed to be stateless - using containers which did application work but didn't need to store any persistent data in associated volumes. As container usage evolves, developers are increasingly building stateful apps inside containers - apps that need to store data in a volume that must be persisted and kept. This is where the world of storage becomes challenging.

The flexibility and openness of containers turns into hurdles and bottlenecks at the storage layer, and simple storage capabilities that we've been taking for granted for years in the traditional application stack (high availability, disaster recovery, backup, encryption) become challenges in the container world. What's worse, what often happens is that each application devises it's own storage strategy, making it impossible to drive standards and data compliance across an organization.

This is why as a best practice we recommend choosing a solution that delivers the Kubernetes-native data services that both cloud native and traditional apps require (since those traditional apps aren't going away anytime soon). This means delivering block, file, and object storage services, in multiple performance classes, provisioned on-demand as Kubernetes requires. It means providing instant data access, protection across all types of failures, the ability to mobilize data between clouds and even to/from the edge, and robust security no matter where an application travels. If organisations do this they will see for themselves why Kubernetes has become the not-so-secret special sauce for modern organisations.

# The future of the workplace is digital: HERE'S HOW TO NAIL IT

With vaccination campaigns rolling out, lockdowns lifting and borders opening, it seems the global economy is slowly getting back in gear. But that doesn't mean that things are going back to the old normal – especially when it comes to the workplace.

**BY SANDER BARENS, CHIEF COMMERCIAL OFFICER, EXPEREO**

IF WE GO BACK 20 odd years, the connectivity aspect of the office environment was pretty straightforward. Companies would have their network spread across a couple of HQs; applications would be hosted in on-site servers; and staff would be connected via the same link.

Today, however, end-users are everywhere. They've gone out of what once was a controlled network environment into multiple locations, creating what we call the 'digital workplace'.

The digital workplace isn't a "place". It's in the connections between locations, people and tools that allow for seamless collaboration across all the different applications used by the organization. The tools need some very specific qualities as to what they do and how they work together. But the people, you've already got. The challenge is to make sure that there is proactive support for the whole adoption of this new, digital way of working.

This involves non-technical aspects such as management practices, organizational processes and team responsibilities. Once that's been taken care of, leaders should assess how well positioned they are to take advantage of everything that collaboration offers – that's when the right tools come into play.

## Moving to the Cloud

Working in the cloud is what fundamentally allows for remote collaboration. By hosting applications in cloud services rather than on-site, enterprises give users access to company resources in real-time, from anywhere, on any device – as long as it's all connected. Compute loads and data bandwidth need constant monitoring, balancing, and capacity management. In order for applications to work in an optimized fashion, you need high bandwidth and 24/7 availability, meaning that the investments made on applications need to be properly supported and underpinned by an effective connectivity solution.

## Making Connections

Distributed sites may rely on an Internet Service Provider or telco, but digital nomads need equally reliable mobile connections on the road, while those working-from-home may use their own broadband connection. All can be addressed by different underlay types of connectivity, such as fixed internet, mobile internet, and / or Bring Your Own Access (BYOA). The overlay, in turn, securely and intelligently steers traffic across multiple sites, adjusting bandwidth where it's needed most.

## Virtualizing Your Business Network

The Internet itself is not one big happy cloud. You need further insights in order to make sure that you've chosen the right local connectivity options and that the routing to your sites and employees is in an optimized way. That's where SD-WAN can help.

SD-WAN is a virtual private network for businesses, acting as an overlay for a company's existing network solution (may it be MPLS, broadband Internet, mobile 4G/5G or a hybrid model). It separates the network control and management processes from the underlying hardware, making them available as software. By using a centralized control function, it securely and intelligently steers traffic across multiple sites, adjusting bandwidth where it's needed most.

SD-WAN is also secure by design, as it runs on encrypted end-to-end tunnels that are built across the whole network. But, with devices, data, and apps far from the corporate HQ, and connectivity taking place on public wi-fi and home broadband, it's no longer sufficient to secure access only to the network, however good a job SD-WAN does of that.

## Getting Sassy with SASE

Security, of course, makes connectivity effective. That's the reason for SASE, or Secured Access

Service Edge. SASE is about secure access at the application level, with the emphasis shifted to authenticating users and devices on an as-permitted basis at the network perimeter, rather than the once-you're-in-you're-in approach of a typical in-house setup – from user identification to logins and permissions, removing penetration risks at every step of the connection between people.

Granting access to specific applications rather than the network as a whole, with IP cloaking making that access invisible even to malware on a compromised device, keeps the network perimeter safe. The door is not merely open or closed; it's a hidden door that only opens to people who know it's there.

## The Value of Smart Edge

Cloud computing, SD-WAN, SASE… all relate to Smart Edge. Smart Edge's purpose is simple: to make sure devices work and interoperate smoothly, so users don't have to waste time grappling with technology. It's not a product per se, but a catch-all term for the managed services you get from a Managed Service Provider, making products from different hardware and software sources work together. After all, it's only when people feel truly connected, wherever they are and whatever equipment they're on, that they truly feel part of the same team. And that's how you nail the digital workplace.

Cloud computing, SD-WAN, SASE... all relate to Smart Edge. Smart Edge's purpose is simple: to make sure devices work and interoperate smoothly, so users don't have to waste time grappling with technology

# Maximising your fibre network **ROI**

The COVID-19 pandemic has helped unify the call for universal broadband to ensure that no family, frontline worker, small business, or community is left behind without connectivity again. The question going forward for communications service providers (CSPs) is which long-term broadband network strategy ensures high-speed broadband is available to everyone.

BY VANESA DIAZ, SENIOR MARKET DEVELOPMENT MANAGER FOR CORNING OPTICAL COMMUNICATIONS

FIBRE BROADBAND NETWORKS provide the best solution and have proven their ability to scale over time well beyond comparable copper or wireless technologies.

Given the lifespan of a fibre infrastructure, CSPs may want to challenge their expectations with respect to the time to recoup their investment, given the lifespan of a fibre infrastructure. An underlying fibre broadband network empowers CSPs and communities to meet the challenge of bringing ultra-broadband capability for both fixed and mobile networks. Significant investment is required to achieve this goal. Through planning, CSPs and communities can maximise the return on this investment, giving their customers the broadband network that best prepares them for the future.

## Build it Once

As CSPs and communities navigate their future a lot can be learned from the past, such as the build-it-once approach. The last thing any CSP wants to do is make an investment in an underlying communications network, only to find out years later that a significant network investment is required to keep pace with innovation and customer demand.

Historically, networks were built in this way – a network was built, only to learn years later that additional investment, upgrades and development was needed to meet the demands of the evolving customer segments and applications. This has often led to multiple interconnected networks which are incredibly complex and expensive to maintain.

Considering the total cost of a build is heavily influenced by installation labour, in some cases upwards of 60%, building a network that requires little to no new outside plant construction after initial build to expand or upgrade will minimise the total cost of ownership. To realise a better ROI, implementing a strategy to build one underlying fibre broadband network that enables various applications is prudent. As fibre technologies advance, the infrastructure they ride on largely remains unchanged. Building an all-fibre network gives CSPs the ability to converge

customer demands on a single unified network. Revenue models multiply as a result, delivering an enhanced ROI.

The build-it-once approach also extends to mobile and smart cities. If communities want to participate in a 5G future, an underlying fibre-rich network is required. Fibre-fed small cells will power both 5G and Wi-Fi, making a wireless ultra-broadband experience possible. Tangentially, the thousands, or even millions, of sensors needed to deliver on the promise of the smart city will need fibre connectivity to function properly.

Whether it's utilising wave-division multiplexing (WDM) on an optical transport network (OTN) segment for 5G transport or employing NG-PON2 to meet increasing residential, SMB, or enterprise demands, CSPs can truly maximise their revenue opportunity from a foundational fibre-based network. Build it once and leverage it several times over.

As CSPs embrace a fibre broadband strategy, there are several factors that can improve ROI. Among all key strategies, noteworthy options include leveraging existing network assets and considering all network design options. If minimising CapEx is a primary concern, considering "lean fibre" architectures may also be an option to explore.

## Leveraging Existing Network Assets

Traditional telcos have network assets that they've learned to leverage for decades. These assets allow long-standing CSPs to bring fibre deeper into their networks, with fibre-to-the-premise (FTTP) networks now proliferating across the country.

Newer CSP entrants, such as electric cooperatives and municipalities, are fortunate to have similar network assets that traditionally were used for delivering electricity. Those same assets can be easily repurposed for a fibre network and go beyond the obvious poles and rights-of-way infrastructure critical to a fibre-build, to include assets like huts, substations, towers, trucks, and technicians.

Technology enhancements that shrink fibre terminals and drops may allow existing handholes or pedestals to be leveraged, reducing upfront costs even further.

## Fibre Network Design Strategies

Planning is necessary to maximise future revenue opportunities. That includes ensuring fibre network designs anticipate expansions, dark fibre demand, or other unforeseen opportunities. This is of interest for those who are pursuing funding which can augment investments in fibre broadband builds.

Funding programs target unserved and underserved territories, but through proper design and planning, CSPs can build networks that are conducive to future expansion into neighboring markets.

Additionally, if smart city or 5G is a part of the vision, network designs should consider higher fibre counts than what is indicated currently. More fibre capacity from the beginning creates potential for dark fibre leasing and other wholesale opportunities, as well as for taking advantage of unforeseen future demand to generate additional revenue opportunities. Fibre is unique because it's the underlying infrastructure that best enables expansion to address market opportunities that may present themselves. Designing a network with these factors in mind is critical for maximising fibre broadband ROI and creating additional revenue opportunities.

## Lean Fibre Architecture

For certain CSPs serving remote territories, turning to a lean fibre architecture strategy can positively impact a fibre network ROI by lowering upfront CapEx and splice labour during construction. A lean fibre strategy relies on distributed split architectures. This approach puts less fibre facilities in the network, whilst providing FTTP services. Ideally, all CSPs would follow a more fibre-rich home-run or centralised split architecture, both of which offer a high degree of flexibility, bandwidth capacity, and room for future network expansion. But these also require higher levels of upfront CapEx, which can be a drag on fibre network ROI – particularly in lower density markets.

Adopting distributed split or optical tap architectures lessens feeder and distribution fibre cable requirements, thus lowering fibre management material and fibre splicing construction costs. This can improve the ROI calculations for less dense markets. However, the trade-off of these leaner architecture options limits bandwidth flexibility to discrete locations and future expansion as a result. This approach should be used in service areas where significant expansion and growth are not expected.
The lean fibre approach can make sense in certain applications. Some of the downside risk can be mitigated by using a higher-capacity main distribution fibre cable that serves as surplus dark fibre. This approach takes advantage of the lower-cost lean architecture strategy, while including additional capacity in the network for future growth.

## The Best ROI Path

Every broadband business strategy should aim to utilise an underlying technology that can best enable all applications, regardless of bandwidth demand or latency requirement. That capability gives CSPs an ROI advantage. It allows them to maximise their revenue opportunity. No other technology can handle the bandwidth and latency requirements of not only today's applications, but tomorrow's as well.
With proper planning and a build-it-once vision, CSPs are not only positioning their companies and the communities they serve for the future, they are also maximising their ROI opportunity. By embracing some of the strategies outlined above, the investment in the network yields returns well beyond what can be found on a balance sheet.

# You have a remote workforce:

# WHAT NOW?

The pandemic caused an abrupt shift to the way we work and is expected to have implications for years to come. The biggest impact was on businesses that had to pivot to remote working.

**BY JAY TURNER, VICE PRESIDENT, DEVELOPMENT AND OPERATIONS AT CONSOLE CONNECT BY PCCW GLOBAL**

UNDERSTANDABLY, at first, many companies and their employees were concerned about moving fully online. Today, the majority of workers believe they are as effective when working remotely and, in many instances, want to ditch the office all together.

As most companies have already leveraged technology that enables flexible working models, such as cloud services, dedicated networks and VPNs, leadership is now looking into making flexible work a permanent operations model. With that in mind, the next step for IT and security professionals on the digital transformation journey, is to avoid unnecessary risks caused by rapid changes in both user and attacker behaviour.

## Post-Pandemic IT Headaches

The changes to the nature of work are likely to stay

for virtually all companies across most sectors. This throws a number of challenges and considerations into the mix for any enterprise IT professional tasked with enabling existing workforce to continue or those joining in to begin to work remotely.

In light of the growing appetite for cloud solutions, there has been a massive uptake in cloud-based business applications particularly for collaboration and communication. Applications such as Microsoft Teams, Slack, and Zoom, and anything that enables unified communication or video collaboration will be in demand and workplaces need to ensure smooth access to those tools.

Maintenance is another potential challenge facing corporate networks as it has been much more difficult to carry out scheduled network maintenance over this period. In addition, employees working longer hours and more frequently online in the evenings and weekends blurs the lines of a 'working day'. All of this has an impact on the network traffic, which previously was more predictable and fit a recognisable pattern. For instance, network managers were able to set expectations around low traffic times such as holidays.

## But That's All in the Past.

Furthermore, Virtual Private Networks (VPNs) are no longer fit for purpose when it comes to the large-scale shift to remote work. As more employees use a VPN to access the access network, more traffic is pushed to the network edge. In other words, the majority of network traffic is now originating from outside of the office, and needs to travel, for example, to the server room then back out to the edge again. The traffic flow increases demand for edge accessibility as well as the load on the hardware responsible for terminating the VPN connections.

## Thriving Despite Challenges

Unsurprisingly, enterprises may start to look at direct connectivity to the applications themselves. Establishing direct and private connections to your cloud provider can help alleviate some of the new demands on a corporate network caused by remote working. Using the power of Software Defined Interconnection® will help address those concerns in an efficient way.

For instance, when struggling with maintenance challenges, one approach could be adopting a network model that is flexible and scalable. This is particularly beneficial in the face of unpredictability and potential disruption with connections and bandwidth that can be set up, flexed up, and torn down in near real-time. That's again where Software Defined Interconnection® can play a role. By using technologies that interconnect networks and clouds in real-time, businesses can self-provision redundant network links – for a day, a week or even a month – that can be used to carry traffic during maintenance.

Furthermore, it is critical to have transparency and visibility over the whole network, allowing IT managers to ensure that business applications are all up-to-date and that employees are running the most secure version. As more and more workers pick up these tools a harsher spotlight has been cast on their suitability. For example, Zoom's initial lack of end-to-end encryption being just one consideration, but then there was concern over the company's privacy policy and what it could do with the data of its hundred million (or more) users. Companies need to have visibility into the traffic to see where the pinch points might be and where traffic is destined to ensure the flow of remote work.

## Empowering Virtual Workforce

Although moving workforces online was necessary for many companies, it is now the right time to think of the IT and business infrastructure that will enable operations in the long run. Businesses need to focus on now to efficiently manage the needs of new remote workers while keeping the network secure.

It is critical to have transparency and visibility over the whole network, allowing IT managers to ensure that business applications are all up-to-date and that employees are running the most secure version. As more and more workers pick up these tools a harsher spotlight has been cast on their suitability

# GENERALISTS vs SPECIALISTS

## Different approaches to IT monitoring

When a company is considering a network monitoring solution they must take into consideration the skill sets and capabilities of their own IT team to make sure they find the right solution that fits their needs.

BY MARTIN HODGSON, HEAD OF UK & IRELAND, PAESSLER AG

QUITE OFTEN IT monitoring is the unloved child of the IT administrator: it is necessary in order to ensure the smooth functioning of his company's IT estate, but they usually don't have any fun doing it. As a consequence, a large number of highly qualified IT experts have surprising knowledge gaps when it comes to network monitoring, this usually comes to light when the monitoring solution used is no longer adequate or, even worse, has become so complex over the years that it is barely operable.

### What Exactly is IT Monitoring?

On a very basic level, classic IT monitoring is the monitoring of availability and performance in IT environments. IT monitoring answers questions such as "Is my server online?", "Does data in my network get to where it needs to go on time?" or "Is my firewall working reliably?".

The basic function of IT monitoring can be broken down into four tasks:

1. Determine and collect data on the availability and performance of IT components
2. Data storage
3. Notifications and alerts based on defined thresholds
4. Data reports

Of course, there are numerous more advanced tasks in the IT monitoring environment. These include root cause analysis to get to the heart of a problem or recognising emerging trends and making predictions based on them. It can also cover monitoring the security of the network such as the function of firewalls or virus scanners or recognising unusual behavior in the network through intrusion detection.

Logging or event log management is also often listed under monitoring and refers to the analysis of log files such as syslog messages or SNMP traps and is often listed under SIEM (Security Information and Event Management).

### Generalists vs Specialists

Time and again, generalists are compared with specialists, but a generalist cannot replace a specialist and vice versa.

### Specialised Solutions

Highly specialised solutions provide specialists with deep insights into narrowly defined areas of the IT estate. The larger a company or the deeper its IT structure, the greater the need for such specialised solutions. DevOps requires detailed information about applications, SecOps needs in-depth insights into security-relevant aspects of network traffic beyond classic tools such as virus scanners or firewalls, while NetOps relies mainly on in-depth analysis of network performance.

In the area of network performance, solutions such as Scrutinizer by Plixer, Flowmon by Kemp or Kentik provide this, sometimes even beyond the narrow boundaries of a specific application area. Flowmon, for example, claims to inform SecOps and NetOps in equal measure. Nevertheless, Flowmon remains a solution for specialists that does not offer a general overview of the entire IT - which is not the claim.

Most specialised tools focus only on a few methods or protocols. In network and application monitoring, this is often flow or what's called 'packet sniffing'. When it comes to security it can be flow or packet sniffing using deep packet inspection, but also event log monitoring. Here, the tools usually deliver outstanding performance, scale even for larger environments,

and offer in-depth data analyses beyond just pure monitoring, some of which also rely on artificial intelligence or advanced algorithms.

On the other hand, this also requires the IT teams to have a certain level of expertise when using the tools. Even if the operation is optimised and designed to be as simple as possible - in order to the maximum value from the tools, the necessary expertise is required to first configure and deploy the solution correctly and then to be able to use the determined data in a target-oriented manner.

However, if the user needs a central overview of the performance and availability of the entire IT - from infrastructure to network to cloud-based applications and perhaps even beyond to areas beyond IT, then even the specialists quickly reach their limits. This is where generalists are needed.

## Generalist Network Monitoring Solutions

SNMP (Simple Network Management Protocol) is often the basis of generalist network monitoring. Even though the protocol is not technically up to date and is regularly declared dead, it is still so widespread that broad IT monitoring without SNMP is not really practicable even in 2021. After all, IT environments are usually not completely overhauled, but continuously adapted to new requirements, older devices and structures coexist with modern systems. Thus, interface-based systems (API = Application Programming Interface) are also becoming increasingly important due to advancing digitalisation. Monitoring tools must still support traditional methods such as SNMP, Ping or WMI, but they must also support interfaces and other contemporary methods such as MQTT or OPC UA.

A central overview with the help of a generalist can be sufficient in smaller companies or limited IT environments. With a few dozen or a hundred devices or IP addresses to monitor, the administrator or ITOps (IT Operations) team usually knows their IT so well that a simple error message is enough for them to be able to identify and solve the problem. They don't need highly specialised tools and usually don't have the time or expert knowledge to use them efficiently. They prefer a monitoring solution that is as broad as possible and provides all the necessary information, is easy to deploy and operate, and alerts when intervention is required.

In larger companies or IT environments with specialised teams that require correspondingly specialised tools, there is s usually also a need for an overarching solution that provides a central overview and allows trends to be identified. This may be the case at the ITOps level or in higher-level management. This is where the generalists come into play. They provide information on all areas of IT without going into too much depth.

Generalists usually support a wide range of protocols and ideally also offer appropriate interfaces to enable the broadest possible collection of data. The data collected is stored and provided with threshold values, which are used to send notifications and alerts via a wide variety of channels. This includes monitoring traffic as well as devices, applications, storage systems, databases or cloud services.

Some generalists offer the possibility of relating data from different areas to each other. These include predefined interfaces to collect and integrate information from specialists. For example, the performance of multiple, redundant mail servers can be combined with traffic data, load balancers, firewalls, databases, storage systems and other components in a single service or process. This enables a management overview from a company perspective across a wide range of IT areas. Set up accordingly, the entire process is displayed as productive and functioning, even if individual components report problems.

As long as the entire process is running, only the responsible employees are informed in order to fix the affected problems before it becomes critical. Only when the entire process is at risk, appropriate alarms are triggered. When a company is considering a network monitoring solution they must take into consideration the skill sets and capabilities of their own IT team to make sure they find the right solution that fits their needs.


# Authentication best practice within a **Zero Trust** strategy

For CISOs around the world it seems at every turn, they are being told to implement Zero Trust.

BY NIC SARGINSON, PRINCIPAL SOLUTIONS ENGINEER, YUBICO

IT'S TRUE that the principle of Zero Trust might be right for this moment; after all, work-from-home and hybrid work policies are becoming the norm, as are cloud applications. However, many organisations may struggle with the reality of what's involved with a Zero Trust framework or infrastructure. Having protected the boundary for so long, it takes a change of mindset to verify every connection attempt.

At its core, Zero Trust should start with strong user authentication and the chosen authentication method should not hamper user productivity. Therefore, organisations need to look first at how users establish their identity and consider the level of trust that can be attributed to that mechanism. The truth is, if authentication is by passwords alone, there is no assurance of security, no matter how robust the rest of the Zero Trust strategy is. Yet despite this, a recent survey of work-from-anywhere cybersecurity practices at companies in the UK, France and Germany revealed that less than a quarter (22%) of respondents had implemented two-factor authentication (2FA). That's a problem for Zero Trust, because going forward with such a model depends on having a strong level of trust in the authentication mechanisms of every user, from every device.

Strong authentication needs to be a foundational building block of the Zero Trust strategy. With that in mind, what are the key strong authentication best practices organisations need to adopt in order to ensure Zero Trust is correctly supported?

**1. Choose Strong Authentication Based on Open Standards**
By decoupling authentication from the identity and access management (IAM) platform, and by choosing an authenticator based on open standards, authentication will work with a wide array of IAM solutions. That way, users are empowered to be productive on a new IAM system, or non-federated access point, using the same authenticator within minutes instead of weeks.

**2. All Accounts Must be Considered**
Service accounts, as well as user accounts, need to be heavily protected, monitored, and properly scoped. Too often, these types of accounts are protected with static passwords. That isn't sufficient, but unfortunately a number of IT and other systems have limitations on authentication options. However, they can often make use of cryptographic certificate-based authentication – private keys that should be stored in hardware security modules (HSMs), dedicated security hardware that come in different sizes, from large physical appliances to small USB devices.

**3. Cryptographically-Based Signing is Key**
It has been possible for quite some time to digitally sign electronic documents and personal authenticators and inexpensive HSMs make this easier and stronger. Cryptographically-based signing, backed by hardware, ensures that content was in fact created by the signer.

**4. Validate Devices**
Strong authentication, such as that provided by a hardware device, supports a Zero Trust approach but it is still very important to validate the device itself to ensure it is not compromised. Attestation validates that the authenticator hardware is from a trusted manufacturer and that the credentials generated on it have not been cloned. Attestation is a key pair that is burned into the device during manufacturing, providing important details such as manufacturer and device model. Attestation concepts are built into the FIDO standard and some vendors also include attestation capabilities for smart card deployments.

**5. Remember Risk**
A trusted strong authentication approach allows for step-up authentication based on risk. This protects the user and the organisation while increasing productivity. Real time risk-based access policies, such as those implemented in a Zero Trust framework, are based on signals and risk scores. A strong authentication solution that is hardware-based, and highly trusted, can elicit a high trust score, thus allowing for higher privileged access.

**6. Purpose Build Phishing Resistance**
The earlier cited survey into working from home cybersecurity tells us that, where companies have implemented 2FA, mobile authentication apps and

SMS one-time passcodes (OTPs) are the most popular.

It's true that these basic forms of 2FA provide higher levels of security than username/password alone, but they are not invulnerable to some threats, such as sophisticated phishing and man-in-the-middle attacks. OTPs via SMS can also fall into the wrong hands as a result of 'SIM-swap' fraud, and employees can be tricked into providing them to a would-be hacker if they're persuaded it's a legitimate request.
A strong, phishing-resistant authentication solution should be purpose-built. It should also reduce, not add to, authentication complexity. A dedicated security-focused device that is simple to use heightens security without an impact on productivity and also allows for easy and consistent monitoring.

**7. Plan For a Passwordless Future**
Achieving secure passwordless login across desktop and mobile requires a rich ecosystem and a consistent framework for authentication. An ecosystem built on open FIDO2/WebAuthn standards is best placed to deliver security and usability, while also satisfying the need for portability, compatibility, interoperability and scale.

Modern multi-factor authentication (MFA) is essential to prevent network access through stolen passwords. Now that the industry is moving away from symmetric based secrets (passwords, OTPs) to more advanced asymmetric solutions bound in physical devices, it's more important than ever to start with strong authentication if Zero Trust is to become a reality.

# Observability is not one size fits all

Roles for data in your application development and infrastructure.

**BY IAIN CHIDGEY, VICE PRESIDENT EMEA, SUMO LOGIC**

OBSERVABILITY HAS BECOME more important for developers and for enterprise IT teams. More companies have adopted cloud infrastructure and microservices to create their applications in order to deliver updates faster and be flexible around digital services. However, these shifts on the architecture and design sides have an impact on how to keep up with what is going on within those applications.

Observability takes data outputs from an application, and then uses that data to understand how the application is performing. The standard definition for observability today is that developers combine their application logs, metrics and tracing in order to get the full picture of application performance. However, it is not as simple as just having sets of data coming in.

Instead, it's worth looking at how you can improve your use of observability data to achieve your goals. This can be more specific than simply gathering data over time, then using it when something goes wrong. You can improve this process by setting up more specific goals, based on the data you have coming in continuously from your applications and cloud infrastructure.

## Getting Your Goals in Place

The first area for observability is around application development and reliability. Using data from your applications' continuous integration and continuous deployment (CI/CD) pipelines, you should be able to see how quickly you are rolling out updates and any problems that come up over time. However, this is still a reactive process. Instead, how can you use that data to spot potential problems - or ways to improve - in advance of a problem coming up?

One of the big challenges for many application development teams and site reliability engineers is, paradoxically, not when something goes down. After all, getting to the root cause here should be easier because the broken component should be obvious - for example, it could be DNS, a faulty change that should be rolled back, a cloud service that is not available, or a network connectivity problem.

However, some of the hardest problems to deal with are not so black and white. Instead, services can degrade over time, working but not at the level expected. To solve this involves some more preparations, based on setting up the right metrics for

each group of components. Once these metrics are in place, you can use your observability data to track that performance over time and then be more proactive on fixing problems.

Another area for a more specialist approach around observability is around Kubernetes. More developers have turned to containers to host their microservices applications, and Kubernetes is the de facto standard for managing those containers. For observability, getting data out of your Kubernetes and containers will help show how those applications are performing. This involves bringing sources of data together including Prometheus, FluentD, Fluentbit, and Falco, or using a framework like OpenTelemetry.

With this data, you can connect any microservice performance issues and errors directly to user experience and then make the right changes. To achieve this, it is essential to understand end-to-end user transactions, uncover latency issues and see which services are impacted. Distributed Transaction tracing provides the telemetry to connect the monitoring of key performance indicators to the real experience of your users.

## Getting a Full Picture

One area that is growing in importance for companies is how to track multi-cloud deployments. Multi-cloud can mean different things to different organisations - from departments choosing their own cloud providers to meet specific needs, through to companies planning their expansion across multiple markets with the most appropriate partner in a country or region, getting data from each cloud provider is only part of the story. Alongside this raw data, it's essential to work on how you normalise this data and get it into one place for tracking.

While there are many cloud services that are either compatible with each other or broadly similar in what they provide - for example, commodity services like block storage or compute - there will always be some differences in how the cloud providers operate, and there will be specific tools available that provide alternative services. Comparing information across cloud providers and getting consistent insight into performance is therefore a task that observability data can support.

Similarly, software development teams can get more value out of their observability data. Like the old tale of the cobbler, software development teams can be better at providing data for others to use before they think about applying data to their own processes. Instead, data from software development tools like Jenkins, JIRA, GitHub, BitBucket, OpsGenie, PagerDuty and others can be evaluated alongside cloud and application component data. Rather than looking at this data solely for alerting or performance, you can use the data from the CI/CD pipeline to look for opportunities to optimise the whole process.

The last area for improvement is around edge and web observability, where we can demonstrate all the hard work that has gone into delivering better services

This is a developing area as application and software teams create multiple pipelines to process their work through the organisation – for enterprises, having tens or even hundreds of different CI/CD pipelines in place is not uncommon. Consolidating data from those pipelines can then show you how well you are currently performing, and benchmark data can be used to spot opportunities to improve performance at the team level.

The last area for improvement is around edge and web observability, where we can demonstrate all the hard work that has gone into delivering better services. As an example, when we order something from an online retailer, we don't consider all the moving parts that go into our purchase - we simply look at whether the service works, and the retailer takes on all the responsibility. When something fails, customers don't think about the third parties involved.

As we create our applications or services, how are they delivered to customers and how can we improve the delivery of that content? How can we ensure that all that good work internally then gets to the customer at speed and avoids potential performance problems? To get this insight, we can use observability data from the content delivery networks and cloud providers that support the delivery of those services. By seeing how services are experienced by customers, we can spot potential problems outside the application itself and take steps to improve.

Observability data holds up a mirror to our systems, and helps us understand what is going on based on the outputs delivered. However, there are more opportunities to make more of the data that we have coming through across our operations in multiple areas, from business or software development processes through to areas like security. By looking at the full picture and consolidating all our data together, observability can help us improve across the organisation as a whole.

# Limiting the use of **Machine Learning & Artificial Intelligence** for better results

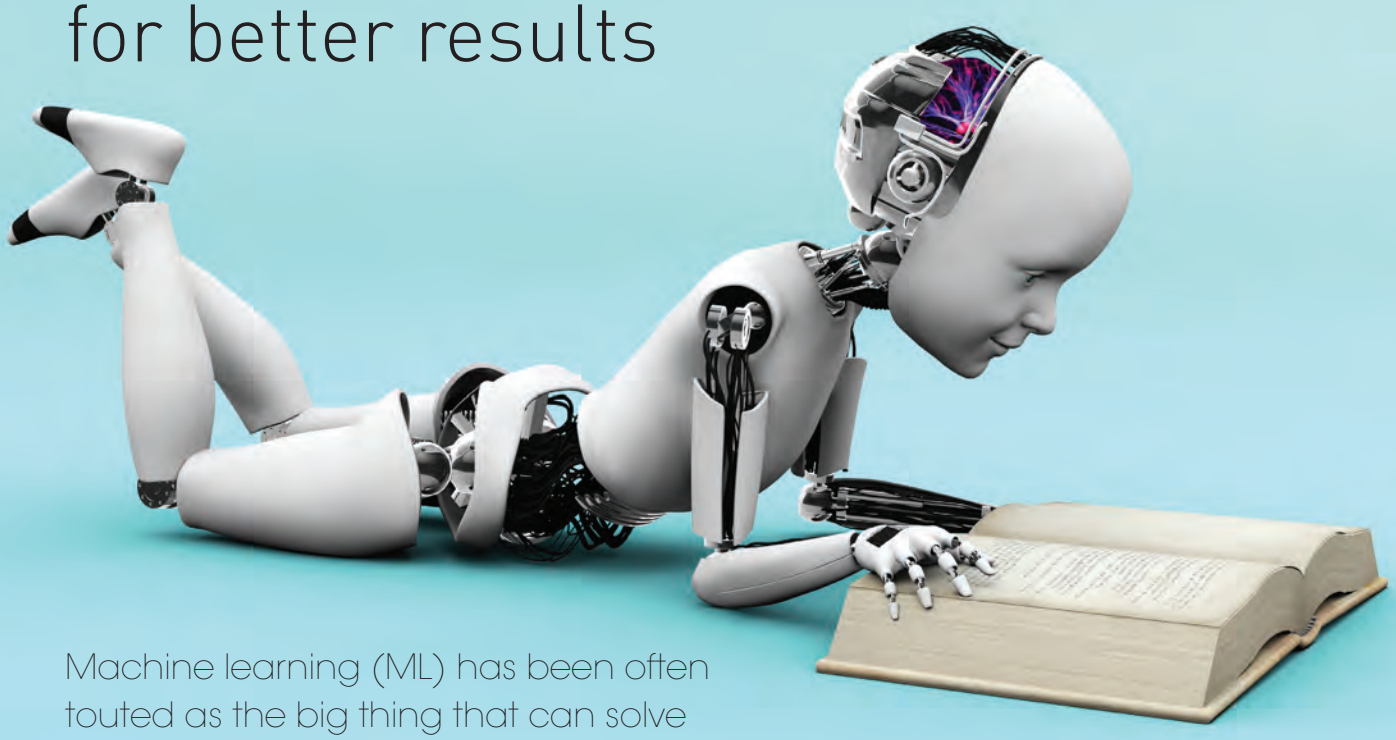

Machine learning (ML) has been often touted as the big thing that can solve nearly every problem. ML hype has washed over all walks of life - from business to academia. However, it has caused tunnel vision as businesses try to apply machine learning to context where other solutions would work better.

**BY JURAS JURŠĖNAS, CHIEF OPERATIONS OFFICER AT OXYLABS.IO**

IN FACT, SOME SCIENTISTS have already issued a warning about the overuse of machine learning with some coining the term “AI solutionism” as the (flawed) view that AI will be able to solve all existing problems in their entirety.

Tech business is no different. Expectations surrounding machine learning are still sky high. It is no mistake - machine learning is definitely an extremely powerful tool. But there are limitations to its application. In some cases, the old school rule-based or manual review approach will deliver better results with less financial burden.

## Statistical limitations in machine learning

Before we get onto the specific applications where machine learning is best aided with other approaches, we should understand that there are certain statistical limitations hidden within the process.

One of the most important limitations of machine learning in business is concept drift. Machine learning models are constantly evolving stochastic (all results are predictions with some randomness, not purely deterministic outcomes) processes. As the data changes, models will not be able to catch up, causing less accurate predictions over time.

Another statistical limitation is associated with data science and insight generation. Machine learning models can be used to “read through” colossal amounts of data between different sets in order to discover correlations. However, as the mantra

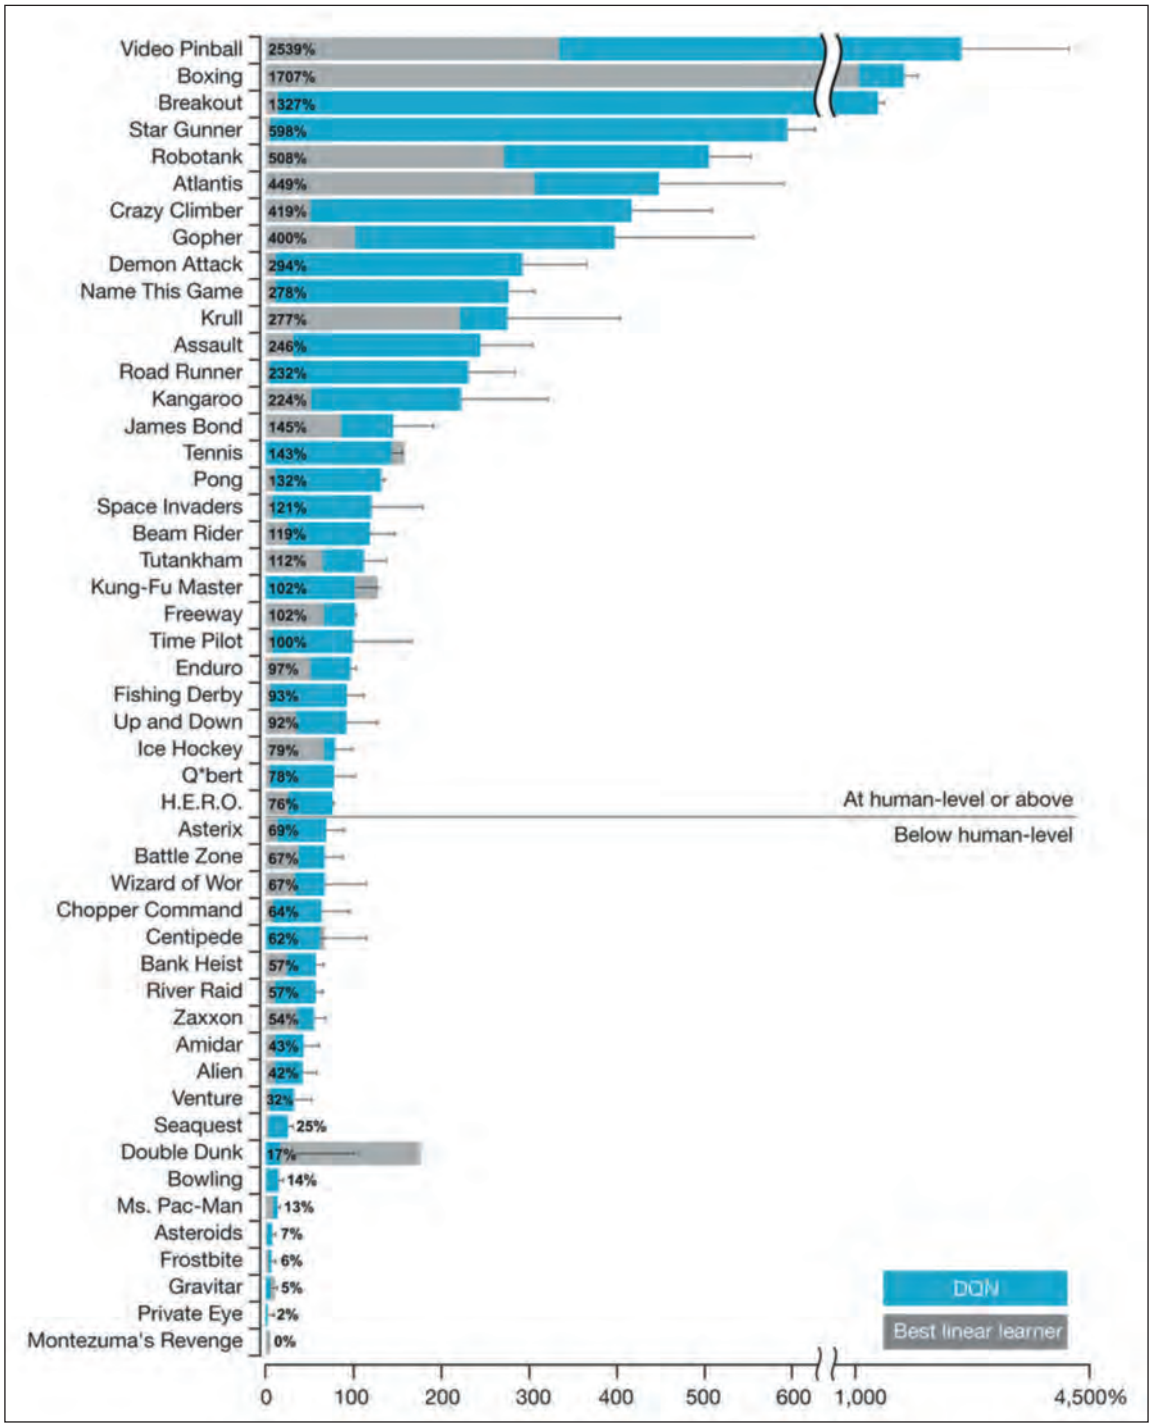


Deep Q-Learning machine learning models completely outmatch humans in certain games. Source.

goes, correlation doesn't equal causation. Without going in-depth into what the machine learning found and attempting to match the correlations with human reason, we would be risking going awry with interpretations, as not all methods are easy to understand, leading to detrimental outcomes.

Finally, due to partly design, partly computational limitations, reinforcement learning (RL) is limited as well. While it has been able to achieve seemingly impressive results in incredibly complex applications (e.g. video games such as Dota 2 or StarCraft), there are still huge drawbacks associated with optimization algorithms. Theoretically, RL models can achieve either a highly skilled application of one particular strategy or a decent application of several different strategies. How these RL limitations might cause future issues is another topic for another day.

## Machine learning limitations in business

Outside of the realm of abstract limitations, there lies the practical issues. In business, we are generally heavily constrained by financial and temporal constraints. Often, we are limited in what, where, and

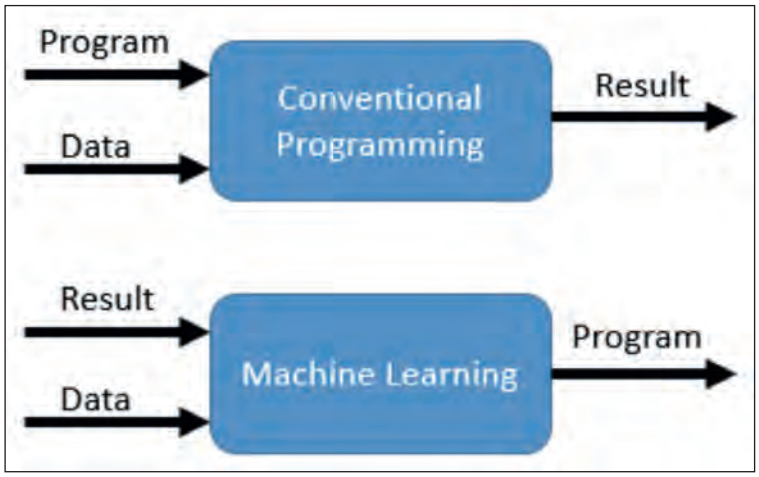


We don't really know what's going under the hood of a particular model. Source

how efficiently we can apply machine learning models. Even if we forego both abstract and practical limitations in machine learning, we would still have reason to apply manual or rule-based approaches. Humans reason in a completely different manner from machines.

Models can usually very accurately and quickly evaluate trends and historical data. They can then provide reasonable accurate statistical predictions about the best business decisions. However, they operate within a narrow domain and through a single layer of logic. Essentially, machines are detail-oriented-small-picture thinkers.

Humans are different. We can zoom out and take into account the entire view of the business. In many cases, a big picture outlook is necessary as business decision making isn't as simple as pulling up a few numbers and arriving at a conclusion. At least in many cases it isn't.

Additionally, machine learning models cannot transfer experience (or information) from one domain to another. While there are ideas circling about transfer learning, we have yet to arrive at a satisfactory way to do transfer skills and abilities between domains in machine learning. Humans, on the other hand, can support their decision making and abilities from a wide variety of domains of expertise through the use of heuristics.

Therefore, we shouldn't be aiming to solve everything through machine learning. Old school solutions are not overshadowed by ML. They are supplemented by it.

## Machine learning against phishing

Let's take a simple, practical example from the usual day-to-day business activities - emails. Everyone with some skin in the game has received a phishing email. However, we don't see most of them as they are blocked automatically.

Phishing emails seem like a great candidate for machine learning. They are plentiful, they look like the original but have key elements different or entirely missing, and they have few to some giveaway signs. All we need to do is label certain fields to check and the model will take care of the rest, right? Such an approach is far from optimal.

For one, phishing emails are a little like airplanes. Failure can yield drastic consequences. Thus, we wouldn't accept a 1% failure rate of an engine. Similarly, we don't want a 1% failure rate of anti-phishing processes as leaking data from just one account can lead to many quickly compounding issues.

Unfortunately, machine learning models are hard to bring over 99.9%. In fact, for most businesses and processes even reaching 99% will be outside the realm of possibility. Only true tech giants and academic institutions should even attempt to reach such heights of accuracy. Add concept drift into the mix and the challenge goes from nearly insurmountable to impossible.

As mentioned above, having a failure rate above 1% would be unacceptable in anti-phishing practices. Yet, training a machine learning model to 99% is already a hard task. Instead of getting too hyped up with machine learning, we should be supplementing it with other approaches.

## Integrating into rule-based and manual review processes

I might have sounded like a true machine learning sceptic. However, it's not all doom and gloom. Going back to the video game example, something I left out was that the developers had specifically set a communication "ping" and a delay between inputs. Machine learning models can achieve greater efficiency in certain areas (e.g. in reaction times) than any human ever could.

If machine learning models are better at some aspects of video games, the same will be true for other areas. Phishing will be no different. Machine learning models

Models can usually very accurately and quickly evaluate trends and historical data. They can then provide reasonable accurate statistical predictions about the best business decisions

HYVE
MANAGED HOSTING

SCAN ME

## The UK's Leading Cloud Hosting Provider

We are Hyve, your cloud experts. Combining our small team ethos with a passion for technology, we provide fully managed, global cloud hosting services.

We are serious about service and take pride in our commitment to personal support culture. Alongside our team of highly-specialised cloud experts, we guarantee scalability, security and unparalleled performance for your business.

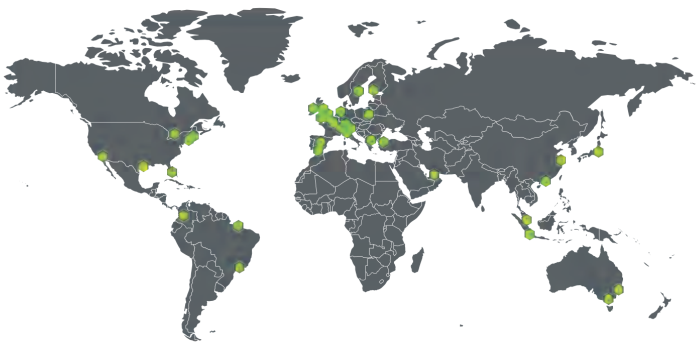

### Private **Cloud**

Created by cloud experts with over 15 years of industry experience, Hyve's Private Cloud provides dedicated resources for your organisation, ensuring the ultimate in security.

### Dedicated **Servers**

A Dedicated Hosting solution with Hyve gives you the autonomy to decide what is, or isn't, allowed on your server. We only deploy the best hardware and the fastest storage available.

### Enterprise **Cloud**

Hyve's Enterprise Cloud is a multi-tenanted cloud, providing the ease of scale and cost savings of a Public Cloud infrastructure, but with added security and monitoring.

### Managed **Colocation**

We provide fully managed Colocation on a global scale, spanning across 35 locations worldwide. You can maintain complete control of your hardware, whilst we take care of the management.

will be able to detect inconsistencies that could be difficult to spot for humans and vice versa. While it may be possible to develop some Skynet-level antiphishing model, businesses usually have other things to consider - mainly, the costs of such an undertaking.

Rule-based approaches have many benefits. For one, they have been utilized in numerous industries for a considerable period of time, allowing effortless access to best practices and domain-specific knowledge. Additionally, implementing a rule-based approach into (returning to our previous example) email protection is going to be tremendously more cost-efficient than training any machine learning model.

These approaches are also more transparent and accessible. Machine learning models are black boxes by design. Figuring them out takes considerable amounts of time. Compare the nature of these models with a rule-based system: even extremely complex rule-based systems can be quite quickly understood by someone experienced in that particular industry..

Additionally, the inclusion of manual reviews might be a necessity for some aspects of business (e.g. anti-fraud). Some regulations might require businesses to comply with consumer requests for manual reviews if something has been flagged automatically.

However, manual reviews aren't just for compliance with the law. They can be utilized to strike a better balance between red flags and letting things go through the system.

As mentioned above, machine learning models are clearly better than humans at a specific set of tasks while humans outperform models at others. Thus, mixing rule-based systems, manual reviews, and machine learning models would allow a business to maximize revenue and minimize costs.

Finally, rule-based systems and manual reviews will come in handy whenever a model needs new training data either due to concept drift or "bad" training data. Both can produce unpredictable results over time, necessitating a return to the tried and true practices.

## Conclusion

Instead of living through the hype attempting to solve every issue under the sun with machine learning, we could just integrate that model into our current processes. Processes that already have rule-based approaches and manual reviews in place. It's a lot like the, hopefully on the brink of the end, recommendations for protection against COVID-19. Taking a single precaution (such as washing your hands) does fairly little but combining all of them together creates a formidable aegis.

PURESTORAGE®

# Experience Pure FlashBlade Now

## Test drive modern storage

Let's Go

# Lessons from the kitchen:

## An introduction to Value Stream Management

The need for value stream visibility is not a nice-to-have, it is a business necessity whether you are running a kitchen or leading a digital transformation.

**BY ANDREW DAVIS, COPADO**

VALUE STREAM MANAGEMENT is the process of visualizing the stages through which work is performed (a value stream) and tracking metrics so you can plan improvements. Every process for creating value includes some activities that add value, other activities that do not add value, and other inefficiencies like waiting time. Understanding your processes in this way gives you the ability to target process improvements to remove wastes such as delay.

As an example, consider something like ordering food in a restaurant. When a customer arrives, they need to be seated. There may be a waiting period while staff clear and prepare a table. After being seated, the customer can order drinks, appetizers, and a main dish. Each request is sent to the kitchen. In the kitchen, one person may prepare raw ingredients, another may assemble and cook them, and yet another person may be responsible for plating them. The food is then delivered to the table, hopefully to the satisfaction of the customer. Customers eat, chat, and eventually leave. The staff can then clear the table and start again.

Every step in the process involves a hand-off between different people. Each step takes a certain amount of time, and the hand-offs need to be optimized to deliver the meal in a timely way, with high quality.
Collaboration, cost and complexity in a value stream
As this example shows, creating value requires materials to pass through the hands of many workers. Later steps in the process depend on earlier steps being completed successfully. And each step is necessary to deliver a final product. Added value – or more specifically – adding the most possible value – does not happen in a silo. It is an inherently collaborative effort.

There is an implied cost to each part of the process. Staff need to be paid, space and equipment need to be available and ready, and if the meal is delayed too long customers might become frustrated or leave. In a kitchen, the whole process is visible and unfolds over a short period of time.

Value stream management is becoming increasingly popular in software development to understand this complex process. But in creating software, most of the process is invisible, many tools or even teams are involved in the work, and delivery of a single working piece of software might take days, weeks, or even months.

Value stream maps are a tool to make this flow of work visible. They enable businesses to understand the flow of value to customers, and waste incurred across a process that may otherwise be hidden from view.
There is one important difference between the kitchen analogy and software development. Restaurants offer a set menu to ensure they have the right ingredients on hand and can create all the dishes with high quality. Since software can be replicated infinitely, recreating the same piece requires no effort. Anyone

To increase the flow of value, businesses (especially enterprises) need to make it easier for development teams to deliver software, eliminate bottlenecks or delays, and cut waste

who needs a piece of software that has been written before can just purchase or duplicate a copy.

The challenge in software development is in creating something that has never been made before. Product development is a unique challenge every time. But the overarching process still follows predictable processes that can understood and monitored to increase the chances of success.

## The Value Stream Management Imperative

When applied to software development, value stream mapping has the potential to drastically enhance productivity, quality, and time to market. Gartner predicts that 70% of organisations will use value stream management to manage their DevOps pipelines by 2023.

It is easy to understand the interest in value stream management when you consider how fragmented software development teams have traditionally been. As businesses become increasingly dependent on software creation, they increasingly seek to improve the effectiveness of technology teams, especially through facilitating collaboration within them (and with the rest of the business).

To increase the flow of value, businesses (especially enterprises) need to make it easier for development teams to deliver software, eliminate bottlenecks or delays, and cut waste.

## Getting Started with Value Stream Management

The first step in value stream management is to create a map or visualization of the process. This can be done on paper, using digital whiteboard tools, or in software specifically designed for creating and managing value stream maps. The goal is to map every major phase or sub-process work must flow through before delivery.

Metrics such as lead time, cycle time, load, and percentage complete and accurate can then be layered onto every stage of the process to help indicate the speed and quality of these processes. Changes in these numbers over time indicate the evolution of the process and may show increasing or decreasing performance. This information is crucial for getting a holistic view of how your development team works together and identifying areas for improvement. This quickly identifies not only the interactions that add most value, but also reveals opportunities for automation as well as bottlenecks to performance. Why don't businesses know this already?

In a world that is now dominated by digital, it is often surprising that businesses do not have this level of insight into their software development already. In a modern business context where software defines so many processes, issues at the 'micro' level of the development team are often amplified and impact the entire business. So how is it that businesses do not already assess software development at this level? This information is often hidden because people are too busy doing their work to take time to observe and understand their processes. There are also gaps in domain knowledge, where the people most familiar with technical systems may lack knowledge about business process improvement or organisational design.

There can also be a lack of precise understanding of where inefficiencies lie. The team and managers may have a 'rough idea' of plans and progress but lack a tool that can accurately identify measures such as delay times, frequency of rework, and the labour required for each stage. There is often little perspective on the entire, end to end process. There can also be cultural challenges, where teams may blame others for difficulties or inefficiencies, or be reluctant to be inconvenienced by changes that would benefit the organization. A series of critical issues, or an overwhelming pressure to improve can force teams to begin the work of understanding their own process and questioning their current understanding. Such challenges can provide opportunities for deep change.

Structural inefficiencies can be particularly acute for businesses with larger IT teams. Larger teams provide more opportunities for clumsy handoffs and are more challenging to fully understand. Realistically, even just within a software development team a business cannot make improvements across all indicators at once. Identifying the single primary bottleneck in a system is key to showing you where to start your efforts.

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. In this context, the need for value stream visibility is not a nice-to-have, it is a business necessity whether you are running a kitchen or leading a digital transformation.

# Navigating the SASE Maze

SASE is causing a buzz right now, and for good reason.

BY JONATHAN LEE, SENIOR PRODUCT MANAGER, MENLO SECURITY

YOU WOULD HAVE to rewind to 1726 to find the world's first office. Originally the result of the British empire expanding and engaging in growing levels of trade, it has become the template from which the vast majority of white-collar work has been completed for the near three centuries since.

Yet, since 2020, the traditional role of the office has been called into question.

The pandemic has had a profound impact that will shape the operations and models of organisations around the world for decades to come. In the case of the workplace, these were previously considered to be vital hubs of productivity; the defining cores of organizational culture and a key factor in securing talent and business.

However, as vaccination campaigns around the world begin to gather momentum, offering up the key to unlocking societies, communities and economies once more, it seems unlikely that many of these buildings will revert entirely to their original purpose. During the UK's initial lockdown, the number of those working exclusively at home rose eightfold from 5.7% in January/February 2020 to 43.1% in April 2020.

And as we muddle through 2021, it's becoming increasingly clear that this "new normal" is, at least in part, here to stay – not just in the UK, but globally. According to Strategy Analytics, the mobile workforce around the world is anticipated to be 1.82 billion strong come 2022, which will account for 42.5% of the total global working population.

It seems, therefore, that hybrid and remote working models may well prevail. But within this shifting physical to digital dynamic, various considerations need to be made to maximise the benefits of such operations.

It is not a case of simply flicking a switch – for companies to continue operating successfully on a physically disparate basis, the right preparations and infrastructure needs to be implemented that can support the challenges of remote working.

## Security is one such challenge, and of paramount importance to any business

Let's consider the context: Many companies were forced to turn to the cloud almost overnight in order to sustain operations remotely, and since have reaped the rewards, their employees able to conduct work and access productivity enhancing applications with ease.

Ordinarily, with so much sensitive information hosted on these platforms, security should not just be a consideration, but rather the front and centre of such a cloud strategy. There are signs, however, that it has all too often been overlooked. Where enterprises have shifted to cloud operating models, they have retained the same security practices and policies that were designed for the physical office setups, often supported by on-premises servers – policies that aren't adequate in a remote arena.

Indeed, a lack of coherent cloud-based security has created vulnerabilities. Somewhat haphazardly, many companies have been relaying network traffic between multiple checkpoints like firewalls, interrupting traffic flow and creating unwanted exposure.

Some have turned to virtual private networks (VPNs) as a resolution, but these are flawed. Not only do

they fail to scale, but they can easily create traffic bottlenecks, adversely impacting productivity and compromising security.

## So, what's the solution?

### Secure Access Service Edge

A term coined by Gartner, Secure Access Service Edge (SASE) refers to the simplification of networking and security that is achieved by delivering both elements as a cloud service to the source of a connection directly, rather than via an enterprise data centre.

It entails the integration of software-defined wide area networking (SD-WAN) capabilities with network security functions such as CASB, Cloud SWG, ZTNA/ VPN, WAAPaaS, FWaaS, DNS and RBI. In doing so, and potentially by integrating with 5G as well, it provides organisations with the opportunity to create frameworks that support today's dynamic, secure access needs.

From a more technical perspective, it relies on a distributed group of cloud gateways, also known as local points of presence (POPs), that receive traffic from other locations running SD-WAN devices. Within each POP, all security functions and policies – be it web and email security or firewall and access control – are implemented.

Within this framework, security becomes an extension of the user – their data and applications – while visibility and control are maintained regardless of location or device type. Unlike VPNs and other legacy solutions, SASE has been built with a cloud-first mindset.

Instead of trying to force a square peg into a round hole, resulting in failed remote or hybrid working models suffering from hampered SaaS adoption, it provides complete, seamless protection while equally prioritising productivity.

## Indeed, the benefits that it offers are noteworthy

Take data, for instance. SASE's integration of networking and security capabilities allows organisations to guard against sophisticated threats while reducing the potential for unplanned data loss. Productivity is also increased, SASE seamlessly enabling employees to work wherever they are and use the applications that they need without delay or failure. In turn, frustrations with potential bottlenecks that could cause employees to attempt to find workarounds will be eliminated, ensuring an organisation is not inadvertently put at risk.

It is said that SASE is becoming much more of a business enabler for this very reason, helping to move security away from its traditional stereotype as a roadblock in many organisations and towards a new status as a facilitator, allowing companies to capitalise on the advancements of cloud and SaaS applications without having to enforce changes to user behaviours. SASE isn't limited to a desktop or laptop either. Mobile devices are often overlooked in security protocols, but it is vitally important that they are protected – these are also potential gateways for hackers. With an integrated cloud security solution that focuses on the device, protective security.

experiences can also be rolled out on more portable devices that have become part and parcel of remote and hybrid working environments.

## Should I consider SASE?

SASE is causing a buzz right now, and for good reason.

Be it bolstered security, enhanced productivity or an improved user experience, enterprises leveraging it have been empowered. There are some challenges which need to be overcome before SASE can be implemented. The cloud platform needs to be smart, dynamic and scalable enough to deliver secure access to resources, for example, no matter where an end user might be located. Therefore, putting the time and resources into ensuring you partner with the right vendor will pay dividends in ensuring a smooth transition.

Further, it is unlikely that the best approach for businesses would be to buy all the different components of SASE from a single vendor. While a limited number of companies do offer an all-in-one SASE solution, such may result in vendor lock in and significant compromises. Multiple vendors may, therefore, need to be consulted in order to achieve an effective, fully integrated SASE-ecosystem.

What is certain, however, is that for many, this transition could be the key to unlocking their business's potential in the new, hybrid or remote working normal. As of 2021, according to the latest CyberEdge Cyberthreat Report, 74% of IT security decision makers are currently adopting technology capable of supporting SASE architecture, with these foundations likely to lead down this path.
While Gartner had originally made predictions that it would take 10 years for SASE to become mainstream, the pandemic has accelerated adoption by between three and five years, meaning it has very much become a reality in the here and now.

Many are looking to adopt it much quicker and embrace its sweeping benefits. Those that do will be well positioned to advance their business on hybrid working platforms for years to come, transforming their organisations into more competitive and exciting prospects.

Those that don't, however, could be faced with an uphill battle against limited security for users and stifled productivity.

# The revolution in IT service management is a big bonus for boardrooms

Powerful new allies have emerged for boardrooms looking to streamline IT service management and reduce costs.

**BY MARK TWOMEY, CEO, XCESSION**

A NEW BREED of high-performance IT service management (ITSM) solutions can deliver 80 per cent of the functionality of tools from the more established names, but at only 20 per cent of the cost.

As enterprises regroup after the pandemic and take a fresh look at costs, they need fresh ITSM toolsets to adapt to the changing models of work. Hybrid and remote working look as though they are here to stay, placing new demands on IT departments which must cope with a more extensive infrastructure and added requirements for greater integration with supply chain partners.

Many enterprises, however, will continue to rely on the same familiar names in the ITSM software industry. They will buy software that has masses of functionality which their IT departments never fully utilise and which performs poorly when organisations seek closer integration with partners' systems. Research

has found that 40 per cent of enterprises have ITSM solutions that remain on the shelf. Many organisations repeatedly pay for software with multiple bells and whistles when what they really need is something sleeker and more purposeful that fits in with how the business world is changing.

## Time for businesses to reassess what they need from ITSM

These changes in business practice – the growth of hybrid working and the increasing interconnectedness of supply chains – demand organisations have the most efficient and cost-effective management of software, hardware, devices and networks available. It is why businesses should reassess how much of their current ITSM software they use and then form a detailed view of what they need, using ITIL IT service management framework checklists or templates as guides. The list of modules an enterprise could implement is long, but the reality is many will never be required. For example, availability management, capacity management, event management and finance management are offered in solutions by the industry's big names, but enterprises seldom use them.

## Focus on integration with supply chain partners for organisational agility

As organisations refocus after Covid, speed of innovation, agility and supply chain integration will be vital attributes, built in large measure on interoperability with suppliers' and partners' systems. A large organisation may have thousands of suppliers in different corners of the globe, so integrating with their ITSM systems should be simple, following one of the SIAM (Service Integration and Management) models.

With established platforms this tends to be clunky and less than intuitive. Newer platforms, by contrast have pre-packaged integrations that plug into a supplier, enabling IT to on-board and off-board partners, suppliers and outsourcers quickly. Plug-in integrations allow tickets relating to printer interfaces or server provisioning to pass swiftly between different organisations in the supply chain without anyone dropping the baton.

The newer, sleeker solutions are also easy to implement. Vendor-agnostic implementers can swiftly knit together new solutions with existing tools IT professionals want to retain, providing a customised platform that fits the desired purpose. This can take just months, whereas platforms from high profile vendors typically require two or three years to achieve full implementation in a large organisation.
It is also much easier to add on functionality relating to HR, facilities or legal. The newer solutions accommodate the need for confidentiality and segmentation of sensitive data, which is especially important when dealing with outsourcers.

## Flexible ITSM pricing better suited to today's businesses

A boardroom may ask how this software can be available at much lower cost than the higher profile brands. In many cases, it is because the disruptor companies have very different pricing models that are far more flexible. Some spend only a fraction on sales and marketing compared with the big-name vendors. A vendor such as 4me, for example, expands revenues by increasing its market share through the network effect, rather than by ramping up fees or through heavy expenditure on marketing.

## Independent advice makes choosing a new ITSM vendor easier and more effective

It is worth remarking that many of the newer companies have outstanding track records. The global technology consultancy Gartner[1] places newer vendors at the top of the ITSM list for customer choice, as revealed in testimonials.

With such a wide field of new ITSM software providers, selecting the right solution may seem difficult. This is best achieved through a vendor-independent implementer not involved in sales-driven contractual arrangements. This is more transparent and avoids the pitfalls of the standardised one-size-fits-all approach, but without over-elaborate and unnecessary customisation. Independent implementers are also more likely to be honest when clients propose sub-optimal integrations or combinations of tools.

Experience in implementation also counts for a great deal, even with the newer solutions that have plug-in integrations. Enterprises that require enhanced employee self-service functionality or automation in their ITSM platform, for example, can easily run into difficulties with data management. And in every organisation undergoing change there are cultural barriers to overcome. An experienced implementer with organisational knowledge, as well as technology know-how is more likely to spot such problems and resolve them before they develop.

Boardrooms may hesitate to switch from a well-known vendor, but they cannot ignore the fact they are in all likelihood paying much more than they need for solutions that are either overladen with functionality or far from best-of-breed in the post-Covid landscape. It is important, therefore, that they seek the most objective and expert advice available so their organisation benefits from new approaches that will save huge amounts of cost and provide them with greater agility and efficiency.

## FURTHER READING

➤ 1 https://www.gartner.com/reviews/market/it-service-management-tools -End

# The new era of Digital Experience Management

Throughout the past year, the series of lockdowns and continued uncertainty have brought new focuses under the spotlight. Indeed, the pandemic has not only led to a preoccupation with public health, but has also dramatically altered the way people work, live and interact, with digital experiences brought to the fore front. Naturally, this has influenced customer behaviour.

BY ASH FINNEGAN, DIGITAL TRANSFORMATION OFFICER, CONGA

iDIGITAL CHANNELs and engagement have accelerated immensely, as companies have had to shift their mode of operations to accommodate the current environment. For example, leading retailers have had to rapidly embrace digital transformation initiatives to replicate the traditional, in-store experience. They have looked to digital solutions to enable customers to shop far more efficiently and effortlessly, with hope of establishing a smooth, end-to-end experience. According to McKinsey, 80 percent of businesses believe that their core business model should now be entirely digital in order to remain economically viable.[1]

In the current landscape, speed is considered a top priority. Customers want products and services in an instant, and the same can now be said for businesses. Within B2B, expectations have never been higher. Companies require solutions fast and expect information even faster. The more successful service providers and customer experience (CX) leaders have reviewed their product offerings with their channel partners, keeping a real-time pulse on customer and business trends. In fact, they have probably had to consider entirely new solutions and completely redesign CX journeys as expectations or circumstances change. The current challenge for companies is to understand which trends will last long term.

## Digital transformation: hype cycle vs reality

Digital transformation (DX) has steadily become an everyday buzzword since the start of the pandemic. Companies have relied on service providers to deliver cloud solutions and help them advance in their digital transformation journeys to ensure they can continue to provide their services to valued customers. As revealed in a recent report by Gartner, 69 percent of boards of directors accelerated their digital business initiatives in response to the pandemic, with the sales process and customer experience being some of the key drivers.[2]

The pandemic has been a catalyst and encouraged this 'hype cycle' with regards to digital transformation, and it has become difficult for customers to separate the hype from reality. As the pandemic has shown, digital transformation is very much a process, not a race. Whilst COVID-19 has accelerated 71 percent of companies' digital transformation plans, only 36 percent are considered somewhat successful. Many of these initiatives are rushed.

Most companies simply approach CX transformation projects all wrong, picking a technology and implementing it at speed, with no real idea of how this will improve their customer journey or digital experience. In fact, vendors often provide a 'unified solution' as opposed to focusing on and providing real solutions for their customers' business. When designing a digital transformation or cloud strategy, it is important to consider the initiative from all angles, as well as consider all parties involved. Unifying an experience by establishing a seamless operation model improves the workflow for teams, but leaders need to consider those vital touch points – those key moments of interaction with the customer – and consider how they can perhaps be improved. This will involve reviewing the operational model, removing any pain points, and identifying the needs of the business, establishing real objectives.

## Future cloud strategies require 'hyper personalisation'

With all of this in mind, vendors and channel partners have a responsibility not to push the latest or most fashionable technology solutions onto their customers, but to actively solve the challenges that they are facing. If improving digital experience management and the overall customer journey is the objective, then the service provider should consider themselves as a consultant, working with their customer to devise a cloud solution that is specific to their business – a 'one size fits all' approach is no longer viable. Offerings should vary from industry to industry, and customer to customer.

Moving forwards, vendors simply need to be more flexible, and perhaps consider tailoring their product offering to ensure that they are responding to their customers' ongoing concerns and needs. Naturally, this will require more technical solutions and perhaps even collaborating with another supplier to ensure their problems are indeed solved. Matchmaking a customer's business with the right technology, regardless of who or what it is, will also be an imperative. Either way, both parties should prioritise a more unified experience for the end user, and ensure all deliverables are met.

## How to approach a CX programme

In reality, when designing a new CX programme, companies need to establish clear business objectives from the start, before considering any complex cloud solution or adopting any transformational technology. This will likely involve reviewing their own operational model and reconsidering all the touch points across the business or sales cycle.

Organisations will need to review CX programmes on a regular basis with their chosen service provider – what works on one occasion may not later down the line – and as COVID-19 has proven, customer expectations can change rapidly. By reassessing their operational model, companies will be better prepared for all outcomes.

With regards to digital experience management, offline and online processes will need to be far more consolidated and aligned. As businesses look to build their omnichannel suites and elevate the customer journey to the next level, each touchpoint or process needs to be considered very carefully, with teams prioritising the most important engagements.

## FURTHER READING


- 1 https://www.mckinsey.com/about-us/covid-response-center/mckinsey-live/webinars/customer-experience-in-the-next-normal-after-covid-19
- 2 https://www.gartner.com/en/newsroom/press-releases/2020-09-30-gartner-says-sixty-nine-percent-of-boards-of-directors-accelerated-their-digital-business-initiatives-folloing-covid-19-disruptions

# Why Low Code is the perfect business transformation cherry picker

In December, IDC predicted that global digital transformation investments will total $6.8 trillion between 2020 and 2023, with 65% of the world's GDP digitalised by 2022.

**BY RUTH WEATHERALL, CO-FOUNDER AND OPERATIONS DIRECTOR, UP3**

IT'S FAIR TO ASSUME that any organisations not going through some form of technology enabled transformation are today firmly in the minority. However, according to McKinsey, 70% of digital transformations fail, suggesting that intent alone is no guarantee of success.

We help organisations scope and implement transformative projects using low code and the challenge we experience is too much slow and steady and not enough opportunism. Most organisations have large scale change projects that think big and have the potential to tackle numerous objectives under the broad banner of 'digital transformation'. But these may be hard to get off the ground, could be slow to adapt and the more stakeholders you add into the mix the higher the risk of inertia. By the time it's complete (if indeed it does even get completed) some business requirements may have already moved on.

Smaller projects may be less exciting on the CV or seem unambitious when getting executive buy-in, but they allow the opportunity for tightly focused programmes that deliver results quickly. An over-focus on strategic transformation projects involving huge scoping and planning can be at the detriment of quick wins needed to build groundswell and momentum.

The fixes that are needed urgently tomorrow get missed, and over time, appetite for transformation wanes while cynicism grows. When the pandemic came along and forced organisations to implement digital applications in days rather than months, it was a lesson for many that not everything you want to change needs to take forever. It also thrust low code and 'citizen development' into the limelight.

Low code platforms, such as ServiceNow and its App Engine solution, absolutely provide the framework required to build and deploy transformative apps quickly without huge amounts of technical knowledge. In fact, low code has the power to enable organisations to 'cherry pick' the crucial quick wins and benefit from them to build more transformative momentum than any top down strategy alone can.

Although big transformation projects are still essential, changing and improving the way people work in small and focused ways gets different business functions thinking more creatively about change and actively contributing more ideas to IT.

The dangers of digital inertia aren't simply that things won't get done. With technology becoming increasingly democratised, there's also a high risk that people will go and find their own solution, leading to a long blind alley of shadow IT that is very hard to reverse out of. Low code platforms provide an antidote to this by giving organisations an approved platform that is easy to use, on the proviso that good governance and planning are in place too.

Although low code platforms offer the means for rapid transformation, they don't find the solution for you. Organisations still need to locate and isolate the problem, however niche, to be able to create the necessary fix. Low code, coupled with business analysis and implementation skills, gives IT the ability respond quickly to urgent needs while still pushing the company-wide transformation agenda forward. It helps IT counter feelings of disenchantment, because the big promise of improvement seems never to happen.

## The ideal low hanging low code project

There are various projects that are ripe for rapid transformation through low code development. They say necessity is the mother of invention, and in this context, this is commonly true. The greater the urgency for genuine change, the more suitable the project (and low code as the means).

COVID-19 triggered many such projects. For one train network operator, as the legislation and responsibility for tracking COVID cases in organisations evolved, we were asked to build a company track and trace system in a matter of days. One person had the unenviable job of manually checking shift patterns to track who members of staff testing positive for COVID had had contact with, and then communicating the need to isolate.

We worked with the customer to scope the process and that thorough discovery allowed us to fully understand what they required. We then utilised applications we had already developed combined with low code to quickly deliver a solution. This removed this stressful personal burden from the individual and ensured there was no backlog in contract tracing. Driven by necessity, low code provided a quick and effective route to transformation.

Another major requirement for low code apps in 2021 could be the change driven by permanent remote and flexible working. Although many quick solutions were found during lockdown, now that restrictions are lifting it has become clear that many people will work either fully remotely or a mix of office and home. This is likely to present further challenges that must be solved quickly to ensure continuity and productivity as organisations grapple with how to effectively manage hybrid working.

Other key triggers include new legislation and regulation changes (anything involving a hard deadline). Exposure to risk is also a common motivator for low code projects, particularly around fraud; financial fraud rose 33% during lockdown in 2020 for instance.

## Speed and effectiveness

While low code can help to deliver transformative apps very quickly, it is crucial that speed doesn't dilute the effectiveness of the project. Low code should be about delivering what is needed and when it is needed. The applications must be robust and support IT governance. Often, projects require super-quick prototypes to demonstrate to the business why funding for an application is required and what can be eventually delivered.

This is often the case when we're contacted by organisations wading through in-depth discovery and strategy imposed by expensive consultants that are simply outsized for the scale of the job. They are asking if there's another way and if low code can provide the answer. We're also seeing far more demand coming from organisations wanting to innovate with apps in stages because time to value is critical. It doesn't have to be perfect first time. Low code is effective in this regard, as you can develop 'lite', deliver value quickly and build on that.

## What low code isn't

For anyone attempting to use low code platforms, it's just as important to recognise what it cannot do. Low code allows skilled technical people to build apps quickly. It's not a guarantee that your app will deliver on its promises. Nor is it something that anyone can jump into on a whim and start building.

Any organisation first needs to implement a low code platform to give it a standardised framework for app development. This is essential as many of the smaller automations and apps draw data from across the organisation so it needs to work within a governed IT environment and infrastructure.

Organisations also need to ensure that they understand the requirements of any app clearly before starting. This is about the age-old IT issue of pinpointing the business problem before starting to build. Rapid business analysis is just as crucial as a low code platform for rapid development. Organisations need partners who have the experience of rapid implementations and know what questions to ask during high energy workshopping sessions.

They also need honesty about what is and isn't achievable, and about price. It's a case of more haste, less speed and ensuring that a premium isn't being paid for every quick win required.

Projects that go awry and don't deliver rapid and effective results are those where the scoping hasn't been appropriate, however good the developer or low code platform is.

# **Integration 101:** How to solve the biggest unsolved challenge of IT

Only with a flexible integration layer built on the principles of API-led connectivity and reuse will today's businesses be primed to become the digital innovation factories that tomorrow demands of them.

BY IAN FAIRCLOUGH, VICE-PRESIDENT OF SERVICES, OFFICE OF THE CTO, EMEA, MULESOFT

IN TODAY'S CONNECTED ECONOMY, software rules the world. Market leadership is increasingly driven by the speed at which organisations can digitally transform to deliver new applications and enhanced experiences for their customers. This trend has been amplified by the COVID-19 crisis and will continue to gather momentum in the years to come. Indeed, IDC predicts within five years, nearly two-thirds of global enterprises will have become digital innovation factories, deploying new application code on a daily basis. However, as the demand for new software experiences increases, IT departments are struggling to keep up.

The Connectivity Benchmark Report 2020 highlighted that 85 percent of global organisations have significant integration challenges, which are stalling their digital innovation drive. Of the almost 900 different applications in use across the average enterprise, just over a quarter are integrated, leaving data trapped in silos. This makes it harder to create connected customer experiences and collaborate with external partners effectively. What's more, just two-in-five global enterprises completed all of last year's digital transformation projects in light of the challenges they faced. Clearly the status quo is unsustainable. Traditional IT operating models are broken, and organisations must find new ways of working to accelerate project delivery.

## IT's biggest unsolved challenge

Integration has been one of the biggest headaches for IT departments since before most of us can remember. The problem goes back to the way organisations connected applications and databases before modern integration platforms existed. The most widely used approach was to create point-to-point custom code, which was embedded directly into the application or database that needed to be connected. This effectively built a bridge between two systems, enabling data to cross from one to the other so information could flow freely across the organisation to deliver digital services.

However, while this worked to a reasonable extent during the early days when IT systems were less

diverse, dynamic and complex, today it creates more problems than it solves. Custom integration code usually results in a tight coupling between applications, databases and devices, creating close dependencies that are difficult and even risky to untangle. If organisations want to introduce a new application or enhance digital experiences, they often need to break apart existing integrations to 'rewire' things, which can impact service functionality in unexpected ways. This makes it incredibly difficult and costly to introduce the change, and almost impossible to achieve the speed that digital transformation demands.

## The only way is API

Custom code is clearly no longer a viable approach to integration. Instead, organisations should embrace API-led connectivity, putting an application programming interface (API) in front of the systems they need to connect. This creates a loose coupling between applications, data and devices, so changes can be made quickly without impacting on existing integrations or the functionality of digital services. It therefore becomes easier to accelerate innovation and deliver new products and services faster, without increasing the risk of business disruption or spiralling costs. However, APIs should only be built where they add value. For example, SaaS applications often come with standard REST APIs that can be used to connect them to other applications and services, removing the need for organisations to build their own custom API.

Organisations can take these strategies even further by encouraging self-service and reuse of both internal and external APIs. This enables them to become much more responsive to regulatory and market changes, and able to innovate faster to unlock new revenues. The best way of achieving this is through API-led connectivity, which creates a composable enterprise with a flexible integration layer that drives the complexity out of digital transformation. IT assets and capabilities can thereby be more easily consumed and reused by employees, partners and third-party service providers, enabling them to become citizen integrators who create their own services and digital experiences without needing to write a single line of code.

## Building a composable enterprise

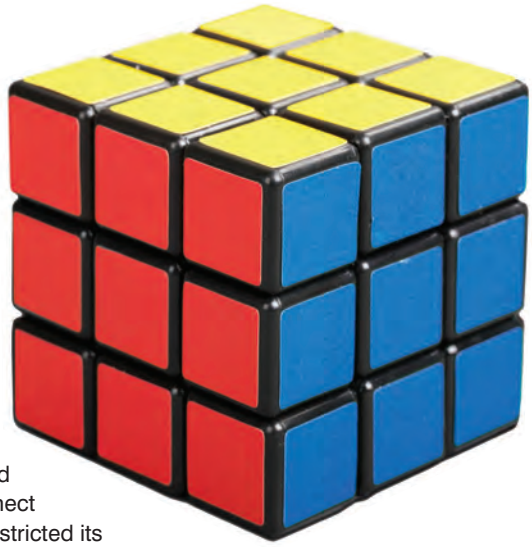

UK mobile-only provider Atom Bank is one such organisation that has realised the benefits of this more sustainable approach to integration. In the bank's early days, it used tightly coupled, point-to-point integrations and custom code to connect its systems, which restricted its ability to create new digital experiences quickly. Atom Bank embraced API-led connectivity to shift towards an agile, scalable operating model. It now deploys half the number of APIs compared to the number it had to develop previously, and because every asset is reusable, its development cycle has been shortened from one month to one day. As a result, Atom Bank is able to innovate at the speed the business and its customers demand.

All the evidence points to the simple truth that the demand for digital innovation is only set to grow. The ongoing crisis has raised the stakes even further, driving greater demand for digital experiences, and accelerating the transformation process within many organisations. Given that they're already struggling to keep up with the existing pressure, it's imperative for IT teams to find a new and more sustainable approach to answering IT's biggest unsolved challenge if they're to put their business in a position to lead and succeed in the future.

The tightly coupled integrations of the past are simply unsuited to the demands of modern digital businesses. Only with a flexible integration layer built on the principles of API-led connectivity and reuse will today's businesses be primed to become the digital innovation factories that tomorrow demands of them.

# Is your network holding up your digital transformation?

If you're wondering why digital transformation still isn't happening and you've interrogated all the usual suspects, maybe it's time to ask a few searching questions about your cloud network. It may be holding up the future.

BY MATT HOAG, CHIEF TECHNOLOGY OFFICER WITH KOCH BUSINESS SOLUTIONS

EVERY BUSINESS will have its own take on what's holding back digital transformation - development costs, skills shortages, disappointing return on investment, legacy systems and vendor lock-in are just some of the usual suspects.

That list hasn't changed in twenty or thirty years, but we need to add a new item: cloud networking – the problem of building end-to-end network services in the cloud. From my perspective, this is the missing piece of the cloud puzzle.

## First a Bit of Relevant Background

Koch Industries is one of America's biggest privately held companies, with annual revenues estimated by Forbes of $115 billion.

The organisation has around 120,000 employees across 700 sites in 70 countries and after several major acquisitions Koch has a complex network infrastructure, including seven global networks, thousands of applications, thousands of routers and firewalls, tens of thousands of access points and hundreds of thousands of switch ports.

As chief technology officer with Koch Business Solutions (KBS), our IT services arm, my job is to make sense of this global spaghetti, including the security of data and the performance of the applications on which the business and its customers depend. Ten years ago, we had a traditional network with MPLS connections between data centres, limited bandwidth to most other sites and all Internet access

centrally controlled. Life was simple and, for many of our users, pretty slow. At Koch, we believe in constant transformation and the need to disrupt ourselves. That's how we stay successful and grew to the global, diversified technology-driven company we are today. Staying true to our transformation vision and realizing the imperative to stay relevant in the digital age, we needed to think and operate our businesses differently.

The KBS team was tasked with preparing our foundational technologies, including the global network, for the rapidly growing digital business domains. We developed plans up and down the IT stack to decouple our businesses from the data centre and to implement cloud-native capabilities.
Our first major decision was to implement a global SD-WAN. This enabled a step-change in network performance – yielding a seven-fold increase in bandwidth across 500 locations.

At this time the cloud was already growing as a force. We settled on AWS as our preferred cloud partner, but soon realised that we weren't going to be able to run production workloads over simple VPN connections. So we developed a vision for cloud on-ramps dubbed "transport hubs" that we planned to implement first as physical connections to the cloud for data centres, then as virtual hubs for improving the experience of branches and remote users, and finally as transport hubs as a service, which would be capable of rapid deployment at scale.

The version 1 hubs provided high-speed connections for data centres to AWS, but approval, installation and testing took 18 months. The hubs also cost more than $2 million to put in place.

In order to be able to break ties with the data centre and improve application performance and user experience, we needed to replicate our physical transport hubs virtually in the cloud and allow users and sites to access those workloads through the SD-WAN.

Version 2 of the transport hub took "only" about six months – a great improvement over the data centre hubs. Extending the network globally would require multiple hubs, and when we wanted to connect to other public clouds, the hub count would start to climb. Even allowing for learning gained with AWS, expanding to other clouds would be months of work in each case.

While we'd solved some business and technology challenges we found that we'd created a highly complex set of solutions. Our network extended to only one cloud and we had no idea how to get to transport hub number 3, our vision of transport hubs as a service. After scouring the market for a solution, we came across an interesting start-up named Alkira. We had history with the founders, Amir and Atif Khan,

In order to be able to break ties with the data centre and improve application performance and user experience, we needed to replicate our physical transport hubs virtually in the cloud and allow users and sites to access those workloads through the SD-WAN

whose previous venture Viptela had been our SD-WAN partner. The network cloud provider's platform appeared to offer the seamless connectivity and full integration of higher-level services we were seeking. As a pure software proposition, delivered as a service, it also promised three huge operational advantages:

1. **Simplicity:**
   Network teams wouldn't have to learn every new capability in every cloud.
2. **End to end visibility and governance:**
   A single "pane of glass" to view all the traffic and transactions wherever they are on the network.
3. **Delegated operations:**
   The organisation doesn't need to build, own and operate network assets and advanced services.

In the autumn of 2019, we got our hands on a preview version of the provider's platform. What had taken us over two years to evolve in our transport hub version 1 and version 2 deployments we were able to recreate with the software in a single afternoon.

We had gone from a mass of complexity and months of work to a dashboard that allowed you simply to draw your network and deploy it in a few hours. This platform is enabling us to greatly reduce the amount of virtual infrastructure we run, manage and pay for. It should also simplify operations in cloud networking environments. We're no longer dependent on a handful of wizards who know where all the magic rings are.

However, the really critical change is the time we're able to save in connecting to different clouds. Building a production-ready on-ramp to Azure would have taken three to six months using the transport hubs in Koch's previous system architecture. With the new solution we were able to stand it up in a single day. The Azure connection was in place before the application teams were ready for deployment. The network team is no longer the long pole in the tent. If you're wondering why digital transformation still isn't happening and you've interrogated all the usual suspects, maybe it's time to ask a few searching questions about your cloud network. It may be holding up the future.

# Moving from a traditional to dynamic **public cloud infrastructure**

Just as there's no "perfect" cloud set up that works for all businesses, there's no single way to move from a traditional to a cloud infrastructure.

BY DARREN PURVIS, CLOUD ARCHITECT, BRIGHTSOLID

THERE ARE SEVERAL BENEFITS to public cloud adoption, but only if the move is done correctly. It's not enough to decide you're going to move to the cloud; you need to consider how you will do so. With some rare exceptions, moving to the cloud using a lift-and-shift method is unlikely to be the right choice. Instead, moving to the cloud is an opportunity to transform and migrate workloads. However, this also brings with it a number of questions to ask, to ensure you do it the best way for your business.

## Is the public cloud an appropriate choice?

With so many cloud options – from public, private, or hybrid cloud to software, platform, or infrastructure as a service – it's easy to assume there is a simple choice. But, for many organisations, the public cloud is the right decision.

When the public cloud was originally conceived, it was designed to work on a dynamic, scalable basis. This is particularly relevant for businesses that are growing and need to add resources on demand: using the public cloud makes it easier to grow without requiring significant changes to operations.

The exception may be organisations in highly regulated industries such as finance or healthcare, where certain security needs must be met and prioritised.

## Is my current setup sufficient?

Using the public cloud efficiently means managing data and business operations differently than in a data centre. A data centre has a certain amount of capacity, versus the scalability of the cloud, and has different security measures. If that works for you – and your plan is to move to the cloud but continue working as though you're using a data centre rather than take advantage of the unique benefits to the cloud – it may not make sense to transition at all.

There is also the question of whether your workloads are in fact able to be hosted in cloud-native infrastructure. If your existing toolkit can't scale or integrate with different data sources, you'll need to first revisit the applications you're using and change them before moving to the cloud.

Even if you are happy with your current setup, it's worth examining if your operations can be made more efficient by a move to the public cloud. For example, a company needs to change the number of workloads running throughout the day based on usage, with someone monitoring each workload and ensuring operations remain stable. But if you don't know exactly how many workloads you need at a given time, it's difficult to plan for this.

The cloud is designed to manage this for you, launching and shutting down workloads and applications as required without having any impact on business operations. You will no longer have to pre-plan your workloads or schedule engineers to monitor them.

But this raises another question:
What does public cloud adoption mean for my IT department?

Moving from a traditional setup to the dynamic cloud infrastructure requires a fundamentally different mindset. This is probably the biggest challenge for IT professionals: in a traditional setup they would be focussed on ensuring operations are managed smoothly and workloads and servers are continuously running – if a server suddenly stops running, it could be damaging for operations.

However, with a dynamic cloud model, IT professionals will be using auto-scaling load-balanced servers or containers and database-as -a-service. This means a server could be legitimately spun down if it's not needed at that specific moment. It's no longer about jumping in when something stops running, but rather trusting in the technology to reduce its load when possible. As a result, monitoring and altering systems and processes need to change – IT teams don't need to know when a server is legitimately spun down, only if it has a negative impact on operations. In other words, your IT department's mindset needs to switch from monitoring at a per server or workload level, to a per service level.

Building cloud native infrastructure is at its heart about building resilient and secure infrastructure. If that's done right, IT professionals will be able to focus on looking at the problem after it's been solved to determine what happened, rather than scrambling to fix it.

IT departments may need to embrace new tools and learn new skills to adopt public cloud and, crucially, they need to be brought on the journey. Some employees may feel threatened by a move from the traditional approach, or feel left behind if they are not involved in transformation or migrations to public cloud services.

What are the business implications? There are other concerns beyond training. Long-term, the cloud is generally more cost-effective (again, assuming it's the right option for your business). Beyond the initial setup costs, there are considerations that will have a financial impact.

The public cloud has a more complex billing structure than a traditional data centre, with meters and charges based on usage. There is a benefit to knowing you're not paying for an application when you're not using it – however, these costs can quickly add up, especially if your applications are not cloud-native. This 'bill shock' is a phenomenon that organisations typically face if cloud-native techniques are not employed.

Budgets and invoicing may also require process changes within organisations when moving to a dynamic cloud infrastructure. Unlike traditional models in which the finance department can approve a fixed cost, public cloud typically works on a monthly variable charging model – one that is usually reliant on company usage and needs rather than a predetermined budget.

## Final thoughts

For most companies, using the public cloud – whether entirely or as part of a hybrid model – is the right move. But it's not enough to decide you want to move away from traditional infrastructure. You must consider how to do it, taking into account your current workload setup, company culture, and the business implications.

If done right, the dynamic and scalable cloud model will set you up with a much more resilient infrastructure that can generally resolve service issues quickly. But that relies on doing the move right from the start and modifying your support and managed service capability.

Put simply: once you decide to move to the cloud, your focus shouldn't be on when to do it. Your key considerations should be around how to do it and what changes you need to implement to make it happen.

# Transforming enterprise system monitoring

Adding new monitored elements is a long and complicated process. Even the most highly trained and experienced teams find it difficult to sustain consistent observation and service levels across the SAP installation with its accompanying dissimilar system elements.

**BY HEIKO MANNHERZ, CHIEF INNOVATION OFFICER, AVANTRA**

WITH SO MUCH TIME and effort expended collecting massive amounts of statistics in hopes of no action being required, it's no wonder that SAP system monitoring is regarded by IT leaders as an unpleasant afterthought.

As painful as it is, high availability monitoring is absolutely critical. For manufacturing companies, the downside of an outage is simply too great. Increasing levels of automation and distributed shopfloor intelligence has made enterprise-wide systems more critical than ever. Unplanned system disruptions can create immense losses in productivity and profitability while impacting customer confidence and trust. With end-to-end visibility through a common platform specifically designed to provide a high level of data analysis and system observability, managed service providers (MSPs) can now meet the growing needs

More advanced systems not only offer all essential functions for monitoring but can also provide self-healing functionality whereby failures can be predicted and prevented. For manufacturing firms, this cutting-edge capability provides exciting options

of the manufacturing industry. This enhanced level of service means manufacturers can have greater confidence in the stability of their critical systems and minimize the risk to their businesses.

## Multiple apps and complex interfaces make IT monitoring inefficient

Advanced Applications GmbH is an innovative SAP partner, headquartered in Germany. It provides all services in the SAP environment to manufacturing companies in multiple sectors, including high-tech, medical technology, automotive, mechanical and plant engineering, plastic and process manufacturing, and pharmaceutical industries.

Having over 250 clients from the industry, Advanced Applications, like many MSPs, offers management options for IT environments on-premise or in the cloud and is required to automate SAP operations and consistently monitor client systems. However, using a monitoring tool based on the SAP Computing Centre Management System, exposure to databases for monitoring was limited.

Using multiple apps for different purposes made issue monitoring overly complex and inefficient, while consuming valuable high-skilled resources on simple tasks.

## Delivering new-age capabilities

Advanced Applications turned to an AIOps-based automation platform for solving success-critical issues. Using AI and real-time monitoring, such systems add new levels of observability, reduce the manual effort required to monitor systems, and can automate and improve the efficiency of maintenance and service-level reporting processes.

More advanced systems not only offer all essential functions for monitoring but can also provide self-healing functionality whereby failures can be predicted and prevented. For manufacturing firms, this cutting-edge capability provides exciting options.

According to Manuel Weggler, Team Leader of SAP Basis at Advanced Applications, “We needed to advance our service from the limited options of using SAP Solution Manager. By employing AIOps technology, we have a system that meets our current needs and reduces the manual work needed to support our clients. It also provides us with some interesting possibilities to provide more proactive monitoring services in the future.” Sophisticated service-level reporting within the monitoring system has significantly reduced the time taken for Advanced Applications to create reports, improving customer experiences, reducing manual errors, and enhancing its image. By automating mundane processes, it has freed highly skilled resources to concentrate on other important areas of the business. Weggler adds, “With the new system in place, we receive fewer queries from customers, and we are all saving time to concentrate on value-add areas of our businesses. It is a win-win scenario.”

## Saving around 30% in costs

Advanced Applications saves around 90 person-hours a month by using its new monitoring software. The company is anticipating its costs being 30 percent lower each year for each system monitored. Using AIOps, Advanced Applications can also quickly analyse the data from SAP systems, enabling swift anomaly identification and response.

## “Working perfectly in the Azure cloud”

Alongside its data centre, Advanced Applications is also increasingly working with hyperscalers for hosting. It already has its first applications running in the Azure cloud using the new monitoring system.

## Summary

A new era of automation has increased the criticality and complexity of IT environments. The cost of losing a production-critical system is enormous and potentially devastating for a business. MSPs need to support this with increased visibility and more efficient and proactive systems monitoring.

Like many MSPs, Advanced Applications was facing challenges using separate tools for SAP and database monitoring and requiring excessive manual overheads to serve its manufacturing clients. Its outdated solution did not provide service-level reporting, and manual compilation of information generated an unacceptable level of errors.

Using an innovative monitoring platform, Advanced Applications has discovered a competitive edge to in the changing times. It now has a standardized solution, effective monitoring through template-based service-level reports, and a system that is both simple and efficiently scalable.

# Achieving the Gold standard?

Pulsant and British Rowing have established a successful technology partnership.

BY PULSANT

ROWING is Great Britain's most continuously successful Olympic sport. The GB Rowing Team has won a gold medal in every Olympic Games since 1984 and has won six Paralympic golds since the sport was introduced to the Paralympic Games programme in 2008. London 2012 was an exceptional Olympics for Team GB and especially for rowing winning nine medals.

This elite sports team has been highly successful in maximising the use of technology for the last 30 years, within the parameters of its funding, in order to keep on making the boats go faster and achieve a truly remarkable legacy of success.

## Going for Gold – Pulsant and British Rowing team up for success

British Rowing operates several mission-critical platforms to provide online services to its 30,000 members and run its data-intensive Olympic and Paralympic team sports analysis.
The organisation previously hosted several of these on-premise. However, maintaining the platforms and ensuring their security and backup levels were satisfactory was neither cost-effective nor sustainable.

In 2017 the organisation turned to Pulsant, looking for a unified platform that could bring those systems together and provide constant availability to the large

amount of data around its membership services and complex training programmes.

Harvey Davies, Head of Technology, British Rowing, said: "Analysing data plays a key role in British Rowing's strategy to continue its long-term success. This applies to information about the Olympic team and studying the pathways that make successful Olympic and Paralympic athletes.

"We are also carrying out detailed analysis into membership, which makes constant access to data extremely important. The integrity of this data is another vital issue and so we also required disaster recovery backup as well as a high level of security. Partnering with Pulsant means we can focus on the delivery of services rather than worrying about platform failure because we have lost power or connectivity on our own premises."

## A resilient and highly available environment

Pulsant designed and now hosts a bespoke private cloud platform with the primary site in its Milton Keynes facility, which British Rowing uses to run the GB Rowing Team's performance data analytics projects. British Rowing also makes use of Pulsant's managed backup service with the aim of keeping its IT solutions secure and scalable.

The infrastructure is resilient and highly available - giving British Rowing the peace of mind that it will always have access to the information it needs to carry out its operations. Pulsant's cloud environment is highly secure and helps British Rowing achieve GDPR compliance.

It also hosts the British Rowing Online Entry platform which supports competitive rowing across England and Wales from beginner to international team selection level. Its two critical functions are competition administration; setting up regattas and head races, and event structure; covering payment and results management, and competition entries from clubs and competitors. Pulsant's infrastructure also supports the new digital content platform British Rowing Plus, which hosts a wide range of content including training programmes and dietary advice, and another newly developed platform, Challenge Hub.

## A win-win partnership gets results for all

Both sides of the partnership have gained tremendously from the project. British Rowing has been able to modernise, scale and raise the standards of its IT systems and improve security and compliance. The new IT infrastructure enabled it to significantly improve online services for clubs and members. Meanwhile, Pulsant continues to be inspired by the association with a globally successful team and has taken away invaluable knowledge. Pulsant will continue to provide a range of hosting solutions for British Rowing until summer 2022, through the rearranged Olympic and Paralympic Games in Tokyo.

Martin Palmer, Regional Sales Director, Pulsant, said: "This fantastic partnership has provided Pulsant with unique insight into how organisations such as British Rowing value their data and how secure and accessible that needs to be.

"The project has also provided a great opportunity for the people at Pulsant to engage with high-performance athletes and witness incredible focus and dedication. This is been a real inspiration to our whole team.

"We chose to go into partnership with British Rowing because we're incredibly privileged and honoured to be associated with such an iconic sporting brand. The success of British Rowing over several decades is legendary in the sport, both in Great Britain and worldwide. For Pulsant it's been a very exciting journey and we look forward to that continuing into the future."

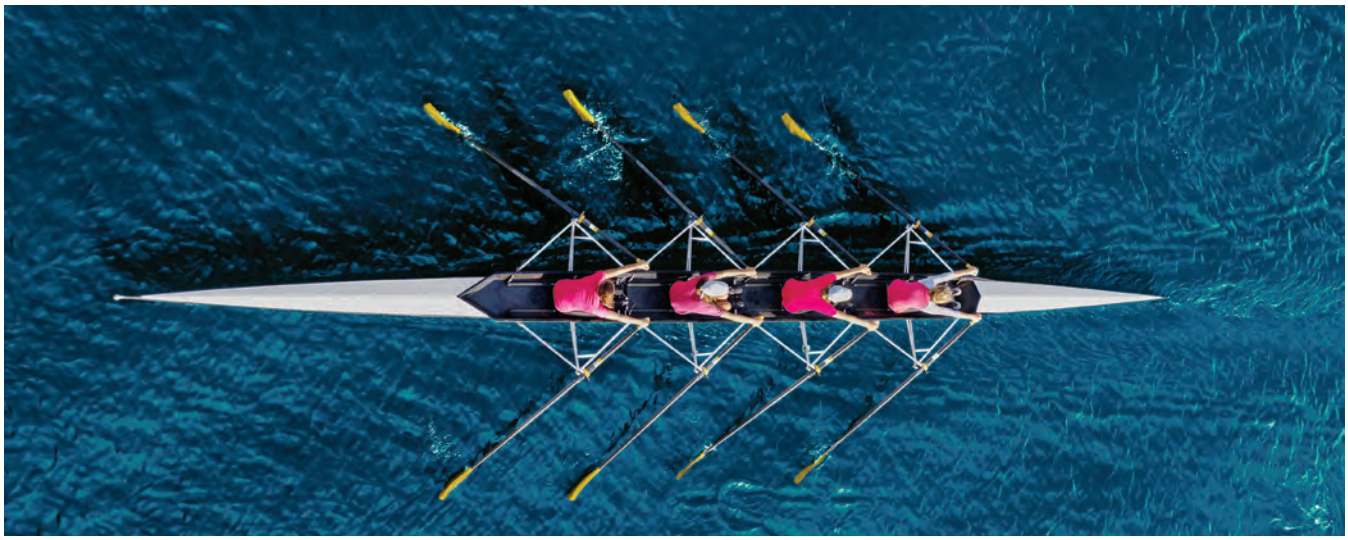

# Canal Extremadura's IT infrastructure revamp with Quantum

Setting its sights beyond traditional radio and TV services, Canal Extremadura is in the midst of a large-scale digital transformation. Headquartered in Mérida, Spain, the company aims to become a public multimedia corporation that produces rich, captivating content for a wide variety of platforms.

BY QUANTUM

TO ACHIEVE its digital transformation, the Spanish television network needed to dramatically improve its IT infrastructure, as it lacked the flexibility to support the evolving business and scale with its fast-growing volumes of multimedia content.

The company's existing archive was a major obstacle as it had reached capacity and expanding it was not financially feasible. As Francisco Reyes, technical chief at Canal Extremadura, explains, "We ran out of room in the tape library and had to migrate some video to a NAS just to free up space."

Canal Extremadura's new archive solution needed to integrate well with its preferred media asset management (MAM) system from Dalet, which is an essential part of the company's media production and post workflow. The new archive also needed to provide a seamless transition from the existing environment, which contained a vast volume of old files in legacy media formats.

## The pursuit for a scalable archive

Choosing a new solution was not a decision Canal Extremadura took lightly. As Reyes explains, "We tend to keep solutions for a very long time—we had been

using the DIVA system for about 12 years—so we needed to be very confident in a new solution before making the selection."

The integration of its Dalet Galaxy MAM system was a key factor in the final decision. Following Dalet's recommendation, Canal Extremadura opted for a Quantum StorNext solution, which includes Xcellis storage servers, an Xcellis metadata array, a QXS disk storage array, and a StorNext AEL6000 tape library. The tape library, which has 400 slots, uses LTO-8 drives—a significant upgrade from the LTO-3 drives that Canal Extremadura was using before.

"The advice and technical information we received from the Dalet and Quantum teams was very helpful. They gave a very clear picture of how the solution would work and how it would be implemented," says Reyes.

Quantum StorNext's networking flexibility has been a great advantage for the company's IT team in supporting a range of client systems. Most notably, the storage environment is configured to offer Fibre Channel connectivity to 10 SAN clients and 10-GbE connections to multiple NAS clients. Meanwhile, the metadata network uses a 1 GbE connection.
To ensure a smooth transition to the new archive solution, the Quantum team provided onsite training for Canal Extremadura. Dalet's implementation team also supported the existing archive's migration to the Quantum environment. "The process took some time because we had a lot of data to migrate, but it was quite smooth," says Reyes.

## Retrieving archived content fast

Thanks to the StorNext environment, some content can remain readily available on disk, meaning Canal Extremadura's journalists, producers, and other team members can now retrieve archived content much faster. "We have more than 100 TB of online storage from Quantum," says Reyes. "So if someone has completed a project six months ago, it will probably still be online."

Even when content has been archived to tape, the IT group can quickly deliver it to users. Upgrading from LTO-3 to LTO-8 has enabled Canal Extremadura to store significantly more data on each tape. Now each retrieval request can be completed without having to load multiple tapes, helping journalists be more efficient in their work and meet their deadlines if they're in a hurry to assemble a new video for a daily news broadcast.

"In the past, users knew they had to wait for content to be retrieved from the archive," says Reyes. "Now it's much faster than before. We have more drives and faster drives with the Quantum archive."

Collaborating with a single vendor has also eliminated some of the compatibility issues experienced with the company's previous multi-vendor environment. Having a single point of contact also simplifies getting ongoing support from Quantum.

## Stepping into a multimedia transformation

With a new long-term scalable archive environment in place, Canal Extremadura is now better prepared for its multimedia future. "If we ever need to expand the archive in the future, we can simply add tapes—it's very straightforward," says Reyes. "With a scalable archive, our company can stay focused on delivering engaging content instead of worrying about where to store it."

# DCA Data Centre Anti-Contamination, Filtration and Cleaning SIG

## An Introduction from DCA CEO Steve Hone

AS THE TRADE ASSOCIATION to the Data Centre sector the DCA understands that it is imperative that key issues affecting the sector have a point of focus. The DCA SIG's (Special Interest Groups) / Working Groups regularly come together over shared interests to discuss issues, resolve problems and make recommendations.

Outcomes result in best practice guides, collaboration between group members, participation in research projects, this includes clarification and guidance for decision and policy makers.

Members find these groups are a great way to ensure their opinions and views are considered in a positive and cooperative environment.

The DCA currently facilitates nine Special Interest or Working Groups. DCA members can join any of the groups (although the Chair has final say) and contribute find out more here: https://dca-global.org/groups

## The DCA Anti-Contamination, Filtration & Cleaning SIG is chaired by Gary Hall, Operations Director at Critical Facilities Solutions UK

THE DEMANDS and growth of digital services has driven radical changes to ICT equipment and this in turn has driven equally radical changes to data centre designs.

This has been caused by wider and greater ranges in temperature and humidity in the data centre together with new technological schemes and upgrades to meet these changes, which in many cases requires a new approach to anti-contamination strategy to ensure the desired reliability and energy efficiency goal of the data centre remains intact.
This group examines the risks posed to data centre facilities of contamination from dust, dirt, airborne particulates that enter data centres. Through a collaborative approach a range of data centre M&E and design experts and several data centre technical cleaning specialists the SIG produces a Best Practise Guide each year. The objective is to provide an independently written guideline for owners and operators to benefit from the collective experience of the industry with the trusted peer review of the DCA.

To request to join this group as a guest or to find out more please contact the DCA - mss@dca-global.org

# SIG Update – Anti Contamination & Filtration Group

## Gary Hall – Critical Facilities Solutions & Chair

ASSOCIATIONS run by individuals, or teams of individuals with solid first-hand experience and an in-depth understanding of the Data Centre industry are of huge value. The DCA Special Interest Group (SIG) for Anti Contamination & Filtration have been raising the profile of contamination control in critical environments since 2013.

The SIG Anti-Contamination and filtration committee examine the risks posed to Data Centre facilities from contamination such as dust, dirt, airborne and gaseous particulates, and other foreign objects that enter mission critical spaces.

The ultimate objective is to provide independent advice and guidelines for owners and operators of mission critical Data Centres to benefit from the collective experience of the industry with the trusted review of the DCA.

Work to Date – Anti Contamination Guide
The first major milestone achieved by the SIG was the creation of the Anti-Contamination Guide that was first released in 2013 with revisions released each year since.

The guide is the product of industry professionals including individuals from Data Centre cleaning and air filtration companies who have contributed years of experience and an in-depth knowledge on the best practices that should be adopted in mission critical spaces.

The guide covers a variety of topics such as how certain contaminates impact operational performance, what level of vetting or experience your specialised cleaning company should hold and how to select and evaluate a company, based on risk, before engagement.

The guide has received exceptional feedback from Data Centre owners and operators and in its current release being referenced in a global standard operating procedure document for a world-wide Data Centre facilities management company which is testament to the credibility of the information it contains. In the annual reviews the SIG evaluates and updates the guide with new standards and best practice, it adds innovative solutions and products and access industry drivers. These are all aimed at helping drive efficiencies through keeping the Data Centre free from dirt, particulate matter, and gaseous contaminates.

In 2019, the SIG introduced a 'risk register' into the Anti-Contamination Guide, this new section of the report listed out in detail all the high-risk contamination elements that could impact operational activities within the Data Centre, and a weighted percentage score was introduced as a representation of what could happen if contamination is not managed correctly. The risk register is under constant review and updated yearly by the group.

The latest release of the Anti-Contamination Guide was released in May 2021 (2021 Release) and can be downloaded from the Data Centre Alliance website.

### Insurance Project

In 2015 the SIG worked together with one of the world's leading insurance companies to create a bespoke insurance package that recognised and rewarded Data Centres that demonstrated best practise with a specific focus on contamination control in the white space.

The SIG visited and assessed four Data Centre's in the UK which varied in size, design, and location. In all cases the broker and underwriters provided cost savings varying from tens to hundreds of thousand pounds. In each case the saving were more than adequate to cover the costs of the audits, accreditations, and all other requirements to quality for the reduction in premiums.

The proposition had market-leading cleaning standards, provided by specialist companies, built into it. The insurance underwriting had been customised to take this into account, as one of the key differentiators, thus raising the profile and ultimately the importance of cleanliness in Data Centre environment.

## Awareness

Specific industry targeted blogs and thought pieces have been carefully written over the past two years by the SIG. One of the highlights was raising awareness of the changes listed in the EU Code of Conduct for Data Centres (Energy Efficiency) Report.

The EU Code of Conduct for Data Centres (Energy Efficiency) report highlighted that air quality is monitored and managed in Data Centre environments to ensure that critical equipment is not damaged by particulates or corrosive elements which might impact both IT equipment and cooling equipment in terms of performance, energy efficiency and reliability.

Section 3.2.12 of the document was listed as an optional practise in early releases of the report until the 2019 release which moved "Monitor and Manage Air Quality" from 'optional' to an expected application to all existing IT, Mechanical, and Electrical equipment within the Data Centre.

This was a significant shift in the report to highlight the importance of Data Centre cleanliness and the SIG felt the need to raise the profile of this change. Additional documents have been produced by the SIG on COVID-19 related measures and the best practices in a very different world we are currently living in, selecting the correct cleaning contractor and the risks of not performing due-diligence, and what the future holds for Data Centre cleaning.

## Look Forward

The SIG have set some ambitious goals for the future as we look to continue to raise the profile of the group. The committees next project is a study and assessment relating to power consumption of a Data Centre to determine 'if' there are savings and 'what' those potential savings are in both kWh and GBP in a contaminated Data Centre versus the same Data Centre once cleaned to an ISO standard by industry experts.

We are trying to evidence that Data Centres would benefit in PUE if they are cleaned and regularly maintained.

The group are currently engaged with owner operated Data Centre managers to progress the idea from paper to reality. This proposal will also see the Anti-Contamination SIG collaborate with the Energy Efficiency SIG highlighting the benefits of an industry body such as the Data Centre Alliance.

Additional plans are to create a bespoke 'Contamination Profile' report which will list all the contaminants that can be found within a Data Centre. The report will look at particulate matter make up, how it is generated and the direct impact they have on the environment.

## Summary

A Data Centre is not a pharmaceutical cleanroom, therefore, contamination from airborne sources is unavoidable. Electrostatic dust, corrosive oxides, volatile organic compounds, solvents, and other contaminants put IT equipment at risk. Even everyday sources of contamination such as pollen, dust, hair, and carpeting fibers can prove to be problematic. The work produced to date by the SIG to highlight these risks is extremely invaluable and concise.

# Datacentre controlled indoor environment with clean air to protect people and operational reliability

## Extract from Camfil Blog

OVER THE LAST YEAR and a half we have all gone through a period of rapidly changing working practices. The increased use of remote home working has boosted the need for datacentre facilities that can be reliable and adaptable to meet these fast changing needs of working.

Currently we are beset globally by Covid19 and the subsequent economics effects. Also global climate change, carbon emissions and the city health effects of exposure to traffic air pollution are increasingly becoming apparent.

Datacentres need to be designed to minimise energy use and operate in a sustainable manner. In one sense they allow people to work more effectively without the need for unnecessary travel and the ensuing carbon emissions but they must not become part of the energy problem.

Currently we are fortunate in the UK to be experiencing a reduction in the Covid19 virus rate of infection. As knowledge of this infection has increased we now know the virus transmission is made principally by breathing in large numbers small infectious airborne particles. These particles are projected into the air by infected people as they breathe, talk, cough and sneeze.

Apart from infectious airborne particles from people inside buildings, ventilation air required for cooling from outside can contain acidic corrosive gases and PM1 fine combustion particles that can be toxic to people damaging to datacentre switching and data storage devices.

Controlling the temperature, humidity and of course cleanliness of the air is therefore critical if consistent working performance and reliability is to be maintained. The same thing could also be said about the people working in these facilities along with their health and wellbeing.

Clean air is of course a relative term and that is where we have to start defining and benchmarking what the limits of acceptability are for exposure to air pollutants. For people the guidance is given in the form of a listed table for particles. PM1, PM2.5 and PM10, PM stands for particulate matter and the number signifies size in microns and below. eg. PM1 is 1 micron and below in size.

# AI and Natural Language Processing

## Faster and more efficient trading

*Connecting* opportunities

www.ipc.com

Protecting datacentre facilities and the people in these building can be best achieved by removing these small particles and acidic gases. This can be best achieved by a effective application of air filtration technology on supply air ventilation and recirculation air filtration in the indoor air space.

**ISA classification of reactive environments (ANSI/ISA 71.04-2013)**

| COPPER REACTIVITY LEVELS (Å/month) | | G1 (MILD) | G2 (MODERATE) | G3 (HARSH) | G4 (SEVERE) |
|---|---|---|---|---|---|
| | | < 300 | < 1,000 | < 2,000 | > 2,000 |
| GROUP | GAS | GAS CONCENTRATION (parts per billion) | | | |
| A | Hydrogen sulfide ($H_2S$) | < 3 | < 10 | < 50 | 50 |
| | Sulfur dioxide ($SO_2$) / Sulfur trioxide ($SO_3$) | < 10 | < 100 | < 300 | 300 |
| | Chlorine ($Cl_2$) | < 1 | < 2 | < 10 | 10 |
| | Nitrogen oxides (NOx) | < 50 | < 125 | < 1,250 | 1,250 |
| B | Hydrogen fluoride (HF) | < 1 | < 2 | < 10 | 10 |
| | Ammonia ($NH_3$) | < 500 | < 10,000 | < 25,000 | 25,000 |
| | Ozone ($O_3$) | < 2 | < 25 | < 100 | 100 |

Over the last few years air monitoring technology has become a lot more affordable and user friendly. This is particularly the case with internet connected devices with the required sensors for the pollutants of concern. This would include PM1, PM2.5, PM10 monitoring for airborne particles, and for gases NO2 nitrogen dioxide, SO2 sulfur dioxide, O3 ozone, H2S hydrogen sulfide and VOC's Volatile Organic Compounds. The better devices available give full data display to screens with adjustable graphic output clearly showing clean air delivery to the indoor space. Some will even give feedback speed control to linked air cleaners and air purifiers.

The acidic gases in indoor spaces have been shown to impair the reliability of datacenter operation so a direct measurement of corrosivity using copper and silver coupons is a tried and tested method by which risk of damage can be assessed. The classifications can be seen in the table below G1 Mild, G2 Moderate, G3 Harsh, G4 Severe. Some monitors can record rates of corrosivity so that the source can sometimes be identified and preventative measures taken.

There are highly efficient air filter solutions for particles and gases that will remove the listed air pollutants and contaminants once they have been identified and the concentrations assessed.

The latest current global test standards for air filter efficiency performance are ISO 16890:2016 for particle air filters and ISO 10121:2013 for molecular gas filters.

To make sure Energy efficient air filter selections are made for particle filters A or A+ Eurovent energy rated bag filters are recommended. This will ensure that the best filtration efficiency is delivered for the lowest energy consumption at a given air volume flow rate.

Using Life Cycle Costing software can allow energy efficient design of filter intake systems which can enable use of free-cooling to give environmental control of datacenter server halls with minimum plant running costs. Multi-stage bag filter systems have been shown to save energy consumption but deliver filtration efficiencies of well over ePM1 90%.

Standalone air cleaners and air purifiers have also been increasingly in datacentre facilities to given localized air cleaning capability. Sometimes this can be in an area that is constantly occupied by people and so needs additional clean air control. Sometimes there is a source or process that needs added control or better protection. Either way the standalone units can provide a flexible and scaleable solution that meets the design requirements of todays modern modular datacentres.

In many way this new and changing role of ventilation air filters and air cleaners to protect people against health risks of Covid19 infection and air pollution fits in well with their traditional role to protect datacentre operational reliability and efficiency. Of course air systems need regular maintenance and inspections to make sure the plant is maintained at optimum efficiency.

Air filtration has proved a trusted and tested technology with effective and accurate international and global technical standards to ensure air cleaning performance is delivered every time.

Links

https://www.camfil.com/en/support-and-services/services/particulate-air-tests/air-image

https://iaq.se/wp-content/uploads/2019/05/Laboratory-services-BROCHURE.pdf

https://www.mynewsdesk.com/camfil/blog_posts/what-are-the-effects-of-corrosion-87313

https://www.camfil.com/en-gb/industries/electronics-and-optics/data-centers

https://www.camfil.com/en-gb/insights/electronics-and-optics/data-centre-air-filtration

https://www.camfil.com/product/documents/dam/25457/Brochures-CamCarb-VG.pdf

https://www.ashrae.org/file%20library/technical%20resources/covid-19/core-recommendations-for-reducing-airborne-infectious-aerosol-exposure.pdf

# The DCA - Anti Contamination, Filtration & Cleaning SIG

## Chaired by Gary Hall of Critical Facilities Solutions

THIS SPECIAL INTEREST GROUP is made up of individuals from various organisations with relevance to this important area of focus for Data Centres.

The group meets on a regular basis to discuss, advise and recommend practical solutions on the control of dust, dirt and contamination. In particular, preventing damage to equipment; loss of data and conservation of energy. The output from this group has included a Best Practice Guide for Anti-Contamination. This was introduced in 2019 with third guide, the 2021 guide now being available.
DCA Data Centre Anti-Contamination Guide 2021 Edition

### Introduction

The demands and growth of digital services has driven radical changes to ICT equipment and this in turn has driven equally radical changes to data centre designs. This has been caused
by wider and greater ranges in temperature and humidity in the data centre together with new technological schemes and upgrades to meet these changes, which in many cases
requires a new approach to anti-contamination strategy to ensure the desired reliability and energy efficiency goal of the data centre remains intact.

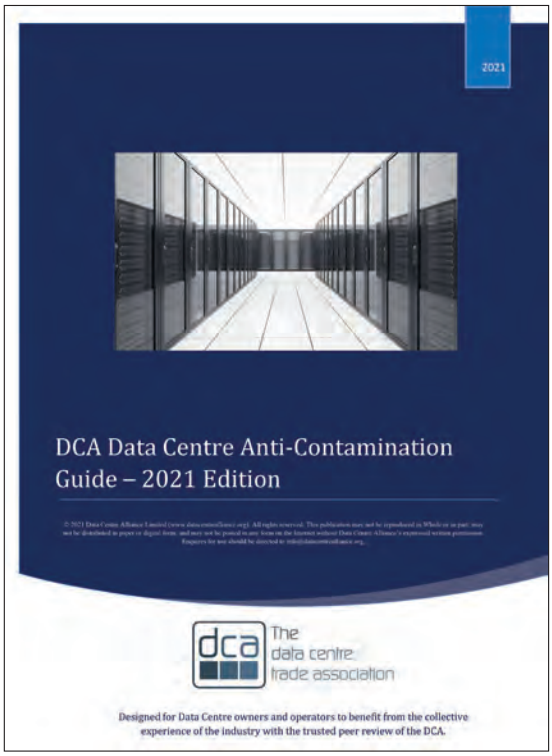


This document examines the risks posed to data centre facilities of contamination from dust, dirt, airborne particulates and other foreign flora and fauna that enter the data centre.

The information provided is the result of a collaborative approach by members of the Data Centre Alliance, an independent industry association. This involved a range of data centre M&E and design experts and a number of data centre technical cleaning specialists.

The objective is to provide an independently written guideline for owners and operators to benefit from the collective experience of the industry with the trusted peer review of the DCA.

### The risks posed to data centres by contamination and dirt

Data Centre owner/operators can benefit from attractive cost savings by implementing modern data centre operating temperatures and humidity guidelines, improved cooling and
ventilation methods and technologies. However, these improvements demand a re-think of the anti-contamination strategy in order to enable these cost savings without compromising the data centres design goal.

### Covid 19 Best Practice Detail 2021

During a global pandemic, Data Centres and operators face huge challenges with risk to staff, engineers and third-party suppliers that are vital to the uptime of the critical environment. Below is a table that details risk adverse recommendations from a cleaning and contamination control perspective.

Download the Full Guide – click here


Angel
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 overal 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, Burnsall Road, Coventry, CV5 6SP. UK
**T:** +44(0)2476 718 970 **E:** info@angelbc.com **W:** www.angelbc.com