DW

DIGITALISATION WORLD

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

JUNE 2020

digitalisationworld.com

```
_mod = modifier_ob.
mirror object to mirror_
mirror_mod.mirror_object =
operation == "MIRROR_X":
mirror_mod.use_x = True
mirror_mod.use_y = False
mirror_mod.use_z = False
_operation == "MIRROR_Y":
mirror_mod.use_x = False
mirror_mod.use_y = True
mirror_mod.use_z = False
_operation == "MIRROR_Z":
mirror_mod.use_x = False
mirror_mod.use_y = False
mirror_mod.use_z = True

selection at the end -add
_ob.select= 1
ier_ob.select=1
ntext.scene.objects.active
("Selected" + str(modifier_
irror_ob.select = 0
bpy.context.selected_obj
data.objects[one.name].se

int("please select exactly

OPERATOR CLASSES ----

types.Operator):
X mirror to the selected
object.mirror_mirror_x"
mirror X"

context):
ext.active_object is not
```

Offloading data centre maintenance to the cloud to support remote working

Forward-looking businesses are offloading data centre maintenance and management to the cloud to keep things running during the COVID-19 pandemic

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

Editor's View

By Phil Alsop

Doom and gloom, or Teams and Zoom?!

READ, watch or listen to the news and, whatever one's political views, it's difficult to become anything other than somewhat depressed as to just how bad things seem to be.

Yes, government responses to the Covid-19 pandemic have varied across the globe, with varying degrees of success, and many countries are beginning to open up after lockdown, but the coronavirus is finding new victims every day, and most, if not all, economies are going to be on life-support for the foreseeable future.

Add to this the political unrest over the 'Black Lives Matter' campaign, and maybe it's time to take early retirement (if you can afford to!) and head off to a peaceful, remote location, should such a place exist…

But wait a minute, history tells us – when we choose to listen – that times of great turmoil and chaos are the perfect conditions out of which are born great new ideas and inventions. And for those for whom the glass will always be half full, the way in which individuals and businesses have reacted to the harsh realities of lockdown, has demonstrated the ingenuity and resilience of the human mind.

And technology solutions have played a key part in this. In our home lives, we've harnessed technology to keep in touch with family and friends, and, along the way, learnt not to take for granted such valuable contact. At work – whether this is in an actual office or from home – Teams and Zoom (other solutions are available) have played a valuable role in keeping the workforce communicating and in focus.

And I've received plenty of stories outlining how many other IT solution have helped all manner of organisations continue to work during lockdown – with managed services and the cloud taking the starring roles.

Furthermore, I'm sure that many IT vendors and suppliers are busy working on more and more new ideas and solutions to help businesses both survive and thrive in the 'new normal' over the coming months and years.

Evolution v revolution is one for the philosophers, but I'm fairly sure that revolutions – however and whenever started – have, in the long term, and in the main, been the catalysts for major, positive changes. Right now, we're on the verge of a major technology revolution. A revolution both in terms of what's being developed but, equally importantly, how it will be used. Properly shaped and harnessed, the benefits could be extraordinary for us all.

DW 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
Sales Manager			
	Jessica Harrison	+44 (0)2476 718970	jessica.harrison@angelbc.com
Director of Logistics			
	Sharon Cowley	+44 (0)1923 690200	sharon.cowley@angelbc.com
Design & Production Manager			
	Mitch Gaynor	+44 (0)1923 690214	mitch.gaynor@angelbc.com
Publisher			
	Jackie Cannon	+44 (0)1923 690215	jackie.cannon@angelbc.com
Circulation & Subscriptions			
		+44 (0)1923 690214	circ@angelbc.com

Directors
Stephen Whitehurst – CEO
Scott Adams – Joint - Managing Director
Sukhi Bhadal – Joint - Managing Director

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

Contents

JUNE 2020

ANALYST

SERVERS + STORAGE

AIOPs

SECURITY

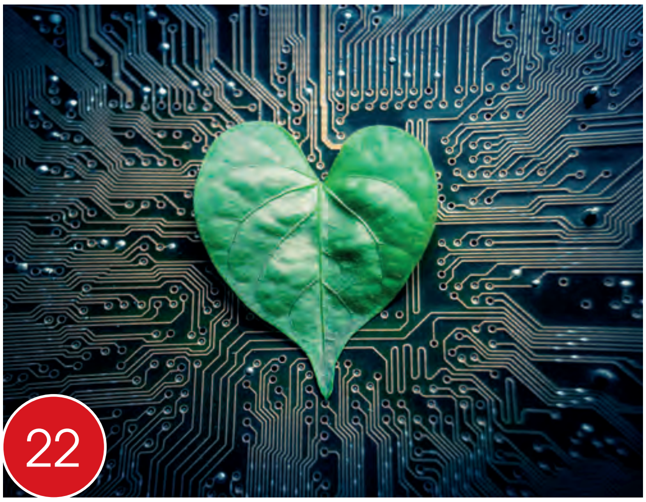

DATA CENTRES

IoT

ANALYTICS

COLLABORATION

EDGE COMPUTING

Ai

Legacy technology and lack of skills stifle digital transformation

VEEAM 2020 Data Protection Trends Report indicates global businesses are embracing Digital Transformation, but struggle with antiquated solutions to protect and manage their data; data protection must move to a higher state of intelligence to support transformational needs and hybrid/multi cloud adoption. As organizations look to transform their business operations and revolutionize customer service, Digital Transformation (DX) is at the top of most CXOs' agendas; in fact, DX spending is expected to approach $7.4 trillion between 2020 and 2023, a CAGR of 17.5%.

However, according to the latest industry data released from Veeam® Software, almost half of global organizations are being hindered in their DX journeys due to unreliable, legacy technologies with 44% citing lack of IT skills or expertise as another barrier to success. Moreover, almost every company admitted to experiencing downtime, with 1 out of every 10 servers having unexpected outages each year — problems that last for hours and cost hundreds of thousands of dollars – and this points to an urgent need to modernize data protection and focus on business continuity to enable DX. The Veeam 2020 Data Protection Trends Report surveyed more than 1,500 global enterprises to understand their approach toward data protection and management today, and how they expect to be prepared for the IT challenges they face, including reacting to demand changes and interruptions in service, as well as more aspirational goals of IT modernization and DX.

"Technology is constantly moving forward, continually changing, and transforming how we do business – especially in these current times as we're all working in new ways. Due to DX, it's important to always look at the ever-changing IT landscape to see where businesses stand on their solutions, challenges and goals," said Danny Allan, CTO and SVP of Product Strategy at Veeam. "It's great to see the global drive to embrace technology to deliver a richer user experience, however the Achilles Heel still seems to be how to protect and manage data across the hybrid cloud. Data protection must move beyond outdated legacy solutions to a higher state of intelligence and be able to anticipate needs and meet evolving demands. Based on our data, unless business leaders recognize that – and act on it – real transformation just won't happen."

The Criticality of Data Protection and Availability

Respondents stated that data delivered through IT has become the heart and soul of most organizations, so it should not be a surprise how important "data protection" has become within IT teams, including not just backing up and restoring data, but also extending business capabilities. However, many organizations (40%) still rely on legacy systems to protect their data without fully appreciating the negative impact this can have on their business. The vast majority (95%) of organizations suffer unexpected outages and on average, an outage lasts 117 minutes (almost two hours).

Putting this into context, organizations consider 51% of their data as 'High Priority' versus 'Normal'. An hour of downtime from a High Priority application is estimated to cost $67,651, while this number is $61,642 for a Normal application. With such a balance between High Priority and Normal in percentages and impact costs, it's clear that "all data matters" and that downtime is intolerable anywhere within today's environments. "Data protection is more important than ever now to help organizations continue to meet their operational IT demands while also aspiring towards DX and IT modernization. Data is now spread across data centers and clouds through file shares, shared storage, and even SaaS-based platforms. Legacy tools designed to back up on-premises file shares and applications cannot succeed in the hybrid/multi-cloud world and are costing companies time and resources while also putting their data at risk," added Allan.

DX and the Cloud

Enterprises know they must continue to make progress with their IT modernization and DX initiatives in order to meet new industry challenges, and according to this report's feedback, the most defining aspects of a modern data protection strategy all hinge upon utilization of various cloud-based capabilities: Organizations' ability to do disaster recovery (DR) via a cloud service (54%), the ability to move workloads from on-premises to cloud follows (50%), and the ability to move workloads from one cloud to another (48%). Half of businesses recognize that cloud has a pivotal part to play in today's data protection strategy; and it will most likely become even more important in the future. For a truly modernized data protection plan, a company needs a comprehensive solution that supports cloud, virtual and physical data management for any application and any data across any cloud.

Allan concluded: "By already starting to modernize their infrastructures in 2020, organizations expect to continue their DX journey and increase their cloud use. Legacy solutions were intended to protect data in physical datacenters in the past, but they're so outdated and complex that they cost more money, time, resources and trouble than realized. Modern protection, such as Veeam's Cloud Data Management solutions, go far beyond backup. Cloud Data Management provides a simple, flexible and reliable solution that saves costs and resources so they can be repurposed for future development. Data protection can no longer be tied to on-premises, physically-dedicated environments and companies must have flexible licensing options to easily move to a hybrid/multi cloud environment."

All change for technology priorities

APPDYNAMICS has released a special edition of its global research study, The Agents of Transformation Report with new findings related to the COVID-19 pandemic. The report reveals the pressures technologists are experiencing as they lead their organizations' responses to the pandemic and how their priorities are changing as the rate of digital transformation accelerates.

The COVID-19 pandemic has required enterprise organizations to shift overnight to an almost completely digital world. Technology departments across the globe are now grappling with surging demand and mounting pressures to accelerate digital transformation strategies. They must deliver high performing digital experiences to customers and all-remote workforces at a time when the survival of the organization is resting on their shoulders.

Technologists are Under Pressure

The latest research from AppDynamics reveals that technologists are experiencing pressure from every angle as they mobilize workforces to operate from home, manage increasing pressure on their networks and applications, and maintain the security of the technology stack, while also taking on new roles and responsibilities.

- 81 percent of technologists state that COVID-19 has created the biggest technology pressure for their organization they have ever experienced.
- 61 percent of technologists feel under more pressure at work than ever before.[2]
- Almost two thirds (64 percent) of technologists are now being asked to perform tasks and activities they have never done before.

Of those surveyed, 66 percent of technologists confirm that the pandemic has exposed weaknesses in their digital strategies, creating an urgent need to accelerate initiatives that were once part of multi-year digital transformation programs.

- 74 percent of technologists report that digital transformation projects which would typically take more than a year to be approved, have been signed off in a matter of weeks.
- 71 percent of technologists point to digital transformation projects that have been implemented within weeks rather than the months or years it would have taken before the pandemic.

Prioritizing the Customer Experience

The research highlights that the majority of organizations (95 percent) changed their technology priorities during the pandemic, and 88 percent of technologists state that the digital customer experience is now the priority.

However, technologists report not having the resources and support they need to make this priority shift, with 80 percent citing that they feel held back from delivering the optimal customer experience because of a lack of visibility and insight into the performance of their technology stack. Technologists list the following areas as the biggest challenges in delivering seamless customer experiences during the pandemic:

- Managing spikes in website traffic (81 percent).
- Lack of unified visibility and insight into performance of the technology stack and its impact on customers (80 percent).
- Managing mean time to resolution (MTTR) with a remote IT department (70 percent).

New Resources and Support are Needed to Rebuild

The report highlights that technologists need specific resources and support from their organizations to meet the challenges ahead. 92 percent state that having visibility and insight into the performance of the technology stack and its impact on customers and the business is the most important factor during this period.13 Other key areas technologists note needing support right now include:

- Clear goals and objectives (90 percent).
- Real time data at the point of need (89 percent).
- Autonomy and accountability (88 percent).
- Freedom to experiment and take risks (87 percent).

Despite the enormous pressure technologists are facing, 87 percent see this period of time as an opportunity for technology professionals to show their value to the business. Already, 80 percent of technologists report that the response of their IT teams to the pandemic has positively changed the perception of IT within their organizations.

Technologists have the opportunity to step up and guide their organizations through the current crisis, but need access to data and insights to make smarter decisions, to be operating within the right internal structures and culture, and with close support from strategic technology partners.

The Urgent Need for Agents of Transformation

In 2018, The Agents of Transformation Report revealed the need for a new breed of technologists, primed to deliver transformation and business impact within their organizations. These elite technologists, known as Agents of Transformation, possess the skills and attributes needed to drive innovation.

At the time, the research concluded that in order for businesses to remain competitive, 45 percent of technologists within an organization needed to be operating as Agents of Transformation within the next ten years. In the midst of the COVID-19 pandemic, technologists report that the target must be reached not within ten years, but now.

With 83 percent of respondents stating that Agents of Transformation are critical in order for businesses to recover quickly from the COVID-19 pandemic, there is an urgent need for technologists to operate at the highest level of their profession.

"Technologists are stepping up in their organizations' hour of need, and it is now the responsibility of business leaders to do everything possible to provide these women and men with the tools, leadership and support they require to deliver first class digital customer and employee experiences," said Danny Winokur, general manager, AppDynamics. "It will be the skill, vision and leadership of these Agents of Transformation that will determine how businesses are able to navigate this turbulent period and emerge stronger on the other side."

Citrix survey shows IT leaders preparing for new work order

75 PER CENT of executives polled say employees won't return to the office as they knew it in wake of COVID-19, 72 per cent accelerating digital transformation to accommodate long-term remote work.

The coronavirus has challenged IT organisations around the world in ways unimaginable. But new research conducted by Censuswide on behalf of Citrix Systems, shows they are rising to the occasion, accelerating their digital transformation efforts to accommodate more flexible ways of working they say employees will demand even after the pandemic subsides. Over three-quarters of more than 3,700 IT leaders in seven countries surveyed believe a majority of workers will be reluctant to return to the office as it was. And 62 per cent say they are expediting their move to the cloud as a result.

Testing their mettle

"COVID-19 has put already stressed IT teams to the test as mandates designed to slow the spread of the virus have forced them to deliver digital work environments with unprecedented speed," said Meerah Rajavel, chief information officer, Citrix. "But as the results of our latest research reveal, they have responded and are stepping up their efforts to accommodate flexible models that will drive work for the foreseeable future."

Flexing their muscle

Over two-thirds of the IT decision makers polled by Censuswide (69 per cent) agree** that it has been surprisingly easy for the majority of their employees to work from home, and 71 per cent say that the technology they have put in place has enabled them to collaborate just as effectively as they can face-to-face. In light of this, they are revving up their digital engines and implementing solutions to support remote work for the long haul.

- 62 per cent of IT leaders say their departments are considering downsizing physical IT infrastructure and transitioning to a cloud model
- 42 per cent anticipate introducing digital workspace platforms
- 44 per cent are looking to public cloud services to facilitate long-term remote working

A rough road

The road to widespread remote work has not been easy. Almost half (48 per cent) of the IT leaders who participated in the Censuswide survey say their organisations did not have a business continuity plan based on the vast majority of employees working from home, and 61 per cent found it challenging*** to make the switch. In addition, the fast and widespread adoption of remote work has opened a new set of concerns and challenges with which they must deal:

- 70 per cent of IT leaders are worried about information security as a result of employees working-from-home
- 54 per cent say there's been a spike in employees installing unsanctioned software.
- 23 per cent say that unscheduled virtual personal network (VPN) shutdowns have been a key problem for their department over the last few weeks.

Taking a toll

All of this has taken a toll on IT teams, with over three-quarters (77 per cent) reporting high stress levels. But there is a silver lining.

"This crisis has thrust IT teams – often the 'unsung heroes' of a business – into the limelight like never before," Rajavel said. "They have worked to deliver secure, reliable work environments that are keeping employees engaged and productive and business moving in extremely challenging times. And in doing so, they will emerge from the crisis more strategic and valued by their organisations than they were going in." More than three-quarters of the IT leaders polled (77 per cent) share this sentiment and say that IT is currently seen as "business critical to their organisation," while 55 per cent believe that their new job title should be "working from home warrior" or "corporate saviour."

Digital transformation falling onto the back burner due to a widening ‘Hesitancy Gap’

IT TEAMS are wasting more than a quarter of their time just laying the groundwork for projects, reveals research from NTT Ltd.

Global Data Centers, a division of NTT Ltd., has released a report revealing that digital transformation projects are stalling due to a ‘hesitancy gap’, as enterprises struggle to navigate the risks and complexity associated with turning innovation from concept to reality. The report, ‘Mind the Hesitancy Gap: There’s No Time to Waste in High Stakes Digital Transformation’, shows that 26% of IT teams’ time is wasted laying the groundwork for digital transformation projects, costing UK enterprises an average of £ 2.01m per year.

Despite a continued focus on digital transformation – with respondents saying they are investing in artificial intelligence (63%), Internet of Things (58%), and software defined networks (51%) – half of UK enterprises admit their projects are always or regularly delayed, as a result of too many barriers to overcome, or too much existing pressure on IT. At the same time, 65% are heavily reliant on multi-cloud services to underpin their projects, creating added integration challenges. In fact, 35% cite the complexity of connecting the range of cloud services and other technologies together as a major barrier to the progression of their digital transformation projects.

“In a rapidly evolving landscape, enterprises can’t afford to drag their feet on digital transformation, but it’s not surprising that many are feeling hesitant,” explains John Eland, Chief Strategy Officer of Global Data Centers. “The complexity of connecting a mix of cloud services and other technologies together, adds a significant challenge to overcome before transformation projects can turn into a reality. Adding further strain, there’s the risk that even just a Proof of Concept could have a negative impact on live production systems, leading to service failures that result in reputational or revenue damage. This is understandably causing enterprises concern, resulting in many projects falling behind and innovation to stagnate.”

Other key challenges to innovation include:

- Concerns over whether a new digital capability could introduce a security risk or leave the company non-compliant with regulation (42%)
- The need to provide significant upfront investment, before even trialling a new transformational concept (37%)
- Limited access to the skills needed to run successful projects with emerging technology, such as AI and Blockchain (35%)
- Concerns that a transformation project could lead to business disruption (35%)
- The time needed to design and build a production-ready environment to run a transformation project or Proof of Concept (PoC) (34%)

Addressing these concerns, the survey showed that UK enterprises estimate they could shave an average of nine months off digital transformation projects, if they didn’t have to spend time building a partner ecosystem and cloud infrastructure and implementing connectivity. A further 94% of enterprises say their digital transformation projects could be ‘supercharged’ if they could test new concepts in a full-scale, production-ready environment, using a multitude of cloud services, partners and connections – without the hassle of pulling everything together by themselves.

“Many of the challenges associated with a project’s viability could be overcome if enterprises could connect with partners and start-ups to test out their concept before taking the plunge,” continuous John Eland. “Having access to a full-scale, production-ready environment where an innovation project can be trailed using, connections, technologies and service providers, would undoubtedly be a significant boon for businesses. By supercharging digital transformation projects and accelerating time to market, enterprises will be well on their way to closing the hesitancy gap.”

The report is based on a survey of 200 IT decision makers in large UK enterprises with over 1,000 people, across multiple industry sectors, conducted by independent research firm Vanson Bourne on behalf of the Global Data Centers division of NTT Ltd. To download the full report, visit:https://datacenter.hello.global.ntt/hesitancy-gap
With its Technology Experience Labs NTT enables enterprises to overcome the hesitancy gap by helping them to develop efficient cloud strategies, optimise their internal IT landscape, and embrace new business models to remain competitive. Enterprises can gain access to a network of service providers, and to a broad range of tools, technologies, cloud services and connectivity.

This reduces the initial investment in complex new scenarios, enables IT departments to measure the impact on IT service delivery, and helps accelerate time to market deployment – crucial to the modern ‘agile’ enterprise.

Global IT spending
to decline 8% in 2020

Worldwide IT spending is projected to total $3.4 trillion in 2020, a decline of 8% from 2019, according to the latest forecast by Gartner, Inc.

THE CORONAVIRUS PANDEMIC and effects of the global economic recession are causing CIOs to prioritize spending on technology and services that are deemed “mission-critical” over initiatives aimed at growth or transformation.

“CIOs have moved into emergency cost optimization which means that investments will be minimized and prioritized on operations that keep the business running, which will be the top priority for most organizations through 2020,” said John-David Lovelock, distinguished research vice president at Gartner. “Recovery will not follow previous patterns as the forces behind this recession will create both supply side and demand side shocks as the public health, social and commercial restrictions begin to lessen.”

All segments will experience a decline in 2020, with devices and data center systems experiencing the largest drops in spending (see Table 1.) However, as the COVID-19 pandemic continues to spur remote working, sub segments such as public cloud services (which falls into multiple categories) will be a bright spot in the forecast, growing 19% in 2020. Cloud-based telephony and messaging and cloud-based conferencing will also see high levels of spending growing 8.9% and 24.3%, respectively.

“In 2020, some longer-term cloud-based transformational projects may be put on hiatus, but the overall cloud spending levels Gartner was projecting for 2023 and 2024 will now be showing up as early as 2022,” said Mr. Lovelock.
“IT spending recovery will be slow through 2020, with

the hardest hit industries, such as entertainment, air transport and heavy industry, taking over three years to come back to 2019 IT spending levels," said Mr. Lovelock. "Recovery requires a change in mindset for most organizations. There is no bouncing back. There needs to be a reset focused on moving forward."

Trends for HR leaders that will impact the future of work

Organizations can differentiate themselves from competitors during the COVID-19 pandemic by leveraging nine ongoing trends, according to Gartner, Inc. These trends are broken into three categories: accelerating trends, new impacts that were not previously part of the future of work discussion, and pendulum swings – temporary shorter-term reactions. "It is critical for business leaders to understand the large scale shifts that are changing how people work and how business gets done," said Brian Kropp, chief of research for the Gartner HR practice. "Then, they must apply this knowledge to their specific organization so they can alter their strategy accordingly."

Gartner recommends HR leaders evaluate the following trends to determine if and how they apply to their business:

Accelerating Trends

Increase in remote work. Gartner analysis shows that 48% of employees will likely work remotely at least part of the time after the COVID-19 pandemic, compared to 30% pre-pandemic. In fact, 74% of CFOs intend to increase remote work at their organization after the outbreak. To succeed in a world of increased remote work, hiring managers should prioritize digital dexterity and digital collaboration skills. HR must consider how the context of remote work shifts performance management, particularly how goals are set and how employees are evaluated.

Expanded data collection. Organizations have increased their passive tracking of employees as their workforce has become remote. According to an April Gartner survey, 16% of organizations are passively tracking employees via methods like virtual clocking in and out, tracking work computer usage and monitoring employee emails or internal communications/chat. In addition, employers are likely to have significantly more access to the health data of their employees. For example, employers will want to know if any of their employees have the COVID 19 antibodies.

"HR leaders should weigh-in on the ethics of using employee data, but also on how to utilize employee monitoring to understand employee engagement across an increasingly dispersed workforce," said Mr. Kropp.

Employer as social safety net. Employers will expand their involvement in the lives of their employees by increasing mental health support, expanding health care coverage, and providing financial health support during and after the pandemic.

Organizations are also considering the question of maintaining compensation for employees, even for those who are unable to work remotely or have been furloughed/laid off throughout and after the COVID-19 crisis.

Expansion of contingent workers. A recent Gartner survey revealed that 32% of organizations are replacing full time employees with contingent workers as a cost-saving measure. Utilizing more gig workers provides employers with greater workforce management flexibility. However, HR will also need to consider how performance management systems apply to contingent workers as well as questions around whether contingent workers will be eligible for the same benefits as their full-time peers.

New Impacts

Separation of critical skills and critical roles. Leaders are redefining what critical means to include: employees in critical strategic roles, employees with critical skills and employees in critical workflow roles. "Separating critical skills from critical roles shifts the focus to coaching employees to develop skills that

Table 1. Worldwide IT Spending Forecast (Millions of U.S. Dollars)

	2019 Spending	2019 Growth (%)	2020 Spending	2020 Growth (%)
Data Center Systems	211,633	0.7	191,122	-9.7
Enterprise Software	458,133	8.8	426,255	-6.9
Devices	698,086	-2.2	589,879	-15.5
IT Services	1,031,578	3.8	952,461	-7.7
Communications Services	1,357,432	-1.6	1,296,627	-4.5
Overall IT	**3,756,862**	**1.0**	**3,456,344**	**-8.0**

Source: Gartner (May 2020)

potentially open multiple avenues for them, rather than focusing on preparing for a specific next role," said Emily Rose McRae, director in the Gartner HR practice. "Organizations should reevaluate their succession plans and may expand the range of roles considered as part of the development path for a given role's potential future successors."

Humanization (and dehumanization) of workers. Throughout the COVID-19 pandemic, some employees have formed more connected relationships, while others have moved into roles that are increasingly task oriented. Understanding how to engage task workers in the team culture and creating a culture of inclusiveness is now even more important. To deliver on employee experience, HR will need to facilitate partnerships across the organization while working with managers to help employees navigate the different norms and expectations associated with these shifts.

Emergence of new top-tier employers. As the labor market starts to return to normalcy, candidates will want to know how companies treated their workforce during the COVID-19 outbreak. Organizations must balance the decisions made today to address immediate concerns during the pandemic with the long-term impact on their employment brand that will span the next several years.

Pendulum swings

Shift from designing for efficiency to designing for resilience. Prior to the COVID-19 crisis, 55% of organizational redesigns were focused on streamlining roles, supply chains, and workflows to increase efficiency. Unfortunately, this path has created fragile systems, prompting organizations to prioritize resilience as equally important as efficiency.

Providing more varied, adaptive and flexible careers helps employees gain the cross-functional knowledge and training necessary for more flexible organizations. Additionally, organizations should shift from trying to "predict" (targeting a specific set of future skills) to "responding" (structuring such that you can quickly course correct with change).

Increase in organizational complexity. Across the next several months there will be an acceleration of M&A, nationalization of companies, and bigger companies becoming even bigger. This rise in complexity will create challenges for leaders as operating models evolve. HR will need to take the lead on shifting to more agile operating models and helping leaders manage greater complexity.

Providing more varied, adaptive and flexible careers helps employees gain the cross-functional knowledge and training necessary for more flexible organizations. Additionally, organizations should shift from trying to "predict" (targeting a specific set of future skills) to "responding" (structuring such that you can quickly course correct with change)

European DX spending keeps growing

According to the latest release of the Worldwide Digital Transformation Spending Guide published by International Data Corporation (IDC), European ICT investments aimed at digital transformation will increase 12% year on year to reach $305 billion in 2020. This represents a significant contraction from the 18% growth forecast prior to the onset of the COVID-19 pandemic.

IN A LOCKDOWN situation with most global economies in at least a mild recession, companies realize that, with the right technology in place, employees can continue working remotely, and operations can be maintained while they adjust as quickly as possible to the new normal. While digitalization investments across industries will suffer, some will be more pronounced than others, depending on the extent to which the given industry has been impacted by the crisis and on the ability of organizations to innovate.

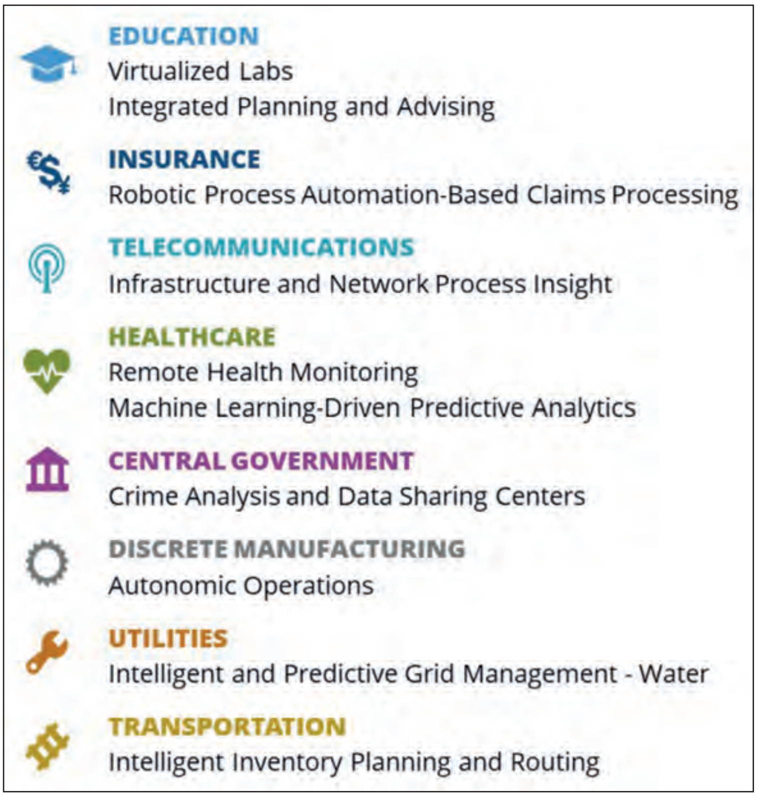

EDUCATION
Virtualized Labs
Integrated Planning and Advising

INSURANCE
Robotic Process Automation-Based Claims Processing

TELECOMMUNICATIONS
Infrastructure and Network Process Insight

HEALTHCARE
Remote Health Monitoring
Machine Learning-Driven Predictive Analytics

CENTRAL GOVERNMENT
Crime Analysis and Data Sharing Centers

DISCRETE MANUFACTURING
Autonomic Operations

UTILITIES
Intelligent and Predictive Grid Management - Water

TRANSPORTATION
Intelligent Inventory Planning and Routing

The biggest drops in DX spending growth in 2020 compared to the pre-pandemic forecast will be seen in transportation, personal and consumer services, retail, and discrete manufacturing. The most resilient sectors are expected to be government, utilities, education, and telecom. In terms of DX spending, healthcare is expected to be the fastest growing industry in 2020 in Europe.

However, even in the most affected industries, some areas of technology investments (use cases) will keep growing, as businesses are looking to resiliency, flexibility, and efficiency to their operations. In discrete manufacturing, efficiency and productivity have never been more important, yet the skills shortage was costing European manufacturing companies billions of euros annually, even before the COVID-19 outbreak.

Although they are considered expensive, autonomous technologies are more appealing to companies than the alternative of downtimes. Technologies like IoT, automation, and robotics also decrease the effect of restrictions related to health concerns (lockdowns) and increase flexibility when reconfiguring production lines. For transportation companies, the intelligent inventory planning and routing use case becomes essential when traditional transportation routes are not available, or destinations and cargo volumes change. In healthcare, use cases like remote patient monitoring are accelerating as more people with chronic diseases can be consulted and treated remotely. In education, uses cases have focused on distance education, which has triggered investments in virtualized labs and integrated planning and advising.

TOP 10 Use Cases for Advancing DX

Investments in the following use cases have accelerated the most compared to the pre-COVID 19 forecast.

IDC COVID-19 Tech Index, 2020 - Worldwide

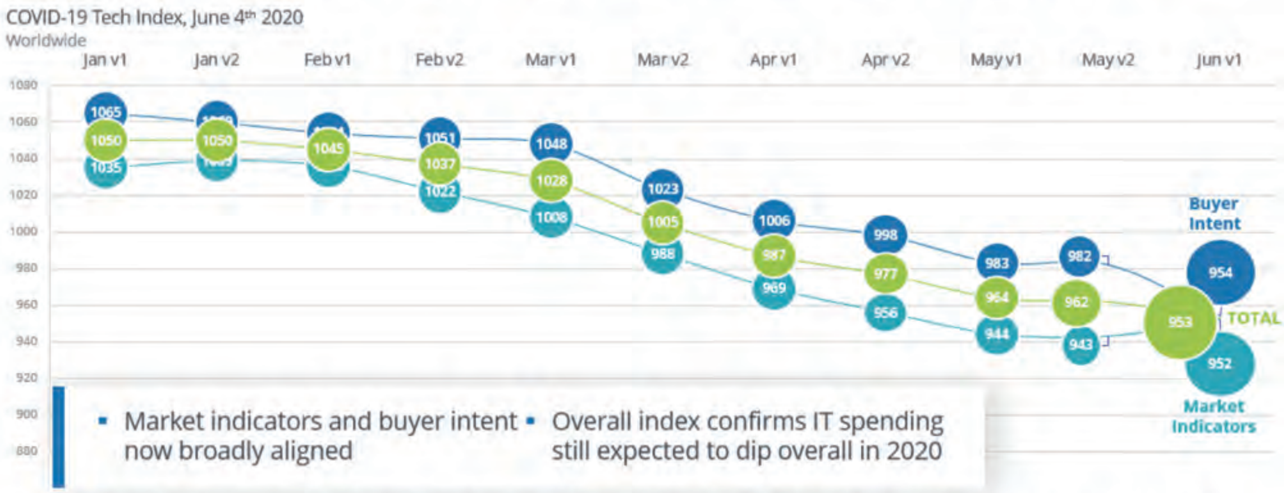

Index based on scale of 1000, where a score above 1000 indicates growth and below 1000 indicates a decline in IT spending; IT buyers were asked to provide guidance on IT spending plans; market indicators reflect a composite of macroeconomic factors; some historical buyer intent data is estimated from analysis of previous surveys and relationship with forecasts published during this timeframe

Source: International Data Corporation

Business confidence in IT spending declines

Business confidence levels declined in the last week of May, according to the latest update to the IDC COVID-19 Tech Index. IT buyers in the US, Western Europe, and some parts of Asia/Pacific indicated that they now expect total IT spending to decline by more than previously anticipated. This is in spite of a general stabilization in other market indicators over the past month, as many countries prepare to tentatively move into a gradual recovery phase.

Confidence levels are still especially weak in the USA, where they have continued to trend down since the crisis began. US firms are a little more confident about the overall economy than two weeks ago, but conversely less confident about their own IT budgets for the year as a whole. Significant spending declines are predicted for traditional technologies including PCs, peripherals, software applications, and project-oriented IT services. Survey results also deteriorated in Europe, especially in France, Italy, and Russia.

"The survey results have diverged with businesses in most countries now expressing less confidence about their own spending than about the broader economy," said Stephen Minton, vice president with IDC's Customer Insights & Analysis group. "This could just reflect the fact that we're still in the middle of the second quarter when the biggest spending cuts are likely to be concentrated and the scale of the short-term impact has been even worse than some firms expected. In fact, survey results are now closer in line with market indicators in terms of the scale of IT spending decline projected for 2020 as a whole."

The COVID-19 Tech Index uses a scale of 1000 to provide a directional indicator of changes in the outlook for IT spending and is updated every two weeks. The index is based partly on a global survey of enterprise IT buyers, and partly on a composite of market indicators which are calibrated with country-level analyst inputs relating to medical infection rates, social distancing, travel restrictions, public life, and government stimulus. A score above 1000 indicates

COVID-19 Tech Index	March	April	May	June
Buyer Intent	1023	1006	983	954
Market Indicators	988	969	944	952
Total Index	1005	987	964	953

Source: IDC COVID-19 Tech Index, 2020

Notes: *Index score above 1000 indicates expected increase in IT spending for 2020 overall; score below 1000 indicates a projected decline.*

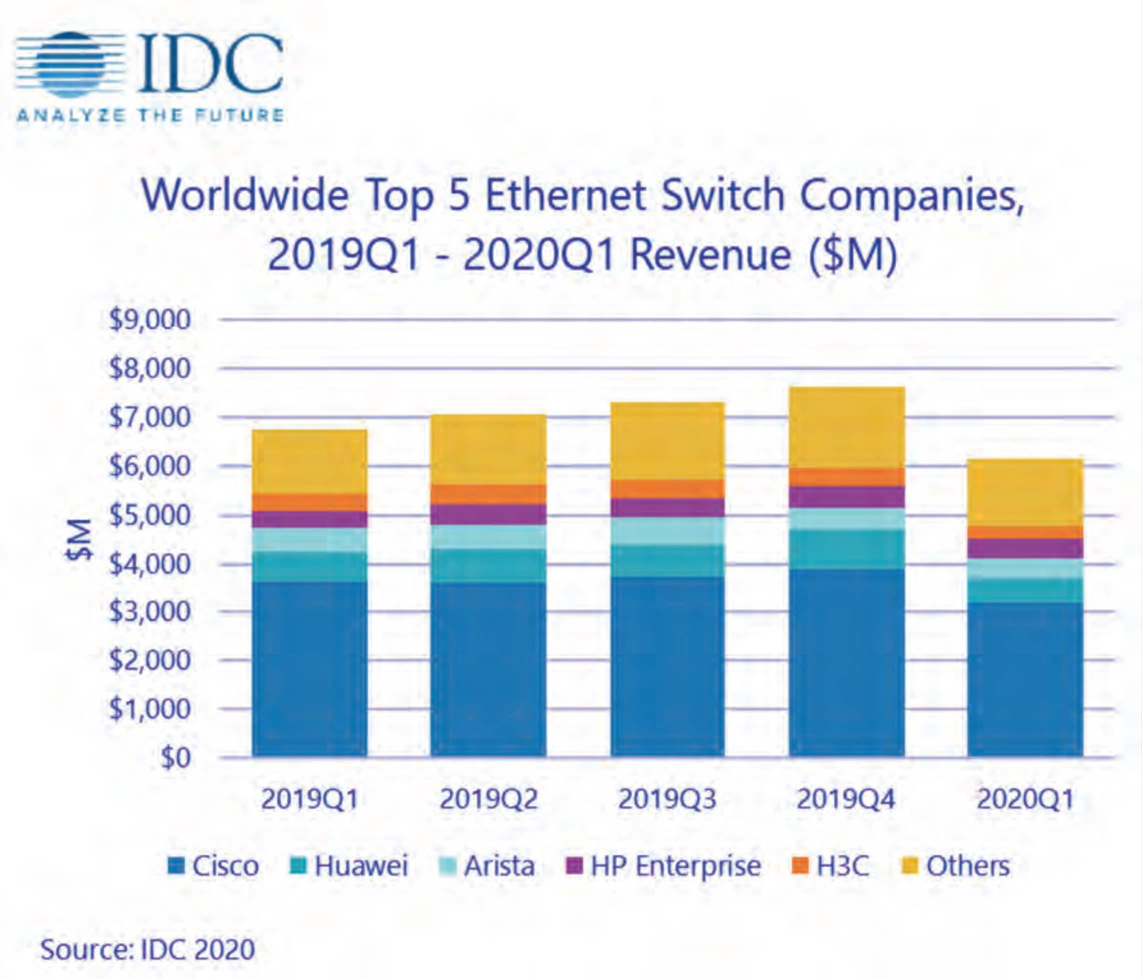

Source: IDC 2020

that IT spending is expected to increase, while a score below 1000 points towards a likely decline.

Business confidence had been improving steadily in Asia/Pacific, but the picture is more complex according to the latest poll. IT spending is still projected to increase in China, where the economy has moved more quickly from a containment to recovery mode, but confidence levels plunged in India and even declined in Korea where moves to ease lockdown measures appeared to trigger some instances of infections increasing again.

“The recovery phase in the second half of the year will be unpredictable and there may be volatility in survey results as businesses react to anxiety around a possible second wave of infections,” said Minton. “The first phase of this crisis was uniformly bad for everyone, but the next chapter will be very localized and dependent on a delicate balance of medical and economic factors.

Not surprisingly, the latest survey results support a sense that IT buyers remain cautious in this type of economic climate and continue to be vigilant in the near term. Moreover, we have now entered a phase where some companies are being forced into bankruptcy or employee reductions, which will have inevitable implications for tech spending in the second half of the year.”

Worldwide ethernet switch and router markets decline

The Worldwide Ethernet switch market recorded $6.16 billion in revenue in the first quarter of 2020 (1Q20), a decrease of 8.9% year over year. Meanwhile, the worldwide total enterprise and service provider (SP) router market revenues fell 16.4% year over year in 1Q20 to $2.99 billion. These market results were published today in the International Data Corporation (IDC) Quarterly Ethernet Switch Tracker and IDC Quarterly Router Tracker. A variety of factors led to the weakening of these markets across the globe.

Despite the Ethernet switch market growing 2.3% for the full year 2019, in the fourth quarter of 2019 the market fell 2.2%, indicating that the market’s slow end to 2019 spilled into 1Q20. The first quarter of 2020 was also impacted by the COVID-19 pandemic that swept across the world throughout the quarter, specifically disrupting supply chains while weakening customer demand. IDC expects the negative impact of COVID-19 on both the Ethernet switch and router markets to continue in the second quarter of 2020.

Ethernet switch market highlights

From a geographic perspective, the 1Q20 Ethernet switch market saw mixed results across the globe. The Middle East and Africa (MEA) region declined 2.9%, with Saudi Arabia's market off 12.7% year over year. Across Europe, growth was uneven. The Central and Eastern Europe (CEE) region grew 3.7% compared to a year earlier, with Russia up 23.2% year over year. The Western Europe market fell 12.9% with Germany losing 10.6% year over year and the United Kingdom off 18.4% from a year earlier.

The Asia/Pacific region (excluding Japan and China) (APeJC) declined 7.0% year over year, with India off 11.3% and Australia declining 16.2% year over year. The People's Republic of China was down 14.6% year over year while Japan was relatively flat with a 0.1% increase compared to the first quarter of 2019. The Latin American market dropped 9.7% year over year, while the market in the United States declined 8.7% annually and Canada fell 7.2% year over year.

"Weakness in the Ethernet switch and routing markets at the end of 2019 continued into the first quarter of 2020, which was exacerbated by the onset of the novel coronavirus and subsequent lockdown of economies around the globe as 1Q20 progressed," said Brad Casemore, research vice president, Datacenter and Multicloud Networks. "Meanwhile, diverging trends intensify in the Ethernet switch market as hyperscale and cloud providers invest in greater datacenter scale and higher bandwidths while enterprises continue to refresh campus networks with lower-speed switch ports."

Growth in the Ethernet switch market continues to be driven by the highest speed switching platforms. For example, port shipments for 100Gb switches rose 52.1% year over year to $5.5 million. 100Gb revenues grew 9.9% year over year in 1Q20 to $1.28 billion, making up 20.8% of the market's revenue. 25Gb switches also saw impressive growth, with revenues increasing 58.9% to $482.9 million and port shipments growing 67.7% year over year. Lower-speed campus switches, a more mature part of the market, saw mixed results in port shipments and revenue, as average selling prices (ASPs) in this segment continue to decline. 10Gb port shipments rose 3.9% year over year, but revenue declined 21.4%. 10Gb switches make up 24.8% of the market's total revenue. 1Gb switches declined 3.8% year over year in port shipments and fell 11.9% in revenue.1Gb now accounts for 39.0% of the total Ethernet switch market's revenue.

Router market highlights

The worldwide enterprise and service provider router market decreased 16.4% on a year-over-year basis in 1Q20, with the major service provider segment, which accounts for 75.1% of revenues, decreasing 16.8% and the enterprise segment of the market declining 15.3%. From a regional perspective, the combined service provider and enterprise router market fell 29.4% in APeJC. Japan's total market grew 8.4% year over year and the People's Republic of China market was off 10.9%. Revenues in Western Europe declined 23.5% year over year, while the combined enterprise and service provider market in the CEE region declined 17.1%. The MEA region was down 5.2%. In the U.S., the enterprise segment was down 12.2%, while service provider revenues fell 19.6%, giving the combined markets a year-over-year drop of 17.5%. The Latin America market declined 15.9% on an annualized basis.

Vendor highlights

Cisco finished 1Q20 with a 12.0% year-over-year decline in overall Ethernet switch revenues and market share of 51.9%. In the hotly contested 25Gb/100Gb segment, Cisco is the market leader with 39.8% of the market's revenue. Cisco's combined service provider and enterprise router revenue declined 28.1% year over year, with enterprise router revenue decreasing 18.7% and SP revenues falling 33.8%. Cisco's combined SP and enterprise router market share stands at 36.3%.

Huawei's Ethernet switch revenue declined 14.0% on an annualized basis, giving the company a market share of 8.4%. The company's combined SP and enterprise router revenue declined 2.1% year over year, resulting in a market share of 28.8%.
Arista Networks saw its Ethernet switch revenues decline 18.7% in 1Q20, bringing its share to 6.7% of the total market. 100Gb revenues account for 73.7% of the company's total revenue, reflecting the company's longstanding presence at hyperscalers and other cloud providers.

OCP Technology Segment Data, 2019 and 2024 (Revenues are in US$ billions)

Market	2019 Revenue	2019 Market Share	202 Revenue	2024 Market Share	2019-202 CAGR
Compute	$13.25	83.1%	$28.07	83.0%	16.2%
Storage	$2.45	16.9%	$5.73	17.0%	18.5%
Total	$15.70	100.0%	$33.80	100.0%	16.6%

Source: IDC Worldwide Open Compute Project Compute and Storage Infrastructure Market Forecast, May 2020.

HPE's Ethernet switch revenue increased 6.7% year over year, giving the company a market share of 6.2%, up from 5.3% market share the same quarter a year earlier.

Juniper's Ethernet switch revenue rose 14.1% year over year in 1Q20, bringing its market share to 3.3%. Juniper saw a 16.0% decline in combined enterprise and SP router sales, bringing its market share in the router market to 10.5%.

"Results in the Ethernet switch and routing markets were fairly consistent across geographies, indicating the widespread impact the Coronavirus has had on both markets across the globe," notes Petr Jirovsky, research director, IDC Networking Trackers. "At the end of 2019, a variety of factors contributed to economic uncertainty, including tensions related to the US-China trade war, Brexit being finalized, and then in the first quarter of 2020, the novel coronavirus. While results from some countries have been mixed in 1Q20, the dynamics across these markets will continue to evolve significantly in the near term. IDC will monitor these trends and their impact across all geographies and segments of these markets."

Worldwide Open Compute project infrastructure market revenue forecast to grow at a 16.6% CAGR through 2024,

A new forecast from International Data Corporation (IDC) shows worldwide revenue from the Open Compute Project (OCP) infrastructure market will reach $33.8 billion in 2024. While year-over-year growth will slow slightly in 2020 due to capital preservation strategies during the Covid-19 situation, the market for OCP compute and storage infrastructure is forecast to see a compound annual growth rate (CAGR) of 16.6% over the 2020-2024 forecast period.

The forecast assumes a rapid recovery for this market in 2021-22, fueled by a robust economic recovery worldwide. However, a prolonged crisis and economic uncertainty could delay the market's recovery well past 2021, although investments in and by cloud service providers may dominate infrastructure investments when they occur during this period.

The Open Compute Project (OCP) was established in 2011 as an open community focused on designing hardware technology to efficiently support the growing demands on compute infrastructure at midsize to large datacenter operators (hyperscalers). Open Compute standards are now supported by market leaders such as Facebook, Microsoft, LinkedIn, Alibaba, Baidu, Tencent, and Rackspace. The OCP encourages infrastructure suppliers, hyperscalers, cloud service providers, systems integrators, and components vendors to collaborate on new innovations, specifications, and initiatives across several key categories.

"By opening and sharing the innovations and designs within the community, IDC believes that OCP will be one of the most important indicators of datacenter infrastructure innovation and development, especially among hyperscalers and cloud service providers," said Sebastian Lagana, research manager, Infrastructure Systems, Platforms and Technologies. "IDC projects massive growth in the amount of data generated, transmitted, and stored worldwide.

Much of this data will flow in and out of the cloud and get stored in hyperscale cloud data centers, thereby driving demand for infrastructure," said Kuba Stolarski, research director, Infrastructure Systems, Platforms and Technologies at IDC.

OCP Technology by Segment

The compute segment will remain the primary driver of overall OCP infrastructure revenue for the coming five years, accounting for roughly 83% of the total market. Despite being a much larger portion of the market, compute will achieve a CAGR comparable to storage through 2024. The compute and storage segments are defined below:

- **Compute:** Spend on computing platforms (i.e., servers including accelerators and interconnects) is estimated to grow at a five-year CAGR of 16.2% and reach $28.07 billion. This segment includes externally attached accelerator trays also known as JBOGs (GPUs) and JBOFs (FPGAs).
- **Storage:** Spend on storage (i.e., server-based platforms and externally attached platforms and systems) is estimated to grow at a five-year CAGR of 18.5% and reach $5.73 billion. Externally attached platforms are also known as JBOFs (Flash) and JBODs (HDDs) and do not contain a controller. Externally attached systems are built using storage controllers.

Buyer type highlights

OCP Board Member purchases make up the bulk of the OCP infrastructure market and are poised to grow at a 14.8% CAGR through 2024, when they will account for just under 75% of the total market. Conversely, non-member spending is projected to increase at a five-year CAGR of 23.2% and will expand its share of the OCP infrastructure market by just over 600 basis points during that period.

In terms of end user type, hyperscalers account for the largest portion of the market at just over 78% in 2019 and are projected to expand spending at a 14.2% CAGR through 2024, although this will result in erosion of total share. Conversely non-hyperscaler purchases will expand 23.8% over the same period, increasing this group's market share by approximately 650 basis points from 2019 to 2024.

DCS INNOVISION2020
DATA CENTRE INSIGHTS + PERSPECTIVES

A very special issue of DCS Magazine dedicated to the data centre industry’s visionary leaders and technology innovators

To herald the launch of the all-new Data Centre Solutions digital publication, we’re producing a very special first issue, entitled InnoVision – providing an overview of the state of the data centre industry right now. 100+ Vendors from across the supply chain are invited to provide their viewpoint on the future and innovation.

How will the data centre industry evolve over the coming months and years, what will be the major drivers and opportunities?

We want your viewpoint!

https://form.datacentre.solutions/innovision

IN ASSOCIATION WITH

DCS DATACENTRE SOLUTIONS

DW DIGITALISATION WORLD

The future of on-premises infrastructure is cloud-defined

DW talks with Nebulon, a storage start-up coming out of stealth mode, with the focus very much on the company's disruptive technology: Cloud-Defined Storage.

DW: *Can you tell us a little bit about Nebulon and why you got started?*

N: As with any startup, we recognized a need in the industry.

When our CEO, Siamak Nazari, was at 3PAR and later at HPE, he regularly travelled to see enterprise customers. During this time, he would often get asked by the CIOs why customers have to buy premium-priced external arrays when they have literally thousands of servers sitting in their data centers, each with a couple dozen slots for disk drives.

Siamak's answer was always the same. In order for server-based storage to have the same enterprise capabilities as external arrays, they would need to have some piece of SW on each server and customers would hate that. That software would need to be different for different OS/hypervisors, would have to be installed and maintained on each server, and their firmware updates and reboots would take storage offline. That was not the answer these enterprises wanted to hear.

But Siamak knew that something needed to be done. After a number of conversations with other IT leaders, and some timely industry developments, our Siamak was able to arrive at a more holistic solution to the CIO challenge shaped in large part by the following themes that came up over and over again:

- Simple API-centric, cloud-based management
- An alternative to the 3-tier architecture
- A modern approach for critical workloads on-prem

So that is why Nebulon was created and what we decided to do. Provide a simple solution for CIO's to leverage their on-prem, server-based storage to handle mission-critical workloads with the flexibility and ease of cloud-managed storage.

DW: *There are storage alternatives on the market. Why can't this problem be solved with existing architectures?*

N: The long and short of it? Arrays are expensive & Hyperconverged Infrastructure and Software-defined Storage solutions are restrictive.

Moving data to the cloud seems like the obvious alternative, right? But the reality is that for many CIOs, their most business-critical data assets often cannot or should not be moved to the cloud for service level, cost or compliance reasons.

These organizations prefer to have external storage arrays and 3-tier architectures move to single-tier server-based approaches similar to hyperscale data centers. These have proven to be cost-effective and easier to use for a specific set of workloads. A

ll agree that array-based capacity remains the highest cost capacity in the industry, with the same SSD sold in a server at a fraction of the cost. But a move to a single tier approach requires a solution without the server software footprint and associated restrictions typical of the SDS/HCI approach.

DW: *You have a big announcement coming up. More generally, where does this announcement fit in the company's overall expansion plans?*

N: Nebulon is coming out of stealth with the announcement of our flagship product: Cloud-Defined Storage!

DW: *Cloud-Defined Storage (CDS) seems to be the name or category of the new Nebulon platform – can you outline this approach to storage?*

N: Cloud-Defined Storage is on-premises, application server-based enterprise-class storage which consumes no server CPU and memory resources and is defined and managed through the cloud.

DW: *Which brings us on to what makes Nebulon different from other storage networking vendors?*

N: There is a lot of innovation in the industry and storage specifically has changed a lot in the last few decades, but not as fast as the connect side of the data center. Ten years ago a company called Meraki Networks, later acquired by Cisco, introduced a line of wireless network devices with the control plane entirely in the cloud to address the complexities of managing network equipment. As we like to say, storage is always late to the party but likes to make an entrance!

With these technologies we have re-factored the array controller into a GPU-like PCIe controller which runs a full enterprise data services software stack for the customer application server, and stripped the management plane from on prem devices and moved it into the cloud.

Existing storage solutions are constantly evolving existing approaches linearly, but Cloud-Defined Storage take a non-linear approach to solve the customer problem.

DW: *CDS offers an API-first approach? What does that really mean?*

N: Well, while many data storage vendors, be it software-defined or traditional, claim to provide some level of this, the reality is that there isn't a single option available that currently offers a true API first approach.

Storage administrators need to automate repeated tasks, streamline complex sequences of operations, or to accelerate time-critical operations in order to keep up with change requests of application owners. So naturally, storage vendors cater to these needs by providing a scripting option for their products but delivered mostly as an afterthought.

> Storage administrators need to automate repeated tasks, streamline complex sequences of operations, or to accelerate time-critical operations in order to keep up with change requests of application owners

Most vendors provide a command line interface (CLI) for automation. However, the CLI is designed for human interaction and using it for automation may jeopardize your data as console outputs are unreliable. Others provide an API that covers only a subset of the array's functionality and is scoped to a single array. What if you have 10, 20 or 100 arrays?

DW: *How many vendors allow you to do firmware upgrades using their API?*

N: In a true API first approach, all devices can be managed from a single API in the cloud and control of every single aspect of your data storage system is done through native APIs that work reliably at scale so every data storage related request from application owners can be fully automated.

DW: *You claim that Cloud-Defined Storage was designed to benefit the IT Administrator and the Application Owner. Can you explain how?*

N: That is exactly right. Cloud-Defined Storage was built for the needs of the IT Admin and the Application Owner.

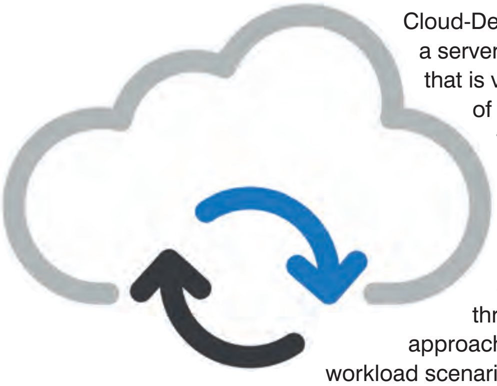

Cloud-Defined Storage gives IT admins a server-based data storage approach that is virtually zero-touch, a fraction of the cost of arrays and SANs– without compromising SLAs.

For the application owner, we transform volume management to simple application provisioning through a template-based approach. We transform a shared workload scenario to one that is entirely workload-aware with the per-application nPod (Nebulon Pod). We also transform the need to query multiple systems individually to simple API call for full visibility across the infrastructure.

App owners get actionable recommendations from AIOps, they 'own' their entire infrastructure by virtue of owning the server (avoiding separate external storage dependencies), and scaling is as easy as adding another server. We like to call it a data center win-win.

DW: *What advantages does Nebulon have over vendors seeking to adapt legacy solutions for the cloud world?*

N: The reality is that for on-prem storage, the only way to achieve true cloud flexibility is if the control plane sits entirely in the cloud. Think of one vendor you know that does that. We can't name any in storage but it is virtually ubiquitous in netowrking.

What benefit does having a control plane in the cloud bring to an organization? Well, it makes their at-scale operations much easier by giving organizatons the following capabilities:

- Simplified "push-button" provisioning.
- Fleet management of thousands of devices.
- Automated software updates.
- At-scale automation from a single IP address in the cloud.
- AI-based insights for operations staff.

DW: *Nebulon recently announced the opening of a Northern Ireland software engineering centre. What was the thinking behind the location?*

N: Establishing our new EMEA Development Center in Belfast was a no-brainer. This is a play we have run in previous companies with terrific results, so the decision was a simple one. Frankly, the engineering talent, both at other tech companies in the area as well as the local universities, make Belfast a terrific and affordable European presence. To sweeten the deal, we partnered with Invest Northern Ireland (Invest NI), an economic development agency within Northern Ireland's Department for the Economy. Through this partnership, we were able to secure funding to kickstart our WW expansion, and have job openings for Belfast's best and brightest in the areas of cloud, security and storage software development.

DW: *Nebulon's Sales HQ is in London. Tell us about the thinking around that?*

N: From the get-go we knew that we wanted to make an impact and have WW operations vs just being focused in the US.

It didn't hurt that our number one pick for the person to lead our sales efforts WW was our friend and former 3PAR colleague, Tim Pitcher who is our vice president of sales and along with Martin Cooper, Nebulon's new senior director of solution architects.

Tim has held senior executive roles at 3PAR, SolidFire and more recently NetApp. Martin also comes with a stellar pre-sales track record at Solidfire and more recently NetApp, where he led the EMEA SE team. We are lucky to convert these two talented guys into #nebnerds to help us bring cloud-defined storage to hybrid IT organizations in EMEA and beyond.

DW: *I can't help noticing that there's a strong '3PAR presence' in the company. How has your team's experiences of working at their previous storage companies, helped to shape Nebulon?*

N: We're very fortunate to be part of a team who has been through this before and working together for such a long time at 3PAR first and then HPE. But we have a lot of people on the team with amazing talent and a great pulse on the industry. Because of this, Nebulon is like no other company.

What makes our team even stronger is bringing in the cloud expertise. We hired a team of cloud engineers in our Northwest Development Center in Bellevue, Washington. It's no secret that with Google, Snap, Twitter, Amazon, Microsoft, etc., the greater Seattle area is the global hub for cloud computing expertise. So that is where we went to tap the world's best software architects and data specialists to develop our revolutionary cloud-defined storage offerings.

Overall, we're excited to recruit new #nebnerds to the team in the upcoming months!

DW: *Finally – any other thoughts or observations, whether they be specific to Nebulon or, more generally, the storage networking and wider IT landscape?*

N: The future of on-premises enterprise infrastructure is cloud-defined and we look forward to bringing this new era of IT to the market in the near future.

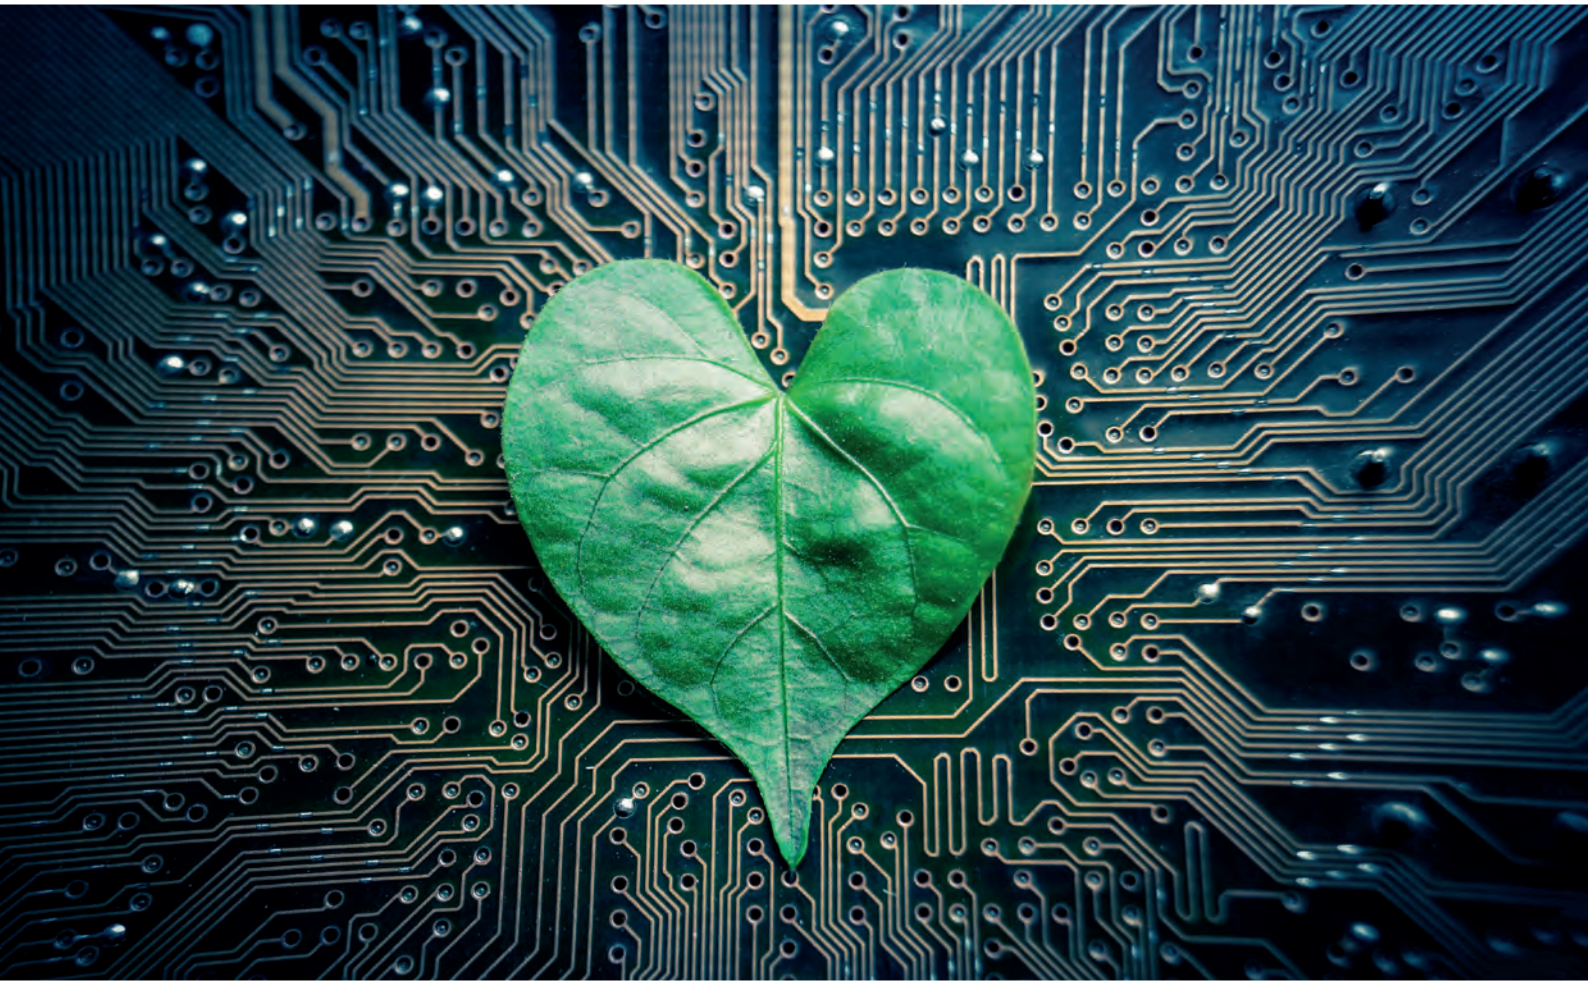

High performance, high efficiency, green computing

DW talks with Michael McNerney, Supermicro Vice President Marketing - covering both the breadth and depth of the company's impressive technology portfolio.

Michael
McNerney
Supermicro
Vice President
Marketing

DW: *Questions for Supermicro for DW and/or DCS Please can you provide us with a bit of background on Supermicro as a company?*

MM: Supermicro was founded and headquartered in USA. The company is a leader in high-performance, high-efficiency server technology and innovation, developing and providing end-to-end green computing solutions to the data centre, cloud computing, enterprise IT, big data, HPC and embedded markets. The Building Block Solutions® approach enables Supermicro to provide a broad range of SKUs to build application-optimised solutions based on specific requirements.

DW: *And what have been the key company milestones to date?*

MM: Supermicro was founded in San Jose, Calif., in 1993 and expanded operations into Taiwan and the Netherlands by 1998. Supermicro now has a global manufacturing footprint across Silicon Valley, the Netherlands and Taiwan, totaling nearly 2 million square feet. Supermicro currently has global operations in over 100 countries and continues to expand.

In 2001, Supermicro created the industry's first dual Intel® Xeon® server which has now developed into a full portfolio of product families including multi-node Twin, Blade, Ultra, GPU and SuperStorage solutions. Recognised as a company of consistent growth, in 2016 Supermicro was named the 'World's Fastest Growing IT Infrastructure' company by Fortune Magazine. Most recently in 2020, Supermicro is listed in CRN's 'Data Center 50: The Hottest Data Center Companies in 2020'.

DW: *And how would you summarise the strengths of Supermicro and its technology portfolio in a competitive market?*

MM: Supermicro's key competitive advantages include internal hardware design, US production and manufacturing and first-to-market hardware technology leadership. A strong partnership with industry leading component vendors enables Supermicro to bring the industry's broadest and flexible portfolio to market ahead of our competitors. As a leader in energy-efficient computing, Supermicro's goal is to promote the adoption and deployment of technologies that can reduce costs as well as the impact on the environment. Resource Saving Architecture continues our tradition of green computing innovation and provides TCO savings for our customers.

DW: *Please can you talk us through the company's product portfolio, starting with servers?*

MM: Supermicro has the industry's broadest portfolio of X11 server systems supporting the new 2nd Generation Intel® Xeon® Scalable processors. Included in the portfolio are the following product lines:

- **Twin Servers:** An innovative, modular computing Twin Architecture which includes multi-node servers in a range of form factors. From the 2U BigTwin (4-node system) for demanding HCI and Storage applications, to the 4U FatTwin for Big Data and HPC applications this range of servers has a variety of options to suit specific requirements.

- **Blade Servers:** High performance servers based on Resource-Saving architecture. Available in 8U, 6U or 4U form factors, Blade servers are designed for a range of use cases with a common goal: extreme density, efficiency and value.

- **Ultra:** Supermicro's highest performance and flexible servers for enterprise applications. This range of servers is available in 1U and 2U form factors and has features such as NVMe, hybrid storage and low-latency optimisation. High flexibility comes from the vast networking and expansion possibilities including Max/IO and Ultra Riser cards.

- **GPU:** The GPU systems at Supermicro are highly optimised for applications requiring maximum performance. With advanced GPU interconnect, these systems offer high efficiency and extreme low latency. Supporting NVIDIA® GPUS, the range includes 1U 4GPU, 2U 6 GPU, 4U 20GPU and 10U 16GPU rackmount servers optimised for AI, Deep Learning, Virtual Desktop, Scientific Research, HPC applications to name a few.

DW: *Moving on to storage?*

MM: Supermicro provides scalable, high-performance NVMe and hybrid storage architectures. The three main product groups are:

- **All-Flash NVMe**
  - Selected Ultra, BigTwin & Blade servers support NVMe
  - 1U Petascale – Capacity optimised system with extreme density, with 32 hot-pluggable drives supporting EDSFF long and short form factors.
  - Storage Bridge Bay servers optimised to deliver high availability to mission-critical storage applications.
- **Top Loading Storage**
  - Easy to deploy 45, 60, and 90-bay storage systems. (NVMe optional)
  - Simply Double: 2U Rackmount servers with double the storage density of a traditional 2U System
  - Cost-Effective 1U Cold storage system, supporting 12 3.5" drives.
- **General Purpose Storage**
  - A wide range of storage solutions designed to meet demands of the toughest storage environments. Flexible form factors include 1U, 2U, 3U or 4U solutions with front and rear loading drives.

DW: *And something which Supermicro calls 'building blocks'?*

MM: The Building Block Solutions® from Supermicro enables the company to provide a broad range of SKUs and deliver application optimised solutions based on specific requirements. The range is made up of Motherboards, Chassis, Rack and System Accessories.

DW: *And then there's an embedded/IoT focus?*

MM: The IoT/Embedded solutions are made up of Fanless/IoT Gateway servers, Compact and Industrial servers, Mini Towers, Rackmount servers and Outdoor Edge Systems.

Dependent on requirements, users can choose from a range of systems which enable multiple edge environments that are not normal environmental conditioned facilities. Building systems for outdoor, customer premises and on manufacturing floors using the building block approach for the system elements not just enables best-in-class, high-end solutions, but

also allows us to include customisation and quick to market systems.

DW: *And networking?*

MM: The Networking Solutions range include Ethernet Switches, Open Networking Switches, Layer 3 Ethernet Switches, Omni Path Switches and Networking adaptors; the range of Blade systems feature integrated switches. With multiple options available, this is an example where the Building Block approach would be taken to ensure a suitable match.

DW: *And, finally, workstations and gaming?*

MM: Supermicro offers a range of SuperWorkstation systems available which have powerful graphics capabilities for rendering, image processing, engineering applications plus more.

Supermicro's gaming solutions are more commonly known as the Supero™ range. This is made up of systems, motherboards and chassis and boasts server grade power leveraging the expertise of Supermicro engineering.

DW: *Solutions-wise, please can you tell us something about Supermicro's approach to the AI and HPC markets?*

MM: Supermicro works closely with partners such as Intel and NVIDIA to produce some of the most powerful and high density AI, Deep Learning and Machine Learning solutions to the market. Offering a custom Deep Learning framework installation, our design enables the end user to directly start deploying Deep Learning projects without any GPU programming. Examples of customised deep learning frameworks for applications include: Tensorflow, Caffe2, MxNet, Chainer, Microsoft Cognitive Toolkit as well as others. Earlier this year, Supermicro announced the industry's broadest portfolio of validated NGC-Ready systems optimised to accelerate AI and deep learning applications. The GPU solution portfolio includes 1U, 4U and 10U solutions which support multiple GPUs, and are optimised to suit the requirements of our customer.

Supermicro is actively innovating in building HPC solutions. From design to implementation, Supermicro optimises every aspect of each solution. Supermicro's advantages include a wide range of building blocks, from motherboard design, to system configuration, to fully integrated rack and liquid cooling systems. Using these tremendous array of versatile building blocks, Supermicro focuses on providing solutions tailored to the customers' need. Supermicro takes pride in building HPC solutions from the ground up.

Supermicro can tailor HPC solutions to meet any variety of workloads: compute intensive, high throughput, or high capacity storage applications used in different industries. Supermicro HPC systems can be bundled with a variety of open source platform and commercial applications, making it a truly turnkey solution.

DW: *And ongoing work in the cloud and virtualisation space?*

MM: The cloud infrastructure required to support today's ubiquitous Public Clouds, SaaS applications, and hybrid IT services provides the backbone of IT infrastructure and is continuing to grow. The growth of data located in the cloud has continued to accelerate, compounded by the adoption of the latest in 5G, IoT, and AI advancements.

Cloud itself has evolved from simple IaaS based on VMs, to PaaS, and now towards a microservices paradigm based on containers often used together with a DevOps software development process. Kubernetes is at the heart of containers and hybrid cloud, together with many other Open Source tools to enable true cloud-native deployments.

Supermicro servers and storage systems have long been used in large-scale clouds worldwide. Supermicro offers the broadest range of cloud-scale products in the industry, designed to provide the highest performance, maximum density and efficient scalability regardless of design patterns. With Supermicro's focus on Resource Savings, customers can achieve lower PUE in their data centres while also reducing acquisition costs and e-waste.

Hybrid cloud solutions from Supermicro and our Open Source partners provide Kubernetes and OpenStack together with the ability to run containers and VMs on-premises and/or off-premises in the Public Cloud for customers.

DW: *And data analytics and enterprise applications are big business?*

MM: There is a tremendous amount of information driven by the ever-changing applications, from structured, unstructured, to semi-structure data. Conventional IT infrastructure is not built to handle the variety, velocity and volume of the data produced by social media networks, mobile applications, machine sensors and scientific researches, etc. For Enterprises, utilising big data analytics is no longer a question of when, it is a question of how. Hadoop, designed for the cost-effective storage and processing of large volumes of data, is born for this purpose. It can linearly scale up to thousands of servers and petabytes of storage.

How to take advantage of Hadoop technology and gain competitive edge is on the mind of almost every corporate CIO. For enterprises, how to deploy the Hadoop infrastructure efficiently means winning or losing in the big battle of market share. Enterprises deploying Hadoop solutions often spend large amount of resource searching for the best architecture and the most capable solution provider. This is where Supermicro comes in to help.

Supermicro Hadoop clusters are a series of optimised big data solutions that provide high performance, high reliability and high scalability. Supermicro Hadoop solutions are fully integrated, fully optimised and completely tested turnkey clusters with flexible support packages available to meet customer specific requirements.

Supermicro Hadoop clusters feature industry proven high density compute and storage servers populated with best of breed components selected through extensive engineering design, validation and testing. Certified configurations take the guess work out of designing and deploying a truly scalable Big Data compute and storage infrastructure that meets the most demanding enterprise IT and data centre environments.

DW: *Presumably, the 5G, edge, IoT focus is expanding right now?*

MM: 5G networks are designed to leverage virtualisation and containers on open hardware platforms – reducing dependence on legacy proprietary hardware. Supermicro brings proven expertise in optimised server hardware and deep virtualisation experience to these growing 5G deployments.

The Radio Access Network (RAN) will be centralised and virtualised into server-based Centralised Units (CUs) and Distributed Units (DUs), slimming down fixed-function hardware into Remote Radio Units (RRUs). Intelligent Edge servers will support real-time applications with computing and AI inferencing. And 4G Evolved Packet Core (EPC) functions will be replaced by 5G Core components running as virtualised network functions (VNFs) in data centres and the cloud.
Supermicro Powers 5G Networks from End-to-End:

- **Data Centres and the Cloud**
  - Disaggregated servers based on Resource-Saving Architecture support data centres in the cloud and the virtualised 5G Core
  - Multi-node Twin architecture and high-density Blade systems for reduced TCO and PUE
  - Deep experience with virtualisation and container solutions
  - Highest-performing GPU servers for data centre machine learning/AI training
- **Virtualised RAN and the Edge**
  - High-performance servers – supporting up to 2nd Gen Intel® Xeon® Scalable Processors – can be configured with FPGA accelerators for virtual RAN (vRAN), including O-RAN
  - Edge AI inferencing with GPU card options improves local decision-making for time-critical applications
  - Systems can be deployed in micro data centres or IP65 ruggedised enclosures
- **Customer Premises and IoT**
  - Compact Edge servers and gateways to manage IoT device deployments from the Cloud
  - Easily reconfigurable networking options to support virtually any installation
  - Wide range of VNFs for SD-WAN, uCPE and other solutions supported with software partners

DW: *And, however mature the market might be, we can't forget data management?*

> Supermicro Hadoop clusters are a series of optimised big data solutions that provide high performance, high reliability and high scalability. Supermicro Hadoop solutions are fully integrated, fully optimised and completely tested turnkey clusters with flexible support packages available to meet customer specific requirements

MM: Yes, data management is a critical component to the successful processing of all forms of data regardless of origin: 5G, IoT, or traditional data centres sources. Supermicro supports a range of software-defined solutions focused on data management, including storage and hyperconverged infrastructure solutions in partnership with a range of software partners.

DW: *Moving on to some recent Supermicro news – please can you tell us something about the recent 27 world record performance benchmarks?*

Supermicro's portfolio of servers optimised for 2nd Gen AMD EPYC Processors has achieved 27 world record performance benchmarks. In addition to the industry's first blade platform, Supermicro's entire portfolio of new H12 A+ Servers fully supports the newly announced high-frequency AMD EPYC 7Fx2 Series processors.

Besides the new H12 SuperBlade and single and dual-socket multi-node Twin A+ Servers, Supermicro is also introducing its next-generation WIO line of A+ Servers as well as a 4U server supporting eight double-width GPUs. With PCI-E 4.0 x16 support, these A+ Servers can deliver 200G connectivity and feature a large memory footprint of up to four terabytes (4TB) per socket running fast DDR4 memory up 3200MHz to deliver record-breaking performance.

DW: *And you've introduced servers for 5G and the telco market?*

MM: Yes, we have moved aggressively to support customer demand and delivered what Supermicro is calling the outdoor edge. We are calling it data centre on a pole, and Supermicro recently announced a new platform a first-to-market with Outdoor Edge Systems – IP65 enclosure-based servers for 5G RAN, AI inferencing, and other intelligent edge-focused applications based on Intel Xeon D processors and 2nd Gen Intel Xeon Scalable processors with broad configuration options. These new systems are ideal for harsh outdoor environments and support the industry's movement toward open-source software and disaggregated hardware.

Supermicro embraces the industry movement to non-proprietary hardware platforms and the growing adoption of standardised system interfaces. Supermicro's membership in the O-RAN Alliance supports its initiative to promote a cloud-native, open 5G RAN architecture for the evolution of 4G to 5G networks. Since joining the O-RAN Alliance, Supermicro has developed reference solutions with several leading telecom operators and software stack providers. This compact pole-mounted solution allows for rapid rollout of adaptable 5G networks with virtually zero real estate.

DW: *And there's been a recent announcement around VMware HCI?*

MM: Supermicro vSAN ReadyNode™ focuses on deploying VMware vSAN, a hyper-converged solution, as quickly as possible. vSAN provides you with the ability to provision and manage compute, network and storage resources from a single pane of management. Working with VMware, Supermicro delivers an alternative to traditional Fibber Channel SAN based virtualisation infrastructure, which is known for its complexity and interoperability challenges. Targeted at a multitude of use cases in tier 1/2 production workloads and Virtualised Desktop Infrastructure (VDI), especially with all flash deployments, Supermicro vSAN ReadyNode™ introduces a new high-performance storage tier optimised for enterprise-class virtual environments that is simple, resilient and efficient that reduces the total cost of ownership. It is a perfect solution for Enterprises ROBO and SMBs to efficiently grow and manage virtualised infrastructure for maximum ROI.

DW: *The new NVIDIA NGC-ready servers sound exciting?!*

MM: Supermicro NGC-Ready Systems are validated for functionality and performance of AI software from NVIDIA NGC. These systems, together with NVIDIA NGC, enable customers to develop and deploy end-to-end AI solutions. The systems come pre-installed with operating system, container, CUDA environment necessary to run NVIDIA NGC software. Supermicro provides NVIDIA NGC software installation as an additional service. Enterprise-grade support for hardware and the operating system is available to help system administrators minimise system downtime, giving users the confidence to support AI workloads. Powered by NVIDIA V100 and T4, the Supermicro NGC-Ready systems provide speedups for both training and inference.

DW: *And what about the COTS systems for the hyperscalers?*

MM: Earlier this year, Supermicro launched its new MegaDC line of servers – the industry's first COTS systems designed exclusively for large scale deployment in hyperscale datacentres. Supermicro's breakthrough MegaDC servers are purpose-built and flexible COTS platforms specifically designed for hyperscale infrastructure deployments. By reducing the component count and optimising the power distribution and backplane designs, MegaDC servers deliver increased cost-effectiveness and reliability.

For better flexibility, these new servers support open standards including OpenBMC for customised control over functionality and versioning, advanced I/O modules (AIOM) that support OCP V3.0 SFF cards, as well as common redundant power supplies (CRPS).

DW: *And clearly you think that the edge, 5G, IoT opportunity is huge?*

MM: Yes, especially as companies and employees rely more and more on virtual conferencing platforms and mobile communications. Latency, and system degradation from an abundance of users can be mitigated by the inherent capabilities of 5G.

DW: *At the other end of the spectrum, what about the hyperscale market, how do you see this developing?*

> The storage landscape is evolving from premium priced proprietary hardware and software solutions to open industry standard hardware and the benefits are significant: reduced vendor lock-in, significantly open innovation with new technologies like all NVMe solutions.

MM: The storage landscape is evolving from premium priced proprietary hardware and software solutions to open industry standard hardware and the benefits are significant: reduced vendor lock-in, significantly open innovation with new technologies like all NVMe solutions. Supermicro's extensive storage platform of choice for leading storage vendors and major hyperscale data centres.

DW: *We seem to be rapidly heading towards a hybrid world, is that where we'll finish, or is there more to come?!*

MM: If you mean solid state hybrid drives and spinning disks for storage we believe this will continue as long as customers demand it, and companies find it to be a profitable market. If you mean hybrid storage (a combination of private and public data centres), we feel this is also going to be an option for companies.

DW: *And there's plenty of talk about ASICs, FPGAs, GPUs and the like – how mainstream will these technologies become over time?*

MM: As evidenced by market trends, and what our customers are asking for, we see all of these components as mainstream, already. Supermicro's Building Block Solutions® are predicated on customisation based on market and customer requirements. Supermicro has relationships with the leading manufacturers of these technologies and can easily create the right solution.

DW: *We've covered a fair amount of ground, if not the kitchen sink(!), are there any other comments you'd like to share?*

MM: We feel encouraged by the combination of leading technologies and a robust market. Supermicro expects to fully participate and offers much more than what has been discussed, including Green Computing, IPMI, Global Servers, Global SKUs and eStore. To find out more information, visit www.supermicro.com and join our virtual events and webinars.

Preventing IT outages and downtime

As businesses continue to embrace digital transformation, availability has become a company's most valuable commodity.

BY DANIELA STRENG, VP & GM EMEA, LOGICMONITOR.

AVAILABILITY refers to the state of when an organisation's IT infrastructure, which is critical to operating a successful business, is functioning properly. However, when an organisation experiences an influx in demand or another catastrophic IT issue, availability subsides and downtime occurs at an alarming rate. One of the biggest challenges organisations face is that availability is difficult to maintain and is indiscriminate, even for the world's largest enterprises.

Companies like British Airways, Facebook and Twitter have all battled through expensive outages in recent years that not only impact their businesses, but also expose society's growing dependence on technology to perform key functions of our daily needs. As technology continues to advance, IT outages will continue to ensue and will affect more than just an organisation's bottom line.

Downtime is still a major issue

Outages occur when an organisation's services or systems are unavailable, while brownouts are when an organisation's services remain available, but are not operating at an optimal level. According to a LogicMonitor survey of IT decision-makers in the UK, US and Canada, and Australia and New Zealand regions, 96 percent of respondents said they experienced at least one outage in the past three years.

Surprisingly, 69 percent of respondents in Australia and New Zealand experienced five or more outages in the last three years, versus an average of 50 percent of respondents in UK, US and Canada respondents who said they experienced five or more outages in the past three years. Only 31 percent of Australia and New Zealand-based IT decision-makers said they experienced four or fewer outages over the

last three years. In comparison, approximately 50 percent of UK, US and Canada respondents said they had experienced four or fewer outages in the same timeframe.

An outage can impact more than just an organisation's finances. The survey found organisations that experienced frequent outages and brownouts incurred higher costs – up to 16-times more than companies who had fewer instances of downtime. Beyond the financial impact, these organisations had to double the size of their teams to troubleshoot problems, and it still took them twice as long on average to resolve them.

The industries most affected

Results from the survey also revealed that the frequency of outages and brownouts is conducive to the industry in which the company operates. Financial and technology organisations experienced outages and brownouts most frequently during a three year period, followed by retail and manufacturing.

According to the survey:

- 41 percent of respondents from financial organisations stated that they experienced 10 or more outages over the past three years.
- 37 percent of respondents from technology organisations said they experienced 10 or more outages over the past three years.
- 34 percent of respondents from retail organisations stated that they experienced 10 or more outages over the past three years.
- 28 percent of respondents from manufacturing organisations stated that they experienced 10 or more outages over the past three years.

These numbers highlight the sweeping nature of outages across the various industry sectors and prove that no company should consider itself immune.

The importance of availability

Availability matters not only to an organisation's customers, but also to the IT decision-makers tasked with maintaining it. In fact, 80 percent of global respondents indicated that performance and availability are important issues, ranking above security and cost-effectiveness. After all, IT availability is essential in the smooth running of IT infrastructure and therefore crucial to maintaining business operations. Availability ensures that airline passengers, for example, aren't stranded due to system outages, food stays at safe temperatures and customers can access their online banking applications.

Despite the importance of availability, IT decision-makers indicated that 51 percent of outages and 53 percent of brownouts are avoidable. This means that organisations could prevent this costly downtime, but do not have the means necessary – whether that involves tools, teams or other resources – to avoid it.

Concerns over the repercussions

With high-profile outages and brownouts hitting the headlines on a regular basis, concerns over the repercussions of experiencing downtime are inevitable. In the UK, 38 percent of respondents said that they will likely experience a major brownout or outage so severe that it will generate media attention, while 35 percent believe someone might lose his or her job as a result of this downtime.

In the US and Canada, 50 percent of respondents said they will likely experience a major brownout or outage so severe that it will generate media attention. Of the same respondents, 52 percent fear someone will lose his or her job. A majority of respondents (63 percent) in Australia and New Zealand feel the same way.
The sector that feared the repercussions of downtime the most was retail, followed by manufacturing. 68 percent of respondents working in retail felt that they would experience a major brownout or outage so severe that it would make national media coverage and that someone could lose his or her job. 67 percent of IT decision-makers in manufacturing felt it would make national coverage, while 69 percent were concerned someone would lose his or her job.

Comprehensive monitoring is key

To combat downtime, it's critical that companies have a comprehensive monitoring platform that allows them to view their IT infrastructure through a single glass panel. This means potential causes of downtime are more easily identified and resolved before they can negatively impact the business. This type of visibility is invaluable, allowing organisations to focus less on problem-solving and more on optimisation and innovation.

Evaluating monitoring solutions can be an arduous but necessary task, and the importance of extensibility cannot be overstated. Companies must ensure that the selected platform integrates well with all of its IT systems and can identify and address gaps in a company's infrastructure that might cause outages. It is also imperative that the selected monitoring solution is not only flexible, but also gives IT teams early visibility into trends that could signify trouble ahead.

Taking it a step further, intelligent monitoring solutions that use AIOps functionality like machine learning and artificial intelligence can detect the warning signs that precede issues and warn organisations accordingly. Ultimately, whether adopting new technologies or moving infrastructure to the cloud, enterprises must make sure that availability is top of mind, and that their monitoring solution is able to keep up. By selecting a scalable platform that provides visibility into their systems and forecasts potential issues, businesses can rise to the next level without sacrificing availability.

This type of visibility will not only prevent downtime and system outages, but also keep organisations from hitting unwanted headlines.

In 2020, DevOps and AIOps go hand-in-hand

Today's DevOps professionals have a lot to monitor and react to in order to keep IT systems running including high volumes of alerts, signals coming from disparate tools and considerable amounts of IT noise. What's more, their workload has only increased since the shift to remote working due to COVID-19.

BY GUY FIGHEL IS GVP & PRODUCT GM, APPLIED INTELLIGENCE AT NEW RELIC.

ON TOP OF ALL THAT, they are expected to continuously improve IT infrastructure performance, problem-solve more accurately and find incident resolutions more readily. So much data and so many alerts can cause response fatigue and make it hard to prioritise issues and know what they really need to act on.

Busy DevOps engineers and leaders are well aware of how AI can help them achieve their goals, possessing a keen 'automation mentality', whereby they identify opportunities to automate away toil by deploying a tool and thus save time down the line. They recognise that the more time they save with AI taking on manual tasks, the more time they have to spend focusing on more complex and higher-value tasks.

Research by New Relic and Vanson Bourne revealed that out of 750 global IT decision makers, 89 percent said they believe AI and machine learning (ML) is important for how organisations run IT operations, and 84 percent also remarked that AI and ML will make their job easier. Plus, findings from Gartner shows use of tools such as AIOps specifically is growing – it predicted their integration in large enterprises will grow from 5 percent in 2018 to 30 percent by 2023.

AIOps tools that detect, diagnose and resolve problems and improve incident response are vital to the success of today's DevOps professionals, but in what ways is AIOps helping them exactly?

Automatic anomaly detection

Some of the latest AIOps tools automatically monitor and detect anomalies via site reliability engineering golden signals such as latency, saturation and traffic. They can then send notifications to IT teams including details about the anomaly. This enables them to quickly and easily assess how to respond, before it potentially causes a problem.

Data-agnostic tools for richer data analytics

Data-agnostic AIOps tools allow DevOps teams to leverage data from numerous sources; standardise it and improve its usefulness with metadata to provide greater context, such as which components are related. This allows users to have a greater understanding of the problem and thus reach the root cause of any issue faster.

Correlation of related incidents to reduce IT noise

DevOps teams are used to noisy environments, but AIOps helps them significantly reduce large volumes of alerts down to manageable amounts and thus avoid alert fatigue. This is possible due to AI establishing relationships between cases of incidents that are alike or related. Some tools also become 'smarter' the more they are used, enabling the user to feedback to the AI, for example, by confirming that it correctly identified alerts were resulting from one issue, training it to spot similar instances in future.

Augmentation of incident management

The use of AI is not to replace those working in DevOps, it's to augment routine activities so that workers can perform better. The two working in tandem together means organisations get the best of both the ability to manage huge datasets accurately from AI, enhanced intuition, and the combined decades of experience of the people that make up the IT team doing their jobs. There are AIOps technologies that include 'decisions builders', which allow users to create their own logic based on event attributes or choose similarity algorithms out-of-the-box to correlate incidents. Tools that are transparent rather than opaque also allow humans to stay fully in the loop with why certain actions were taken so they can stay in control of the process and avoid missing critical signals.

Accurate routing of incident for ownership and actioning

AIOps tools can automatically suggest where to route incidents based on data about the issue and enable DevOps professionals to improve the process by which tasks are distributed among the team. For example, they can mark cases related to a specific application to be sent to a dedicated group, and if they already have too much on, go to another team member with relevant experience and the capacity to own it.

Those in DevOps today may be experiencing the busiest work lives they and their colleagues have ever faced in their careers right now, particularly since the shift to remote working. At the same time though, they have the most advanced technologies at their disposal to deal with the high volumes of alerts and disparate signals successfully.

AIOps tools truly go hand-in-hand with DevOps. This means IT professionals possess quicker and easier ways to identify issues, create diagnoses and find the right resolutions to issues, both after and before they cause problems.

DataOps: making the most of your data faster and more effectively

Adam Mayer, Senior Manager Technical Product Marketing at data strategy company Qlik talks about the challenges facing companies who want to get the best out of their data, and how introducing DataOps can help them keep up with the speed of innovation.

IN TODAY'S ECONOMY, speed is a key driver of competitive advantage. How quickly you get your data ready for analysis can directly impact your business's success. But infrastructure and processes are not always able to handle these demands, as many of our current systems and processes were not built to meet today's demands.

This is of course a technical challenge, but technology alone will not solve it.

Getting to best practice data use requires effective processes as well as tools. So where can companies start when it comes to overhaling their processes and maximising the value of their data?

The trouble with data: where and how are companies struggling?

The explosion in data volume across nearly all parts of work and society is creating expectations of availability, speed and readiness. However, most existing infrastructure was not built for modern data volumes, and as a result many companies continually find they are battling with needlessly slow data applications.

Traditional infrastructure using batch and extended cycles are often not up to the task. Neither are the legacy processes and siloed approaches that some organisations have become accustomed to, where data scientists and analysts are separated from line-of-business teams.

As a result, businesses everywhere are suffering from a data bottleneck, resulting in a huge disadvantage for enterprises as they cannot derive the insights they require in time to act upon them. As speed of insight and analysis becomes a vector of competitive advantage, companies with faster access to their data will be more successful and gather even more relevant data, resulting in a virtuous cycle that could be catastrophic for those who cannot keep up.

How can DataOps help?

While many enterprises are suffering from a data bottleneck, there remains huge untapped potential in their raw data. Improved processes are just as important as technology for accessing this value, which brings us to DataOps. While it has been around for a while, DataOps has seen a spike in uptake recently as data governance and analytics has become a vital source of competitive advantage. As a methodology rather than a product or a platform, DataOps is a set of practices for building data and analytics pipelines to meet business needs quickly.

As these pipelines become more complex and development teams grow in size, organisations need better collaboration and development processes to govern the flow of data from one step of the data life cycle to the next. Quicker processing from data ingestion and transformation to analysis and reporting is becoming essential to operations. DataOps can make this happen faster by streamlining the process. The underlying idea for DataOps is inspired by the DevOps movement, from the software engineering world, which bridges the traditional gaps between development, QA, and operations so that technical teams can deliver high-quality output at a faster pace. DataOps brings together stakeholders across the data landscape, from data architects and engineers to data scientists and IT operations workers who build, maintain and model the data infrastructure.

How streamlining your data will benefit the bottom line

The key commercial benefit DataOps affords businesses in the data-driven era is speed, where consumers expect real-time experiences and where business advantage can be measured in fractions of a second. By automating and simplifying data delivery, DataOps acts as a great leveller for teams skilled at different levels of data literacy.

Users are freed to ask deeper questions sooner, thanks to automated, repeatable requests and AI technologies offering visualisation options for a given data set. It also reduces processing costs as teams can reach insights more quickly, helping them to stay ahead of the competition.

In this way, using DataOps opens up the use of data across the business, which helps encourage a culture where data can drive decisions across even large enterprises. These decisions will also be more collaborative, with the responsibility for data insights no longer concentrated among data engineers or data owners. Making teams more aligned on data-driven decisions also improves the quality of decisions, which can have benefits from customer satisfaction to cost savings.

The value of faster data insights

In every era, speed has given businesses a competitive advantage. But in our current data-driven era, where consumers expect real-time experiences and where business advantage can be measured in fractions of a second, it’s more valuable than ever. Companies who are not able to analyse what their customers are looking for and their own abilities to deliver against this will fall behind.

Bringing together people, processes, and technologies to optimise data pipelines in line with the considerable demands of the modern economy helps get data ready for analysis, monetisation and productisation. Using DataOps to streamline data processes and drive collaboration is helping open the door to better customer intelligence and new business opportunities, offering companies the chance to stand at cutting edge of their industry.

> As these pipelines become more complex and development teams grow in size, organisations need better collaboration and development processes to govern the flow of data from one step of the data life cycle to the next. Quicker processing from data ingestion and transformation to analysis and reporting is becoming essential to operations

Securing an ever-evolving platform: Cybersecurity challenges in the cloud

Cybersecurity is a worry for cloud users. Market research company Vanson Bourne, in conjunction with Nutanix, found that 60 per cent of companies cite security as the biggest factor impacting future cloud strategies. Increasing complexity and evolving technology promise to exacerbate security worries for cloud users, but a strategic approach to security and partnerships can keep data and applications safe.

BY CRAIG TAVARES, HEAD OF CLOUD, APTUM.

FOR MANY CLOUD USERS, a single cloud service provider (CSP) isn't enough. According to a 2019 IDC survey of nearly 300 enterprise IT decision makers, 93 per cent use multi-cloud infrastructures, which can include public CSPs, private hosted (single-tenant) systems, and on-premises systems using cloud technology for flexibility. Many (62 per cent) use multi-cloud for specific capabilities that a single provider can't service. Another driver is political, with different business units specifying their own providers.

More cloud platforms mean more complexity and vulnerability. Multi-cloud users must manage data security across not just one cloud environment, but several. These new worries will exacerbate those existing cybersecurity concerns.

Data breaches and exposures will be among the biggest fears for companies grappling with multi-cloud infrastructures. A simple misconfigured Kubernetes server can give attackers control of your container

infrastructure, for example, while an S3 bucket or ElasticSearch database exposed by an uneducated user can make millions of sensitive records publicly available.

Multi-cloud data management isn't just about security; it also concerns availability. The broader their cloud infrastructure, the more susceptible companies are to DDoS attacks.

Visibility and control are key

Companies hoping to mitigate these risks face a visibility challenge. One of the cloud's promises is also one of its biggest problems: it shields users from the complexities of the underlying environment. It can be difficult enough seeing what's happening in a single cloud infrastructure that abstracts data and applications away from the hardware. Multiple clouds amplify that problem.

A lack of visibility leads to poor control. You can't manage what you can't see. This is where one of the cloud's biggest benefits is also one of its biggest challenges. Cloud infrastructures were built to be flexible and to empower their users. You want a new development server? Sure, spin one up. You need a persistent storage resource? Here's a database for you. But what happens when people misconfigure those resources or deploy them with sensitive information and then don't manage them?

A detailed security policy is a critical part of any approach to controlling any cloud solution. It serves as a baseline for secure operations and compliance with industry security and privacy regulations. There are various policy frameworks to choose from, including ISACA's Controls and Assurance in the Cloud using COBIT 5, and NIST's draft Cloud Computing Security Reference Architecture.

Having a policy isn't enough, though. IT environments are malleable and always evolving. Companies that don't monitor and control operations in the cloud risk one of the biggest cloud security dangers: configuration drift. This is where new resources and configurations move operations away from what the policy demands, creating vulnerabilities and regulatory violations.

Tooling up for cloud security

Companies can solve this problem by using cloud management platforms that give them more control over their operations. Single cloud infrastructure users might get away with that service provider's native cloud management solution, but multi-cloud customers will need a monitoring and control system that gives them a single-pane-of-glass view across all their cloud environments. A security information and event management (SIEM) system should ingest the logs from these monitoring tools, making it available for deeper analysis and long-term trending.

This rich data layer forms the basis for a capable cloud security solution that covers not just multiple clouds, but multiple functions. An integrated security technology stack will handle cloud security needs beyond data platform management and visibility. It will ingest threat intelligence to support deep threat analysis. It will draw on up-to-date product and service vulnerability data to conduct regular vulnerability analyses that companies can use to prioritise patching and change management.

Future attacks

This enhanced readiness will become increasingly important as attackers continue to innovate. We're only just beginning to glimpse the opportunities that artificial intelligence (AI) create for online criminals. One example is the capability for automated attacks. The DARPA Cyber Grand Challenge demonstrated this in 2016, pitting AI-powered attackers against defenders in an all-machine hacking competition. AI also makes it easier to mount social engineering attacks by scanning social media for information about potential targets. Thanks to generative adversarial networks (GANS) that generate fake audio, it's now even possible for AI to impersonate specific voices. Experts already suspect that criminals have used deepfakes in 'whaling' attacks, where malicious actors impersonate senior executives on the telephone and fool employees into transferring money.

AI will become an increasing part of the battle against attackers, too, as defenders fight fire with fire. We are already seeing machine learning tools mining oceans of network traffic and user log-on data to detect anomalies, alerting security analysts to potential problems. An integrated set of multi-cloud security tools helps cloud security teams to correlate seemingly unrelated incidents across multiple cloud environments, helping them to prevent security issues instead of reacting to them.

Over time, this technology will become a standard part of the cybersecurity technology stack. It has to, because attackers will continue to innovate. The smart money already knows this - the defenders in the DARPA Cyber Grand Challenge were also AI-powered. Building solutions to protect yourself in a multi-cloud environment might seem like a daunting task, but not everything needs to happen at once. Most imperative is multi-cloud cybersecurity being built into a cloud strategy from the beginning rather than reactively bolting tools together.

Cybersecurity was never about 100 per cent security or passing and failing grades. It's a probabilistic discipline in which your commitment to the cause and willingness to iterate contribute directly to your overall success. By taking a strategic view and using risk assessment to prioritise your cybersecurity investments, you can begin a cycle of improvement that will evolve your cloud security and serve you for years to come."

Peeling the onion

The concept of a Zero Trust Networking is gaining in popularity. However, many organisations still think of security as protecting the perimeter – complete with layered security technologies resembling the layers of an onion.

BY SCOTT GORDON (CISSP), CMO, PULSE SECURE.

HOWEVER, moving from a legacy position to a more progressive approach to cyber security need not be a big-bang project. Instead, some organisations are embarking on more manageable phased transitions that move over key functions with little disruption and with appropriate investment.

The fundamental idea around ZTN is not new, but the terminology has made a comeback in recent years. A number of technologies including Virtual Private Networks (VPN), Mobile Device Management (MDM), cloud access security and Network Access Control (NAC) practice tenets of the philosophy, but where ZTN has evolved is the notion of tying all of these elements together by aligning access controls across users, devices, applications, and resources, both in the cloud and within data centres. ZTN is a logical response to the reality of cyber-attacks that tend to find and exploit the weakest layer in the onion of technologies protecting today's extended and permeable perimeter. A look at the respected 2019 Data Breach Investigations Report shows an incredible diversity in the root cause for cyber-attacks that includes malware, misconfiguration of systems, insider threat and privilege misuse.

Understanding ZTN

But before getting caught in the "magic bullet" euphoria that is so common within the IT security industry, we should examine the architecture in more detail. ZTN is based on the concept of continuous verification and authorisation. It ensures that only authenticated users with compliant devices, whether corporate, personal or public, can connect to authorised applications over any network, whether on-premises or in the cloud.

This approach may sound less glamorous than adding more bricks to "an impenetrable wall", but in practice, it's more effective for administration, cost and defence. Looking at the constituent parts such as VPN, NAC, MFA in isolation does provide an idea of how – when working in concert – the practice of ZTN can significantly reduce the risk of a cyber-attack turning into a business crippling incident.

Experts in the cyber security industry, including several government agencies, such as UK's National Cyber Security Centre and The National Institute of Standards and Technology (NIST), have aligned around the basic principles of ZTN:

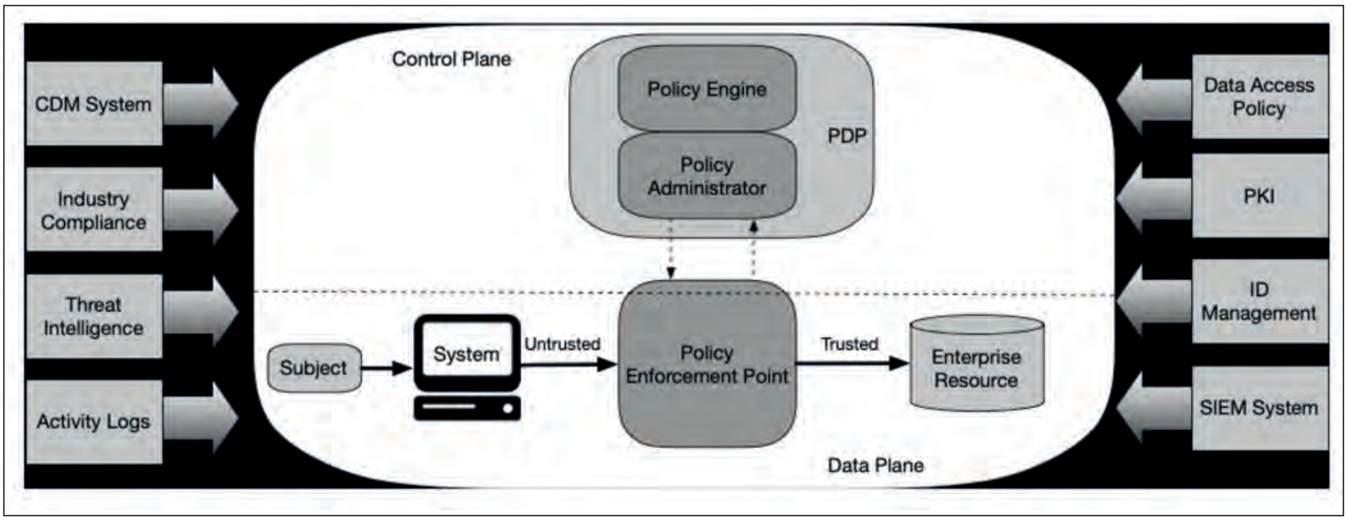

Fig 1 from NIST (SP 800-207(Draft)) document: https://github.com/ukncsc/zero-trust-architecture/

- Know your architecture including users, devices, and services
- Create a single strong user identity
- Create a strong device identity
- Authenticate everywhere
- Know the health of your devices and services
- Focus your monitoring on devices and services
- Set policies according to the value of the service or data
- Control access to your services and data
- Don't trust the network, including the local network
- Choose services designed for zero trust

In order to align with the model of ZTN, there is a set of principles to help organisations align with ZTN tenets without throwing out existing investments.

- Continuous authentication of identity, devices, application, and security posture – before and during any authorised connection
- Centralised authorisation, policy enforcement
- Separated control and data planes
- Granular segmentation based on per application, per-user, and per-device connectivity
- Significantly reduced threat surface by mitigating numerous APTs, malware, DDoS attacks and rendering resources "dark"

To accomplish this, organisations need to review their current secure access solution stack, determine how to orchestrate controls and identify gaps to close depending on access compliance and data protection obligations.

It is critical to centralise policy enforcement so that every user – and each of their devices – is governed by a granular policy based on role, resource and application and other attributes, such as location, to be accessed. It authenticates every user and device security state before the connection is made, ensuring that unauthorised users or devices are only able to see and access authorised resource. Moreover, it also re-verifies user and device security posture during a connection to determine if the security state is no longer acceptable. In such cases, the connection can be terminated, resource access can be reduced, or devices can be quarantined or remediated -

> It is critical to centralise policy enforcement so that every user – and each of their devices – is governed by a granular policy based on role, resource and application and other attributes, such as location, to be accessed. It authenticates every user and device security state before the connection is made, ensuring that unauthorised users or devices are only able to see and access authorised resource

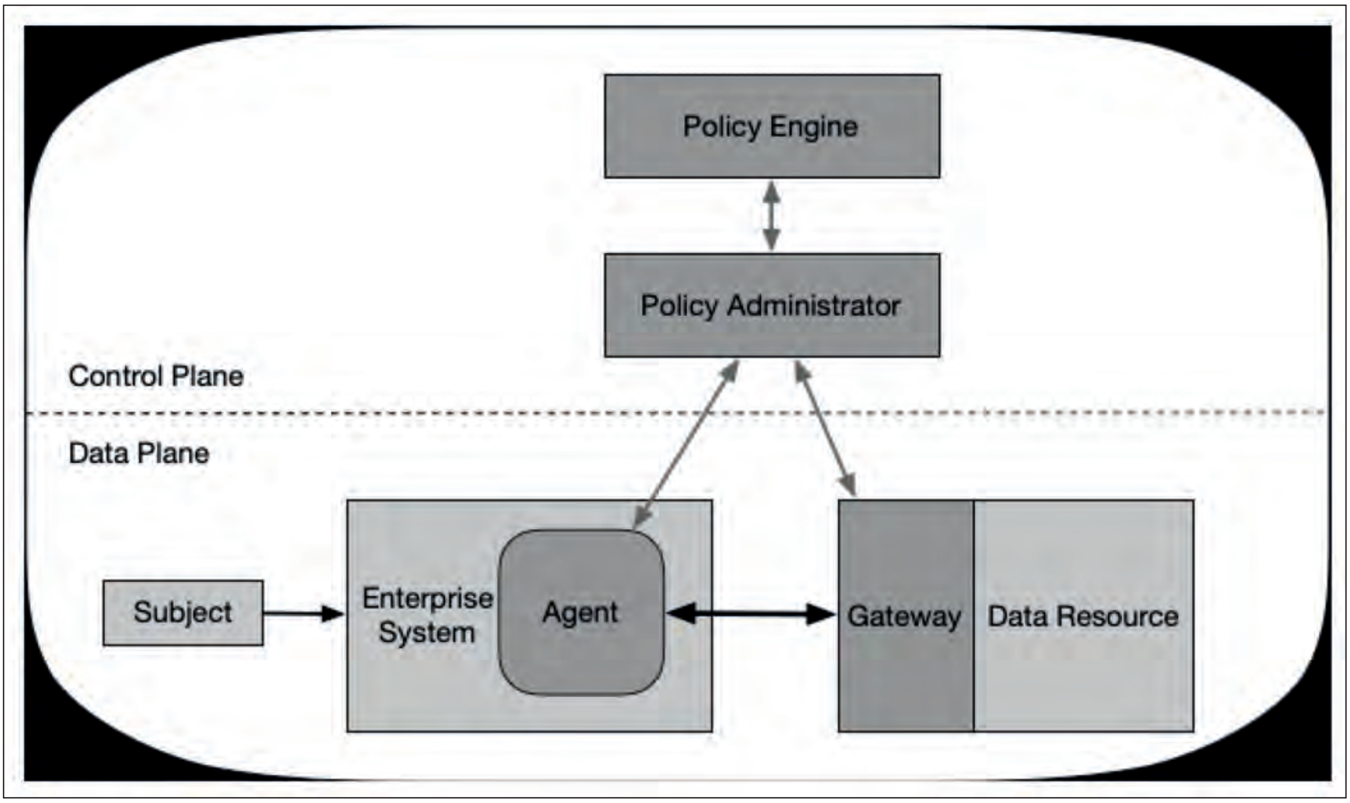

Fig 2 from NIST (SP 800-207(Draft)) document: https://github.com/ukncsc/zero-trust-architecture/

depending on a policy set by the administrator. Finally, resources should be rendered "dark". In other words, no DNS, internal IP address, or visible port information is communicated until proper authorisation takes place. So, unauthorised users can't traverse the network "looking" for resources to infiltrate.

This reduces the attack surface significantly by mitigating or eliminating numerous threats like APTs, man-in-the-middle and malware risks. When moving towards implementing a ZTN model, it's important for organisations to include these controls.

Industry perspective

At an industry technology level, ZTN has gained a lot of coverage and several security vendors have begun to implement the concept. However, only a very small number of vendors have a complete end-to-end zero trust solution that includes gateway, agent, policy administrator/enforcer, and policy engine that spans both physical and virtual environments.

ZTN is a model. As such, it will require organisations to align technologies and orchestrate controls in support of ZTN model tenets. At first glance, this appears to be a major staff endeavour and at a time where organisations are struggling to recruit, train and retain cyber security professionals. However, by prioritising and breaking down the task into key elements that support a new business initiative or a major potential security exposure, ZTN can become more approachable and achievable. Given the increase in cyberattacks and data breaches, the longer-term view is that moving to a ZTN model will lead to less day-to-day security alert firefighting through a systematically improved secure posture and reduced attack surface. A case of short-term pain for longer term gain.

A recent survey found that 72% of organisations plan to assess or implement Zero Trust capabilities in some capacity in 2020 – with larger enterprises being the keenest to take on the effort.

Although moving away from the onion method may seem like a big step for some, the current deluge of security breaches that shows no signs of letting up, with related reputation impact and compliance fines, will prompt more organisations to take decisive ZTN actions – and those that don't want to peel the onion may find that they end up crying anyway.

The Hybrid Cloud:

Avoiding DR and high availability pitfalls

This article examines the hybrid cloud from the perspective of high availability (HA) and disaster recovery (DR), and provides some practical suggestions for avoiding potential pitfalls.

BY DAVID BERMINGHAM, TECHNICAL EVANGELIST AT SIOS TECHNOLOGY.

THE PRIVATE CLOUD remains the best choice for many applications for a variety of reasons, while the public cloud has become a more cost-effective choice for others. This split has resulted – intentionally or not – in the vast majority of organizations now having a hybrid cloud. But there are many different ways to leverage the versatility and agility afforded in a hybrid cloud environment, especially when it comes to the different high availability and disaster recovery protections needed for different applications.

Buyer beware in the Cloud

The carrier-class infrastructure implemented by cloud service providers (CSPs) gives the public cloud a resiliency that is far superior to what could be justified for a single enterprise. Redundancies within every data center, with multiple data centers in every region and multiple regions around the globe give the cloud unprecedented versatility, scalability and reliability. But failures can and do occur, and some of these failures cause downtime at the application level for customers who have not made special provisions to assure high availability.

While all CSPs define “downtime” somewhat differently, all exclude certain causes of downtime at the application level. In effect, the service level agreements (SLAs) only guarantee the equivalent of “dial tone” for the physical server or virtual machine (VM), or specifically, that at least one instance will have connectivity to the external network if two or more instances are deployed across different availability zones.

Here are just a few examples of some common causes of downtime excluded from SLAs:

- The customer’s software, or third-party software or technology, including application software (e.g. SQL Server or SAP.)
- Faulty input or instructions, or any lack of action when required (which covers the mistakes inevitably made by mere mortals.)
- Factors beyond the CSP’s reasonable control (e.g. carrier network outages.)

It is reasonable, of course, for CSPs to exclude these and other causes of downtime that are beyond their control. It would be irresponsible, however, for IT professionals to use these exclusions as excuses for not providing adequate HA and/or DR protections for critical applications. Accommodating Differences Between HA and DR Properly leveraging the cloud’s resilient infrastructure requires understanding some important differences between “failures” and “disasters” because these differences have a direct impact on HA and DR configurations. Failures are short in duration and small

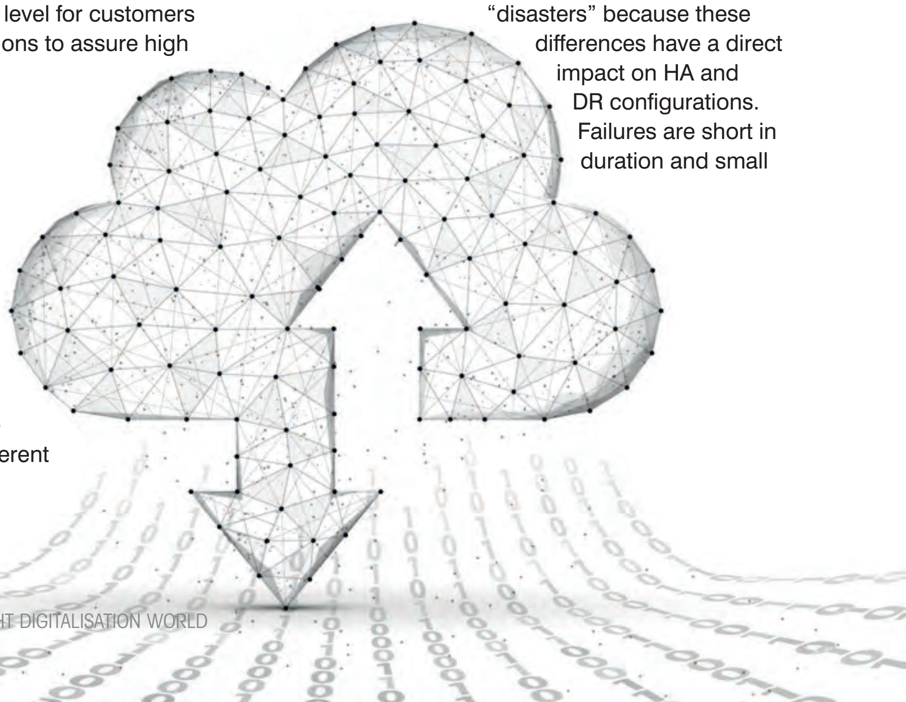

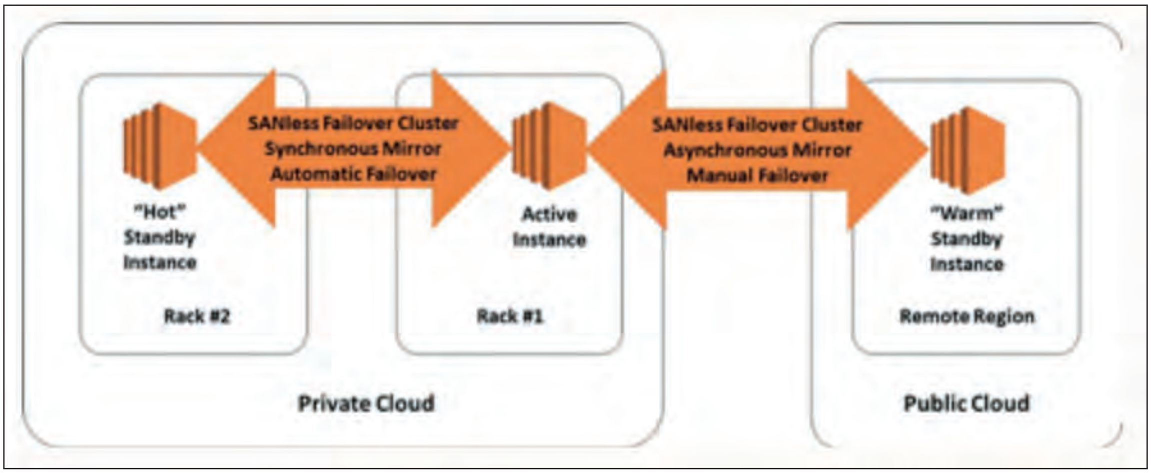

in scale, affecting a single server or rack, or the power or cooling in a single datacenter. Disasters have more enduring and more widespread impacts, potentially affecting multiple data centers in ways that preclude rapid recovery.

The most consequential effect involves the location of the redundant resources (systems, software and data), which can be local – on a Local Area Network – for recovering from a localized failure. By contrast, the redundant resources required to recover from a widespread disaster must span a Wide Area Network. For database applications that require high transactional throughput performance, the ability to replicate the active instance's data synchronously across the LAN enables the standby instance to be "hot" and ready to take over immediately in the event of a failure. Such rapid, automatic recovery should be the goal of all HA provisions.

Data is normally replicated asynchronously in DR configurations to prevent the WAN's latency from adversely impacting on the throughput performance in the active instance. This means that updates being made to the standby instance always get made after those being made to the active instance, making the standby "warm" and resulting in an unavoidable delay when using a manual recovery process.

HA in the Cloud

All three major CSPs accommodate these differences with redundancies both within and across data centers. Of particular interest is the variously named "availability zone" that makes it possible to combine the synchronous replication available on a LAN with the geographical separation afforded by the WAN. The zones exist in separate data centers that are interconnected via a low-latency, high-throughput network to facilitate synchronous data replication. With latencies around one millisecond, the use of multi-zone configurations has become a best practice for HA. IT departments that run applications on Windows Server have long depended on Windows Server Failover Clustering (WSFC) to provide high availability. But WSFC requires a storage area network (SAN) or some other form of shared storage, which is not available in the public cloud. Microsoft addressed this issue in Windows Server 2016 Datacenter Edition and SQL Server 2016 with the introduction of Storage Spaces Direct. But S2D has its own limitations; most notably an inability to span multiple availability zones, making it unsuitable for HA needs.

The lack of shared storage in the cloud has led to the advent of purpose-built failover clustering solutions capable of operating in private, public and hybrid cloud environments. These application-agnostic solutions facilitate real-time data replication and continuous monitoring capable of detecting failures at the application or database level, thereby filling the gap in the "dial tone" nature of the CSPs' SLAs. Versions available for Windows Server normally integrate seamlessly with WSFC, while versions for Linux provide their own SANless failover clustering capability. Both versions normally make it possible to configure different failover/failback policies for different applications.

More information about SANless failover clustering is available in Ensure High Availability for SQL Server on Amazon Web Services. While this article is specific to AWS, the cluster's basic operation is the same in the Google and Azure clouds. It is worth noting that hypervisors also provide their own "high availability" features to facilitate a reasonably quick recovery from failures at the host level. But they do nothing to protect against failures of the VM, its operating system or the application running in it. Just like the cloud itself, these features only assure "dial tone" to a VM.

DR in the Cloud

For DR, all CSPs have ways to span multiple regions to afford protection against widespread disasters that could affect multiple zones. Some of these offerings fall into the category of DIY (Do-It-Yourself) DR guided by templates, cookbooks and other tools. DIY DR might be able to leverage the backups and snapshots routinely made for all applications. But neither backups nor snapshots provide the continuous, real-time data replication needed for HA. For databases,

mirroring or log shipping both provide more up-to-date versions of the database or transaction logs, respectively, but these still lag the data in the active instance owing to the best practice of having the standby DR instance located across the WAN in another region. Microsoft and Amazon now have managed DR as a Service (DRaaS) offerings: Azure Site Recovery and CloudEndure Disaster Recovery, respectively. These services support hybrid cloud configurations and are reasonably priced. But they are unable to replicate HA clusters and normally have bandwidth limitations that may preclude their use for some applications.

HA and DR in a Hybrid Cloud

One common use case for a hybrid cloud is to have the public cloud provide DR protection for applications running in the private cloud. This form of DR protection is ideal for enterprises that have only a single datacenter and it can be used for all applications, whether they have HA protection or not. In the enterprise datacenter, it is possible to have a SAN or other form of shared storage, enabling the use of traditional failover clustering for HA protection. But given the high cost of a SAN, many organizations are now choosing to use a SANless failover clustering solution.

The diagram b

elow shows one possible way to configure a hybrid cloud for HA/DR protection. The use of SANless failover clustering for both HA and DR has the additional benefit of providing a single solution to simplify management. Note the use of separate racks in the enterprise data center to provide additional resiliency, along with the use of a remote region in the public cloud to afford better protection against widespread disasters.

This hybrid HA/DR configuration is ideal for enterprises with only a single datacenter.
This configuration can also be “flipped” with the HA cluster in the cloud and the DR instance in the enterprise datacenter. While it would also be possible — and even preferable — to use the cloud for both HA and DR protection, this hybrid configuration does at least provide some level of comfort to risk-averse executives reluctant to commit 100% to the cloud.
Note how using SANless failover clustering software makes it easy to “lift and shift” HA configurations when migrating from the private to a public cloud.

Confidence in the Cloud

With multiple availability zones and regions spanning the globe, all three major CSPs have infrastructure that is eminently capable of providing carrier-class HA/DR protection for enterprise applications. And with a SANless failover clustering solution, such carrier-class high availability need not mean paying a carrier-like high cost. Because SANless failover clustering software makes effective and efficient use of the cloud’s compute, storage and networking resources, while also being easy to implement and operate, these solutions minimize ongoing costs, resulting in robust HA and DR protections being more affordable than ever before.

Experience zero hassle with Zero-Touch network security automation

Dania Ben Peretz, Product Manager at AlgoSec, discusses the steps needed for organizations to achieve Zero-Touch automation in their network security.

It takes significant manpower and time to get network projects from A to B, especially as organizations migrate to cloud or hybrid cloud environments. With the hundreds or thousands of devices and connectivity options, there is a huge amount of manual effort required to manage networks and the workflows within them, while having absolute confidence that they are secure and free of misconfigurations. One small mistake – even something as simple as a typo when creating a firewall rule – and you could create critical security holes or application outages.

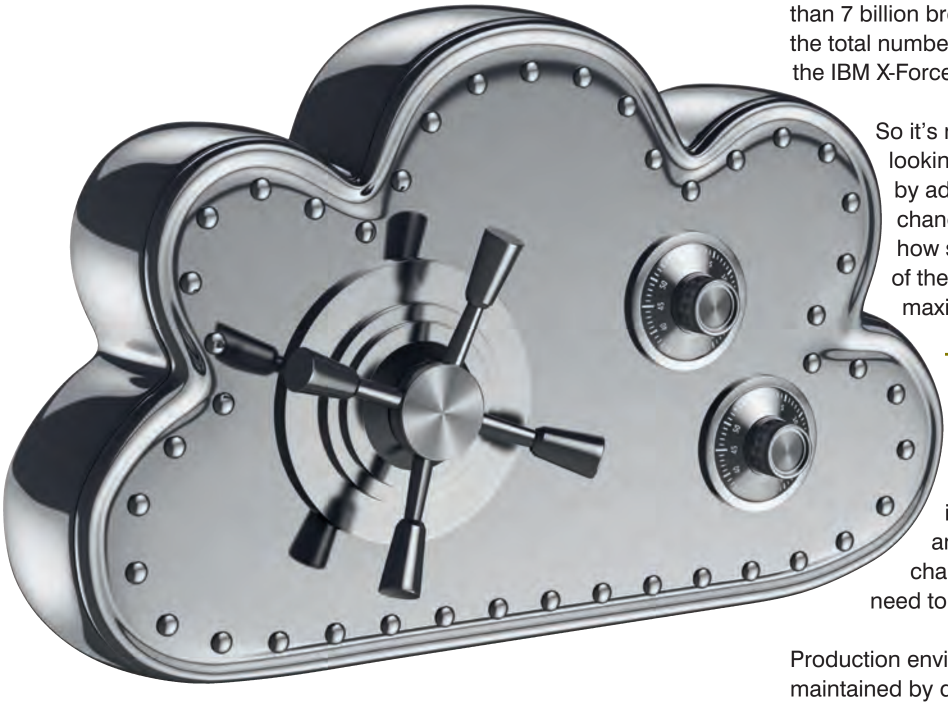

A simple error made during a routine change management process could open up a vulnerability that an attacker can exploit. And given the pressure from the business to make changes quickly – such as spinning up new servers or resources rapidly to serve a business need – those errors are all too easy to make. In 2019, misconfigured systems and cloud servers were responsible for the exposure of more than 7 billion breached data records, or over 85% of the total number of compromised records tracked in the IBM X-Force Threat Intelligence Index.

So it's no surprise that organizations are looking to reduce or eliminate these errors by adopting solutions to automate network change processes and management. But how should they approach the deployment of the solution, to ensure they get the maximum benefits from it?

The barriers to automation

Ideally automation should enable faster delivery, have a positive cost driver, and come with zero failure. Achieving true Zero-Touch automation in the network security domain is not an easy task. In reality, there are many challenges and barriers that organizations need to overcome.

Production environments in all organizations are maintained by different teams, for example DevOps, maintenance, cloud security, IT and more. Not all of these teams are educated in security matters, and some may feel that security is a constraint that slows their work. Conflicts between teams can be frequent, which means that automation is not always adopted.

And return on investment remains a priority for business leaders. Every department manager wants to be able to present to their superiors how automation helps cut costs. This is something that takes time to achieve.

A well-connected network map

To get Zero-Touch automation up and running, a detailed network connectivity map is a must. Establishing a comprehensive network map helps you to understand all of your network components and easily validate that the entire network is complete.

We developed a network map completeness score to help give organizations an indication of where they stand. If the map score is relatively low, the map completion tools can show what gaps are detected and provide a recommendation of what needs to be done.

Once your network is fully connected, automating each step of the workflow will help you to move towards the ideal of Zero-Touch automation, whether you are running on cloud based, hybrid or on-premise networks. As you come to trust automation, you will be able to automate more phases of the workflow until you reach Zero-Touch.

A typical network change workflow

There are several steps that organizations need to take towards complete network security automation, from a simple change request through to implementation and validation. Let's take a look at the most common steps in establishing automation for a simple change request -

Step 1 - Request a Network Change
Every step towards automation begins with a request for a change. A business application owner, not necessarily aware of the firewalls, network map automation or security constraints, will see the current workflows and use this to make a change request in a language he understands.

Step 2 - Find relevant Security Devices
Once this request is translated, the change automation platform will handle the request and implement the changes to hybrid networks. The network admin will be able to see the firewall and routing devices involved in the change request that was applied by the application owner. He will also see the network devices organised on the map to help with understanding which are blocking traffic and which allow it, and where the change is required.

Step 3 - Plan Change
Change automation platforms know how to deal with different firewall and device vendors' specific settings and how to implement the requests in an optimal way to avoid creating any duplications.

Step 4 - Risk Check
The administrator will get a 'what if' analysis, which risk checks the change. In this phase, the decision as to whether to confirm the change and expose the network to the risk mentioned is in the hands of the network admin or security manager, depending who is handling this phase

Step 5 - Push Change to Device
Once planned changes are approved, the 'magic' happens. The change automation platform implements and pushes the changes to the desired devices automatically, either through APIs or directly to the device (CLI).

This is a fully automated action that can be conducted on multiple devices, whether cloud based or on-premise. The push can be done in a scheduled manner, in your maintenance window or on demand. Without this, pushing the changes to multiple devices individually, sometimes on an hourly basis, can be a tedious task.

Step 6 - Validate Change
At the end of each request, the solution will check that the request was successfully implemented across all devices.

The solution also provides ongoing audits of the whole process, enabling easy checking of each stage.

Step 7 - Documentation and Logging
Network security automation platforms have the ability to provide you with a full, automated audit trail. Documentation happens on-the-go, saving IT and security teams time and accelerating tedious network compliance management tasks.

Another case is troubleshooting, which may be required from time to time, for example at the validation phase. Without a documentation trail, it can be very cumbersome to try and reverse engineer steps.

Put your trust in network automation

Over time, you can let the automation solution run handsfree, as you conduct more changes using it and gain trust through increasing automation levels step by step. Soon, you will have reached the automation ideal of full, Zero-Touch network change automation – eliminating network 'grunt work' and the risk of accidental misconfigurations, and the business disruptions and security holes they can cause. Zero Touch also ensures that the balancing act between security and business continuity is maintained, reducing risks while enhancing overall speed and agility.

The impact of innovation on cybersecurity

Digital transformation needs security at heart, says Jonathan Whiteside, Principal Technical Consultant at Dept.

TECHNOLOGY is released into the hands of consumers through two distinct ways. Arguably, the right way is to take time to polish the product before it hits the market, however, the most popular way is to release early, release often. Led by tech giants Apple and Microsoft, this philosophy is tied to building a brand following by putting innovation at the forefront. It sounds like a win-win; businesses are able to launch more products, more frequently and consumers receive digital breakthroughs as they arise. If there is a bug in the software, an update is launched to fix it and, likewise, if there is new feature requirement. The business is able to essentially conform to the user's needs as they evolve. This all sounds well and good, but is there a trade-off?

By now, we've all heard the horror stories tied to data breaches and cyber-attacks that have resulted in many European companies receiving penalties in connection with GDPR. As companies subsequently take proactive measures to improve cybersecurity, simultaneously, the number of organised hacking groups are increasing and their tactics are getting bolder. When innovation takes priority over user's security, people are put at risk. In this new digital age is the fast-tracked route to market still savvy enough?

Apple as Food for Thought

In late August 2019, it took Apple a week to release an emergency fix to a vulnerability allowing malicious hackers to take control of all Apple desktop and laptop computers, mobile devices (iPhone, iPad, and iPod touch) and also TV set-top boxes that are running the latest version of the company's software. Hundreds of millions of users internationally were placed in a compromising position by Apple. There is an ethical and legal obligation for all products, tech or otherwise, to be fit for purpose. Whether its hardware or software, if its material or codebase is faulty and in due course puts users at risk, the onus is on the manufacturer to rectify any wrongdoings.

More innovation means more threats

Now is the peak of technological disruption and this exciting period is expected to last throughout the next decade, as new innovation rapidly emerges and gets introduced into society. Let's take a look at some of the latest advancements: artificial intelligence; virtual reality; cloud computing; strategic automation; internet of things; voice search; facial recognition; 3D printing; robotics; drones; blockchain; autonomous vehicles; smart buildings... the list goes on and on. Innovation is improving lives and transforming how business is conducted. As technology changes, so does the realm for hacking. When cloud storage was released, hackers rejoiced as more valuable details were accessible on the internet, making their 'jobs' easier. All of these recent tech innovations provide new gateways for hackers to connect and explore the ins and outs of its users. Without layers of security, that's on par with the capabilities with these hackers, users' data will be an open book.

Cybersecurity skills gap

There is a general digital skills shortage globally, and cybersecurity skills are a particular challenge, since the role profile constantly changes to reflect breakthroughs in new technology and user requirements, as well as laws and legislation. This means the cybersecurity workforce needs to constantly be re-educating themselves and tweaking their approach to mitigating risks before they arise. And what one organisation deems cybersecurity, another will weigh heavily on the other side of the spectrum. The terms 'cybersecurity' and 'threats models' can often be subjective. There aren't any formal qualifications for cybersecurity or trade governance and, like most of the tech industry, there is a lack of diversity.

Security Software Developer

Malware-attacks reached an all-time high at 10.52 billion last year, according to the 2019 SonicWall Cyber Threat Report. And with this many threats, it's no surprise that organisations are being breached at an unprecedented rate. Up until recently, many leading tech companies were solely reliant on their coding teams to build resilient systems. However, as hackers become more proficient and the consequences of launching vulnerable systems more stringent, the need to add an extra layer of cybersecurity is essential. System architects can no longer keep with the pace of innovation and be on the defence. Specialist developers and analysts are increasingly being introduced to utilise security-friendly scripting language skills and have a good level of knowledge around APIs. For each phase of the software development lifecycle, they include security analysis, defences and countermeasures so as to end up with strong and reliable software. They're also actively upskilling architects to code in new ways and approach solutions differently.

The role of an ethical hacker

Cybercriminals are not just tapping into loopholes, they have sophisticated skills and a high level of intelligence capable of decrypting some of the world's most advanced systems. Their coding concepts are lightyears ahead of the average coder, and they're fuelled by criminal gain. The best way to outwit a hacker is to join them, or at least think like one. Ethical hackers know how to find and exploit vulnerabilities and weaknesses in various systems – just like a malicious hacker. In fact, they both use the same skills however, an ethical hacker uses those skills in a legitimate, lawful manner to try to find vulnerabilities and fix them before the bad guys can get there. An ethical hacker's role is similar to that of a penetration tester, but it involves broader duties. They break into systems legally and ethically. This is the primary difference between ethical hackers and real hackers—the legality. Ethical hackers are usually brought in to review systems after they've been hacked to showcase vulnerabilities, or before a product is launched to ensure it is fully ready for the public.

As digital leaders, we're in a position of power and with that power comes responsibly; to manage the risks that come with the rewards of innovation. At Dept, global digital agency, we understand advances in digital technology, particularly in the fields of AI, machine learning and IoT, will continue to unlock a wealth of new services, industries and business models. With the change, however, comes a need for trust. Digital transformation is built on a foundation of trust of which cybersecurity is an important part. If cybersecurity is considered at the start, digital transformation can actually improve a company's security posture and not detract from it.

Cloud:

Keeping transformations secure

Most mid-sized to large enterprises have already moved some of their infrastructure, data, and workloads into the cloud for better agility and efficiency. Nearly three-quarters of businesses are running a hybrid and/or multi-cloud strategy today, according to Forrester Research.

BY KUNAL ANAND, CHIEF TECHNOLOGY OFFICER AT IMPERVA.

The rise of cloud transformation

While many organisations are moving to the cloud, they may not be ready to move their data. One reason for this may be due to the lack of understanding of the security mechanisms and capabilities organisations need when they make the migration.

Cloud migrations are often part of larger corporate digital transformations that include the adoption of DevOps strategies, microservices, APIs, containers, and more. Security is rarely the driver — though it may be the moast important passenger. To make cloud transformations as efficient and successful as possible, companies must remain secure and compliant throughout.

To do so, organisations must standardise security practices across cloud, hybrid and multi-cloud assets, use modern security platforms built for the cloud automation era and use Defense-in-Depth to protect APIs, applications and data, wherever they reside.

What to ensure when securing cloud migrations

Every company's business transformation is different and performed at a different pace. The environment sometimes dictates your security tools. When an organisation has a choice, it can be quicker to achieve standardised controls through a comprehensive solution. This way, organisations can enable complete visibility across their enterprise.

Today's cloud-enabled enterprises strive to be agile, collaborative, highly automated, and efficient. Manually moving workloads and technologies to the cloud is a step backwards, being slow, labour-intensive, and error prone. This can ultimately lead to more security vulnerabilities, as well as wasted time and money. Therefore, modern enterprises are turning to rebuilding or refactoring business applications on microservices and cloud technology. To protect cloud infrastructure, security solutions must protect critical APIs and manage access to them by applications and users, including privileged insiders.

One of the benefits of an on-premises-only infrastructure is the ability for security teams to lock it down and minimise the attack surface. There is a massive cost to the business, though, as this can greatly hamper employees' productivity, and their ability to innovate, partner, and act on business opportunities. If this is not executed securely, migrating to the cloud can cause organisations' threat surface to balloon, exposing them to a potential explosion of attacks and leading to breaches whose financial damage outweighs all of the cloud-earned gains. To stay ahead of threats while protecting cloud migration, organisations need a multi-layered security architecture that provides autonomic Defense-in-Depth.

Learning from mistakes

In 2019, we learned some hard lessons about securing cloud migration. We announced a security incident that affected a subset of our Cloud WAF customers. We conducted a thorough investigation with internal security teams and outside forensics specialists to determine the root cause.

Our investigation identified an unauthorised use of an administrative API key in one our production AWS accounts in October 2018, which led to an exposure of a database snapshot containing emails and hashed & salted passwords. From this experience, we gained some valuable insights. Firstly, when an organisation responds to a security incident, organisations should operate honestly and transparently with all stakeholders. When this is communicated quickly and early in the investigation process, customers are able to make informed decisions and act on the security measures recommended. From our experience, this openness and sincerity was appreciated by our key customers and partners.

Second, organisations should focus on being fact-driven in their communications to employees, customers, partners and the community, which means organisations must confirm findings and assessments (and take actions to protect customers) in order to responsibly share additional details.

Third, organisations should establish security incident workflows and processes adapted to the hybrid cloud environment. Finally, organisations should take the time to understand the shared responsibility of deploying and managing applications and data in Infrastructure as a Service (IaaS) solutions.

As a company, we regret that this incident occurred and have been working around the clock to learn from it and improve how we build and run Imperva. Security is never "done" and we must continue to evaluate and improve our processes every single day. Our vision remains the same: to lead the world's fight on behalf of our customers and their customers to keep data and applications safe from cybercriminals. Now, more than ever, we commit to our vision, where data and applications are kept safe.

> One of the benefits of an on-premises-only infrastructure is the ability for security teams to lock it down and minimise the attack surface. There is a massive cost to the business, though, as this can greatly hamper employees' productivity, and their ability to innovate, partner, and act on business opportunities

Managing security in uncertain times – how to deal with constant change

When Heraclitus said that, "The only constant is change," he wasn't to know that his words would resonate more than ever, 2,500 years later. In fact, the pace of chance has continued to get faster and faster over time.

BY MARCO ROTTIGNI, CHIEF TECHNICAL SECURITY OFFICER EMEA, QUALYS.

AS COMPANIES SPEND MORE and more on digital transformation – around $6 trillion over the next four years, according to Eileen Smith of IDC – the rate at which change affects business IT teams will grow too.

For security teams working in these environments, keeping up with change can be difficult. There are several trends at work that can affect security, all of which can come into play at any time:

- **Business is more global** – companies are looking beyond their own home markets and country borders for growth opportunities earlier in their development. Rather than taking years to get into the position to expand worldwide, companies can sell worldwide much earlier thanks to the Internet. This in turn supports faster development of local offices that are closer to customers around the world.
- **Business is more dynamic** – alongside organic growth, mergers and acquisitions to support growth and market expansion can support that imperative for expansion into new markets. Each of these can lead to fundamental changes in IT that will affect security.
- **IT is more dynamic too** – supporting digital transformation involves more changes in approach. Digital transformation plans rely on adopting new and more agile approaches to IT like cloud computing and software containers, but these moves also have to be kept secure over time.

With more ephemeral IT in place and constant evolution taking place, security has to plan ahead in order to lead around all this change.

Approaching vulnerabilities

To manage this more effectively, you have to look further into the future. Instead of being gatekeepers and the owners of risk management, security departments have to embed their best practices across all the stages involved in the software development lifecycle. This can help developers, testers and operations teams ensure that whatever changes are made can be tracked and checked for potential vulnerabilities.

The first element to consider here is technical. Whenever any vulnerability is found across the organisation's IT assets – whether these are more permanent assets like PCs and servers, those out in the cloud, or ephemeral assets like containers that can be launched and torn down automatically based on demand levels – that vulnerability will have to be flagged for investigation. Having a complete and continuously updated asset list is the most effective starting point for all your operations, whether these are internal, external, or spread across multiple countries.

The sheer volume of issues that could come up makes prioritisation more important than ever too. While it might be simple to say that every asset everywhere should be patched and kept up to date, the reality is that this might not be practical to treat every vulnerability the same. Some issues are incredibly risky and widespread, and therefore need urgent attention; others are so niche and hard to exploit that they can be put to the back of the queue. In between, there will be many threats that have to be analysed and understood in context, as well as misconfigurations that can lead to insecurity.

To work on these vulnerabilities, security teams and software engineers should look at the risk levels associated with issues that are discovered, how widespread those problems are, and how difficult they are to exploit. These measures can be then used to create risk profiles and approaches that are personalised for the organisation, showing which issues should take priority based on specific requirements. This can create a workflow that teams can follow, providing information to people based on their own requirements for operational efficiency.

The data involved can also be used to inform management teams what risk levels are at any given point, as well as demonstrating where additional resources may be required. Based on this understanding, you can think more about your response over and above the vulnerability detection and management phases. This can then be used to automate some of the steps involved in handling vulnerabilities as they are discovered, further improving those workflows.

For companies with rapidly growing services, or that have multiple business units in place around the world, this level of insight can be particularly valuable. Not every location will have technology specialists in place – some may be small locations with tens of staff, while others might be large offices with hundreds or thousands of people. Nevertheless, each location and business unit will have its own IT assets and priorities – tracking these effectively and keeping them secure will rely on good data that is continuously kept up to date.

Change or be forced to change

The second element to consider here is process. Typically, security has been a department within IT that has had its own goals to meet around managing risk and keeping data secure. With more companies moving their approaches online and going through digital transformation initiatives, this silo approach is no longer viable. Instead, security should be embedded within all new software engineering processes so that it is a de facto standard.

For security teams used to acting as gatekeepers and arbiters of how things should function, this is a very different mindset. It relies on collaboration and understanding of how to achieve goals jointly, so that all elements of the software development lifecycle include security as standard. To achieve this, security teams have to work on getting their tools and approaches adopted into the workflows and processes that developers, testers and operations teams use. Rather than getting in the way, this allows security to be part of overall improvements in approach.

This also allows security teams to help those teams to help themselves around problems. By making security scanning and information on vulnerabilities available earlier, any issues can be flagged for fixing and scheduled quickly. This means that developers have ownership of any bugs or problems in configuration that could cause software security issues, rather than this being an issue between departments that lead to standoffs or longer times between detection and remediation. To complete the circle, security teams can also use best practices like the Center for Internet Security (CIS) Benchmarks, so that all activities can be validated over time.

Embedding security into how software is produced does mean that the way in which security teams collaborate with other departments will change. It involves more discussion early in the process, more insight into what demands are being placed on the teams involved, and more work to make security part of those processes in as easy a way as possible.

However, the resulting changes will make it easier to keep the evolving IT stack secure, whatever new demands are placed on it. It makes it easier to track and respond to vulnerabilities based on your specific technology choices and risk profile. And it makes it easier for security teams to be leaders in these times of change, rather than struggling to keep up around decisions others have made.

While change might be constant, our approach to change is what we decide it to be. For security, leading change is more effective than being forced to follow.

What is fluid security and how does it manage risk in hybrid networks?

The reality of hybrid networks within enterprises will likely remain in place for many years. These large, sprawling and fragmented environments need to be protected by a strong security strategy that enables them to effectively manage risk throughout the organization.

BY AMRIT WILLIAMS, VICE PRESIDENT OF PRODUCTS AT SKYBOX SECURITY.

THIS STRATEGY needs to be as flexible as it is forward-thinking; it needs to ensure that technology, processes and personnel do not fall into redundancy. Fluidity is the central concept that props up successful hybrid security strategies; fluid security can expand to meet new demands and apply to new technologies and to new vendors, all without a great upheaval at each change. Achieving fluidity is particularly important in the current era of digital transformation. New initiatives are being spun up all of the time, often deployed as rapidly as they're conceptualised. To ensure the security of any digital transformation project, security teams need to work more fluidly. They need to be able to secure their seat at the table and establish security as the foundation of any new initiative.

Data managed properly

When organisations achieve fluid security, it means that IT teams are able to integrate data from a variety of sources, environments and vendors into a central hub for normalisation and amalgamation into clean datasets for analysis. Proper handling and management of data, such as the removal of duplicates, is the first step in simplifying and centralising complex and fragmented networks. After that, teams can create an always up-to-date model of the hybrid network infrastructure, security controls and vulnerabilities and threats to reveal the relationships between IT systems and unify teams with a comprehensive overview of the attack surface.

Operational silos eradicated

A hybrid environment model created as part of a fluid security programme enables security professionals to eradicate operational silos. When individual teams are responsible for different areas of the network, there is a potential for processes and information to become disconnected and therefore failure to fully secure the network. For example, operational silos present in a DevSecOps team could mean that a change made by someone responsible for one part of the network could have a knock-on effect on another part of the network managed by another team.

Of course, everyone has their own day-to-day tasks, but these processes must be aligned to a common goal: securing the business' network effectively. With thorough visibility of the attack surface, any changes to new or existing services can be monitored for how their regulatory compliance or risk status might alter. A full view also allows security teams to discover and analyse new vulnerabilities as well as test accessibility, security tags, cloud firewall rules and cloud configurations. This means everyone can be kept 'in the know' and any nasty surprises can be avoided. An offline model updated frequently via an API, security and operations personnel do not need administrative access to the cloud and can go about their tasks without significant disruption to cloud deployment teams. If a security risk or violation arises, security and operations teams can report back to DevSecOps and work in tandem with them to eliminate the issue.

Ongoing cloud cyberhygiene processes undertaken
Fluid security also involves ongoing cyberhygiene processes that help IT teams reduce risk and avoid situations of compliance violation. That is because cloud services usually have short lifecycles and so may be subject to a 'set it and forget it' mentality by IT teams when deploying them. Unfortunately, with DevOps founded on replication, there is a strong chance that risks are also replicated. Compared to on-premise assets, the risks could be duplicated quicker and on a far wider scale. With fluid security, cloud services will be treated with the same vigilance as other infrastructure, so no matter how hybrid the network, it can still be secured.

By approaching security with a fluid strategy, IT teams can lay the groundwork for a proper security programme – one that not only secures complex networks today but also tomorrow when that complexity and fragmentation has further evolved. Security teams can ensure data is handled properly, model for a range of security management processes and unify teams to gain a complete overview of the attack surface. Cloud may be a 'must-have' now, but in five years another technology may supersede it. A fluid security model will allow businesses to stay secure now and prepare to remain secure in the future, no matter what changes are made to their hybrid networks.

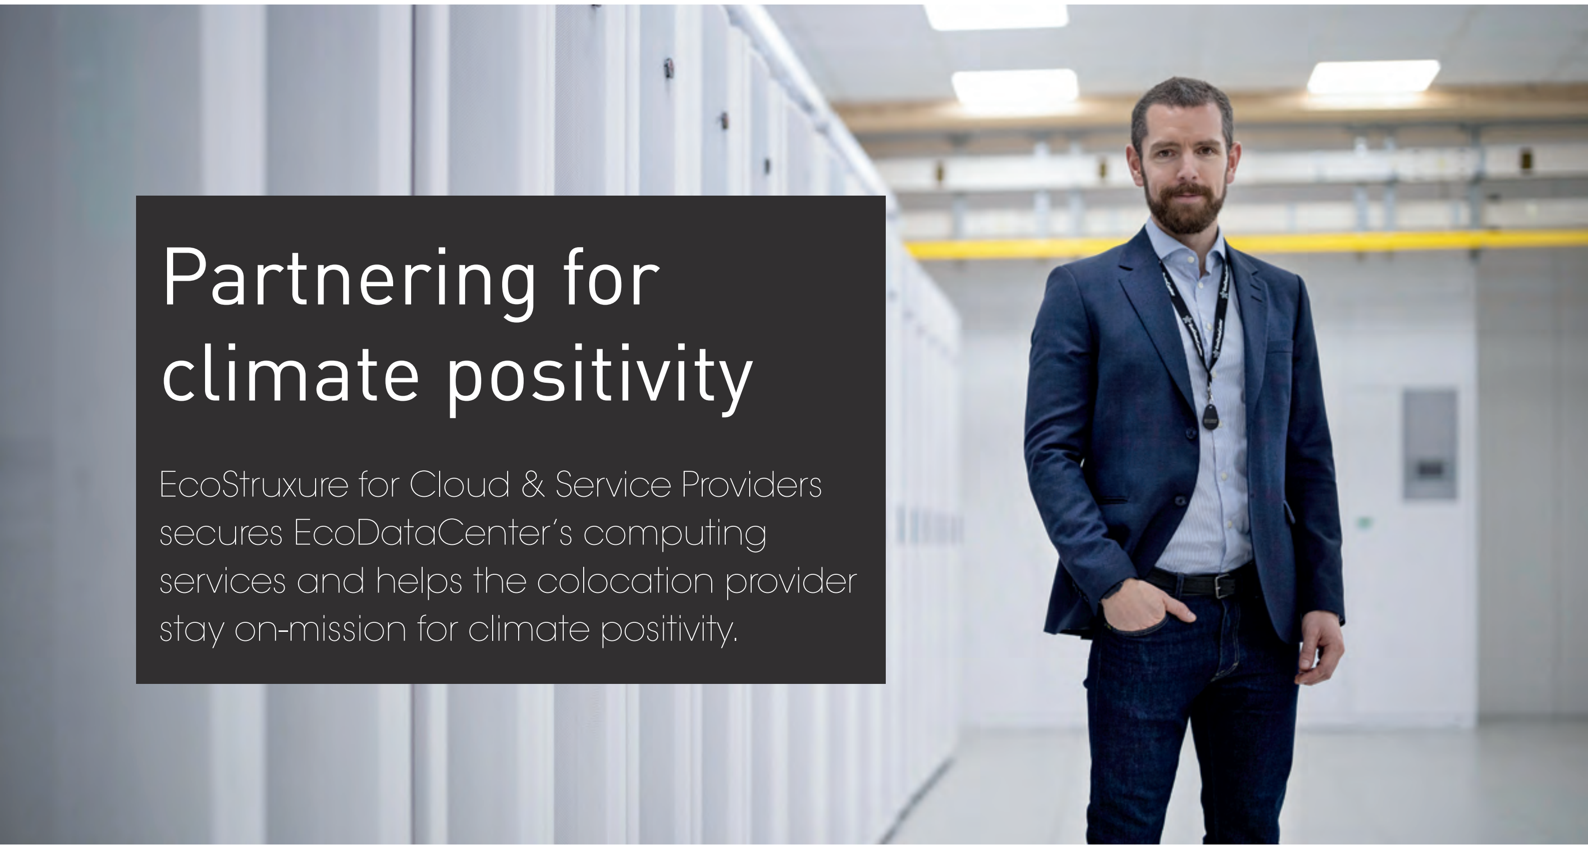

Partnering for climate positivity

EcoStruxure for Cloud & Service Providers secures EcoDataCenter's computing services and helps the colocation provider stay on-mission for climate positivity.

SCHNEIDER ELECTRIC partners with EcoDataCenter to build an ultra-low-carbon-footprint data centre at the heart of their HPC colocation in Falun, Sweden. With the goal of being one of the most sustainable data centres in the Nordics, Schneider Electric's EcoStruxureTM Building Operation, Galaxy VX UPS with lithium-Ion, and MasterPact™ MTZ are just some of the solutions ensuring the colocation facility stays on-mission for climate positivity. Additionally, Schneider's Connected Services Hub provides remote monitoring and 24/7 access to EcoDataCenter's critical facility.

In need of reliable power management that would ensure customer-server uptime and energy efficiency, EcoDataCenter partnered with Schneider Electric. "With the solutions coming from Schneider Electric, we expect to achieve a PUE of 1.15 and combining that with renewable power, we will make sure that we are one of the most sustainable data centres in the Nordics and hopefully in the world," says Mikael Svanfeldt, Chief Technology Officer for EcoDataCenter.

Sustainable digitization with EcoStruxure™ solutions Schneider Electric integrated EcoStruxure for Cloud & Service Providers into EcoDataCenter's new build. With EcoStruxure as an open architecture, EcoDataCenter was able to seamlessly connect existing hardware, software, and monitoring through the platform's analytics and services. EcoDataCenter now has connected sensor and meter data generating analytic reports on the data centres operational efficiency — and its sustainability index. Connected Services Hub now remotely monitors critical sensors, and also gives EcoDataCenter 24/7 expert remote monitoring and troubleshooting by Schneider Electric's team of service engineers. When building a data centre, EcoDataCenter considers the well-being of the community. With the monitoring, efficiency, and connectivity that EcoStruxure provides, EcoDataCenter is able to recycle the facility's heat waste to local utilities.

EcoDataCenter uses that efficiency to give back to the local grid, allowing them to go beyond net-zero emissions and achieve climate positivity.

Two Uniflair™ Turbocor Chillers add to facility climate control, supplementing the usually low Nordic temperatures, especially during summer's humidity. Within EcoStruxure's architecture, four Galaxy VX UPSs support 1,250 kW of customer load in its 99% efficiency ECOnversion Mode. This backup power frees up energy that can be diverted to clients.

"Schneider Electric is honored to partner with EcoDataCenter to deliver its vision to be the world's first climate positive colocation data centre," said Christina Backlund, VP Secure Power, Nordic and Baltics at Schneider Electric.

"Working together, EcoDataCenter and Schneider have proved it is entirely practical for a well architected data centre to be both efficient and resilient whilst meeting sustainability goals, in this case benefitting the environment as well as the local community."

Designing data centres for efficiency yields tangible benefits for Enterprise IT

When deploying IT infrastructure to support an enterprise organisation, there are several factors to consider. All of the decisions are based on the specific needs of the business, it's type, the customers it services and the applications it needs to support.

BY NICK EWING, MANAGING DIRECTOR, EFFICIENCYIT.

FOR MANY, the first decision however, is whether to own or outsource?

Do the needs of the business require predominantly commercial applications, which can be delivered via the cloud? Or does the organisation depend so crucially on local IT assets for performance, data sovereignty or application speed, that it is more preferable to keep their resources on premise? As always, cost, mission criticality and performance are the ultimate factors determining the decision, but if it is decided that control of its own IT assets is essential, then an organisation's task will become focused on designing and building a data centre tailored to meet these demands. In most cases this leaves one final question. How does one deploy the most resilient, secure and operationally efficient data centre solution, in the most cost-effective way?

Determining the value of your assets

Total Cost of Ownership (TCO) is the cost of an IT installation over its working life, including the cost

of acquiring and deploying all critical infrastructure components including Power, UPS, IT and Cooling. Whereas Operating Expenditure (OpEx) is the cost of running, maintaining and servicing the plant throughout its life.

Taking a longer-term and strategic view may be more advisable for an organisation who are looking to extract the maximum value possible from their IT assets. Some businesses may opt to turn them into an investment, if you will. However, this means that the organisation may also need to adjust their calculations and speculate at the outset in order to accumulate a better return during operations.

For many enterprise organisations, building a strong case for investing in new and efficient technologies at the design and specification stages can yield tangible benefits to the bottom line over the long run.

Design and build considerations

When deciding to build a data centre on one's own premises, the first consideration is the real estate that an organisation possesses. Does it have sufficient space on site for its IT needs? Is the space available in a single hall or room, or are there multiple locations across which equipment may be deployed? Is it wise to keep all IT assets in the one place or can they be efficiently dispersed throughout one's campus?

The answers to these will naturally be informed by the structure of the organisation. For example, a research organisation needing access to High Performance Computing (HPC) clusters to process large volumes of complex, algorithmic data calculations will likely require high-density racks, large volumes of storage and a heavy cooling requirement that would not be needed for more general-purpose servers. Isolating such resources within their own space, alongside the specialist support infrastructure they require might be more cost-effective from the points of view of acquisition and operation.

Choice of cooling solution

Power and cooling account for much of the operating costs of a data centre. However, Life Sciences in particular, are continually evolving and so do its technology requirements. For HPC clusters to keep up with the demands of new GPU-based systems, the owner may, for example, choose to deploy a form of liquid cooling in order to drive efficiency and lower OpEx.

Liquid cooling has fast made a comeback as a way of maintaining optimal operating temperatures with far lower power consumption than air-cooled systems. Whether using fully immersed methods or direct-to-chip options, liquid cooling can reduce operating costs by up to 14%. There is, of course, a penalty to pay in terms of costs to acquire and install such features. However, for HPC racks, the pay back over time is significant.

For general purpose data centres, the choice of cooling system may be affected by the constraints of one's existing facility. It may be cost-effective to make use of an existing raised floor if one is available, or if there is sufficient head room in the server room to have such a floor installed. Similarly, hot or cold aisle containment systems, which improve the efficiency of air-cooled installations, should always be considered if there is sufficient space for them.

Power and management systems

The power requirements, especially the uninterruptible power supply (UPS) will be determined by several factors including the criticality of the systems under load, the quality of the existing power supply and of course, the cost. Today, data centre operators may examine the trade-off between traditional UPS systems based on Valve-Regulated Lead-Acid (VRLA) batteries, and Lithium-Ion (Li-ion) powered models.
The latter offers numerous advantages, including a much smaller footprint, longer lifetimes that incur lower service and operating costs, and greater power densities.

What's also important is to ensure the user right-sizes the UPS system. This may also mean they need to spend more in the initial stages by choosing a modular UPS, but the payback will result in long term dividends, allowing them to save larger costs in terms of energy usage.

Underpinning all of the efforts to optimise the efficiency of a data centre, whatever type of hardware assets are deployed, is the Data Centre Infrastructure Management (DCIM) system. Or more recently, next-generation DCIM systems that offer increased visibility, with remote monitoring and management capabilities via Artificial Intelligence (AI).

With the IoT technology now built into all data centre hardware assets including racks, UPS, IT and cooling systems as standard, there has never been more insight available into the operating status of a facility. These new systems also provide the user with the ability to manage an entire portfolio of data centre assets inside a single-console, viewing all of the critical systems and components from anywhere, at any time.

Newer systems such as Schneider Electric's EcoStruxure IT™ can alert the user to maintenance schedules and leverage Machine Learning to identify performance issues before they become problematic, increasing resiliency. Moreover, they have the ability to collect performance data for analysis and facilitate continuous operational improvements so that efficiency can be maximised and operating costs lowered.

According to Andy Lawerence of the Uptime Institute, "The average power usage effectiveness (PUE) ratio for a data centre in 2020 is 1.58, only marginally better than 7 years ago." However, the metric remains an important aspect of IT operations today, as any improvement in PUE and thereby energy efficiency will have a positive impact on OpEx.

For many Enterprise IT operators, a trade-off has to be made regarding mission criticality, as the more infrastructure an organisation deploys to increase the resiliency of its IT systems, the more impact there will be on energy usage. DCIM, however, remains an essential component of all data centre operations today.

Conclusion

As with any business, the adage that what gets measured gets managed applies particularly to data centre operations. The drive for greater efficiency is motivated by many factors including environmental concerns, running costs, government regulations and indeed the need to reduce operating expenses.

As always, the Total Cost of Ownership, driven ever lower via a greater investment in efficient technologies during the design stage, will inevitably impact positively on an organisation's financial bottom line, whilst ensuring their critical IT performs resiliently and exactly as expected. That, of course, is as compelling a business case for efficiency, as any.

For more information on DCIM and remote monitoring, please click here or visit - https://efficiencyit.com

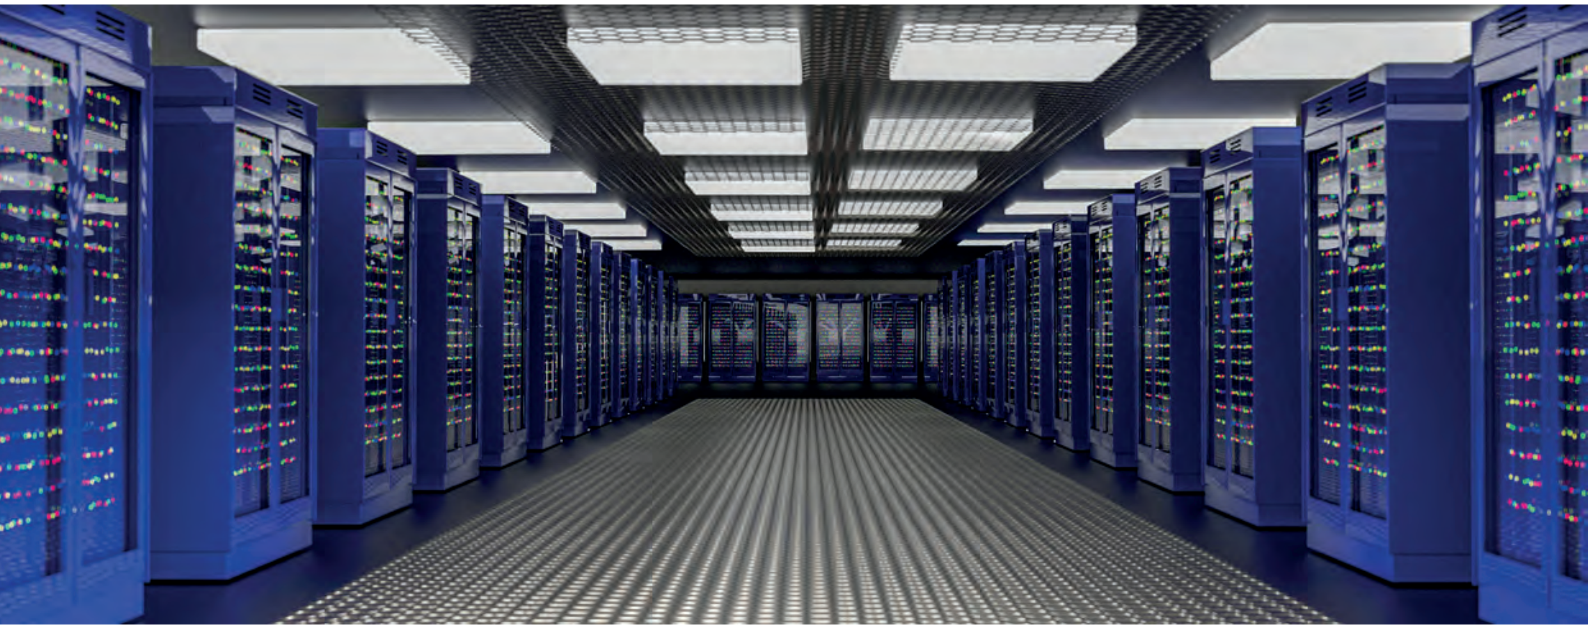

Offloading data centre maintenance to the cloud to support remote working

Forward-looking businesses are offloading data centre maintenance and management to the cloud to keep things running during the COVID-19 pandemic.

BY RAJA RENGANATHAN, VICE-PRESIDENT AT COGNIZANT AND HEAD OF CLOUD SERVICES BUSINESS

FOR YEARS, proponents have urged businesses to better enable employees to work from home, citing benefits like increased productivity, less commute time, better work-life balance and enhanced preparedness for business continuity, should a localised disaster strike, such as a tornado, hurricane, earthquake or flood.

Overnight, the COVID-19 global pandemic made the final argument for work-from-home a reality for millions of workers – ready or not. Many global enterprises must suddenly support more and more people working remotely, whether they are equipped to deliver and support workloads at scale or not. This has sent businesses scrambling to quickly embellish digital channels and platforms, increase bandwidth, add virtual private networks (VPNs), provision more laptops, and offer thin-client applications to their employees and customers to improve operational collaboration and enforce social distancing.

A proper business continuity plan followed up with precision execution can ensure that enterprises deliver such capabilities. However, what happens to the on-premises data centre where a physical presence is required? Even with workplace virtualisation technologies like remote consoles and "out-of-band networks", which reduce the need for on-site data centre operations staff, the fact is, physical boxes in on-premises data centres still need to be managed, guarded and secured by people.

Take the February 2019 data centre meltdown of a major U.S. bank, which crippled the organisation's online and mobile banking capabilities. The company needed to shut down one of its data centre facilities due to a smoke condition. It took two days to bring the facility back up, and only with significant effort, which required the physical presence of data centre staff. Imagine if this happened during the COVID-19 crisis. The time taken to fix the issue would increase exponentially due to a lack of people resources and hesitation to collaborate in-person. Even physical security could become compromised, which raises grave concerns.

An increasing dependence on the foundations of IT

The fact is, as our dependence on IT intensifies, data centres have become the substratum of how we live, work and play. From banking to insurance to 24×7 news, everything is supported by cloud infrastructure housed in virtual data centres. If these data centres go down, critical business functions, financial networks and in some cases our whole way of life become threatened. As a result, virtual data centres need to be continuously supervised and constantly cared for. The Uptime Institute's 2019 Data Centre Survey puts this into context, some of the findings include:
The staffing problem affecting most of the data centre sector has become a crisis. Sixty-one percent of respondents said their organisations had difficulty retaining or recruiting staff – up from 55% a year earlier.

Outages continue to cause significant problems for operators. Just over one-third (34%) of all respondents had an outage or severe IT service degradation in the past year, while half had an outage or severe IT service degradation in the past three years. Ten per cent of all respondents said that their most recent significant outage cost more than $1 million. (The study authors note that "most recent" could have been at any time in the past.)

Why transfer maintenance responsibility to the cloud?

Forward-looking businesses are taking a different approach – they are offloading data centre maintenance and management to the cloud. One reason for this is scale, as cloud service providers (CSPs) have mastered the art of managing scale. In addition to proactively planning for capacity, businesses can leverage auto-scaling features to rapidly meet any unplanned surge in demand.
is also highly automated and allows for the creation of scaling policies that set targets and add or remove capacity in real-time as demand changes. Thus, utilisation and costs are optimised, while the need for having more people on the ground is reduced.

Most CSPs now provide a multi-tenant architecture that allows different business units within an organisation or multiple organisations to share computing resources. This allows organisations to optimise their resources and staff vs. having their own data centres. Lastly, the physical security in and around CSP data centres tends to be more robust and proven than what enterprises can individually afford. Most CSPs have rigorous and ongoing processes for assessment and mitigation of potential vulnerabilities, often performed by third-party auditors.

Planning for disruption

With many proven methods and tools, cloud migration is involved, but it is not difficult. For businesses that take a meticulous optimisation approach, the cloud can be more cost-effective than CapEx-hungry data centres. Even in the context of coronavirus, digital platforms running on the cloud can unleash cost and operational advantages via centralised control while meeting bandwidth challenges that flare up during peak usage periods.

Perhaps it is our hyper-connected world, but severe disasters seem more frequent than ever. As businesses respond to the many challenges COVID-19 presents, they also need to keep their eye on the horizon to prepare their data centres to withstand any disaster that strikes in the futures.

Transformative powers of IoT data mining for digital facilities

The smart building market is expected to register a CAGR of over 23% between 2020-2025. Growing energy concerns, increasing government initiatives on smart infrastructure projects, are the drivers of the market's positive growth.

BY MICHAEL C. SKURLA, CHIEF TECHNOLOGY OFFICER FOR BITBOX USA.

AUTOMATED MULTI-SITE FACILITIES' operators and managers are all too familiar with data aggregation challenges across all the building systems across distributed facilities. While the proliferation of data from IoT devices offers massive automation opportunities, harnessing IoT data mining across geographically distributed sites is critical for any enterprise. The actionable analytics can help drastically increase facilities' efficiency while reducing energy usage and operational expenses.

IoT proliferation is constant

By 2025, Business Insider forecasts there will be over 64 billion IoT devices in operation. In 2020, 90% of cars will be connected to the Internet. And by 2023, Ericsson Mobility Report puts Cellular IoT connectivity at 3.5 billion.

IoT–or the Internet of Things–is not one, isolated thing. It is a wide range of interrelated computing devices. It resides in your refrigerator, your phone, your thermostat, but also in mechanical and digital machines. With no humans or computers to mediate vast information floods, these devices continuously transfer data and information over networks.

Today's commercial facilities and buildings use a combination of subsystems - all operating in autonomous silos. With IoT technology emergence challenging data collection and monitoring, enterprise facility operators and owners have a great opportunity to reap greater outcomes by integrating new technological advances, like IoT devices, into their current systems. Enterprises can gain heightened benefits if they leverage across multiple ecosystems within the infrastructure and tap into data from multiple trades.

Once this is achieved, then the systems working in unison can simplify O+M over the life cycle of the facility and drastically lower capital expenditures while at the same time increase technological services within a commercial facility.

Don't ignore valuable IoT data

According to Cisco, by 2021 the flood of IoT generated data will reach 850 zettabytes. But only enterprises using tools to collect, organize, and deliver the data generated from IoT connectivity can reap actionable analytics and make sound business decisions.

But what's missing from legacy Industrial IoT, BMS, SCADA, and monitoring solutions is the capability to collect and normalize data across large geographical footprints in a simple and scalable fashion. For multi-site commercial facility owners needing to capture untapped IoT and facility data from distributed sites and IoT devices, this is critical.

A network of sensory devices plays a pivotal role in data collection from building management systems and individual control solutions. From lighting fixtures to you HVAC and security, the sensory devices offer critical data valuable to other trades. And when viewed together, enterprises can gain insight into their collective operations.

IoT platforms' critical role

Automated, digitized facilities are the future of our facilities' world. Maintaining market competitiveness and profitability requires facility owners and operators to harness valuable, critical and actionable data from all their IoT devices.

The IT market's solution for collecting diverse data at a massive scale is the easy to install IoT platforms. These extract data from various in-building protocols and subsystems by using a nimble setup and an EDGE appliance wired to a port. IoT platforms' open communications with a cloud infrastructure allow for remote management, provisioning, and monitoring. With no need for designated on-site staff, facility IT operators enjoy time and cost savings while gaining efficient, cost-effective operations.

Commercial building and facilities' operators and managers deploying IoT platforms are free of costly on-site commissioning of monitoring systems and gain a single pane of glass approach to all their edge subsystems and technology. Operators have access to a single source of truth of data on a portfolio for analytics, alarming, AI, reporting, or custom applications from an easy to access cloud location.

With today's enterprise infrastructure bombarded with multiple subsystems, and with an increased number of IoT sensing devices and other technologies, facility operators can use an IoT platform to leverage data between diverse systems. This means they can eliminate duplications since one sensor actually provides all the data without a separate monitoring system, wiring or programming.

Once all the data is collected across multiple sites, and stored in the cloud, facility operators can easily leverage the value of all their aggregated data using Micro-service Analytics and visualization engines. The organized data-lakes, provided by the IoT platform, easily transform the data into context-specific outcomes and actionable recommendations allowing operators to reap the ROIs of the IoT platform.

This includes using available real estate space for additional revenues, traffic patterns for supply chain, to reducing energy use – all gained via the IoT platform layered atop existing and siloed systems enhancing the performance competency of the silos.

Leveraging and storing all the relevant data, collected from various sources, allows for a consolidated "truth" for outcome-based analytics tailored to specific business objectives.

IoT platforms maintain a transformative pulse of facilities technologies and operations, providing facility data for a greater purpose-generating meaningful outcomes, beyond just the physical building operations. By aggregating valuable, and previously inaccessible information, facility operators are empowered with improved bottom line and more efficient management and operations of multi-sites armed with analytics and visualization.

On the ball

How ESPN uses BI and analytics to give sports fans the ultimate viewing experience.

BY SIMON HAYWARD, VICE PRESIDENT EMEA AT DOMO.

SPORT CONSUMPTION has seen a huge increase in recent years. With record viewing numbers for the likes of Premier League last season, and a novel willingness from consumers to pay for exclusive online content. As a result, sports is a lucrative source of revenue for network and programming brands.

Alongside this increase in popularity, modern television offers fans a more up close and personal experience at a fraction of the cost compared to attending live sporting events. Viewers can watch from the batter's eyes, and encounter the huge right hand made by MMA fighters. In many ways, the viewing experience has evolved to become what Doug Kramon, ESPN's senior director of fan support and customer care views as a 'virtually there' experience. However, in an era when TV watchers have more options than ever, and with beautiful camera angles not enough to retain viewers, broadcasters are looking for innovative ways to keep their audience engaged.

The battle for viewers

For networks like ESPN, part of the challenge of keeping viewers engaged comes down to understanding their expectations and experiences. Sports conjure up a range of emotions you simply don't feel when you buy a pair of trainers. Fans feel elated when their team scores the winning try and devastated when they miss the final penalty. At the same time, if a fan misses any part of the action this can create an emotional reaction, that often leaves ESPN's Fan Support inundated with inbound customer service issues. This information can come from a plethora of sources, including email, chat, calls, SMS, or even indirectly complaining via social media channels. The result, a vast amount of data to sift through to uncover the underlying problem.

For ESPN, the key barrier was being able to cut through the noise, to understand the live conversation and pinpoint key issues their fans were experiencing. On top of that, ESPN needed to be able to understand the severity of these issues. Was it limited to one fan, or thousands of fans?

Whether you're a sports broadcaster or a cosmetics brand, providing exceptional service is crucial to retain customer loyalty. However, this can be tricky when feedback is hitting you from every angle. So what are the key steps ESPN implemented to help provide the ultimate experience that keeps fans coming back?

Tackling the challenge

Tuning into sentiment

Understanding the sentiment of data is pivotal to identifying key issues experienced by fans. With the help of Domo and RXA, a leading applied artificial intelligence and data science company, ESPN deployed a solution capable of analysing sentiment. This worked to extract the real-time feelings coming from conversations across different customer channels. Using artificial intelligence (AI) to classify the positive, neutral and negative opinions fans were expressing.

Unifying the voice of the fan

It's not only conversation data that hits the contact centre, there is also a mass of discussion happening on other online platforms, from Twitter feeds, and personal blogs to online forums such as, Reddit. Using Domo's platform, ESPN is able to connect siloed data from conversations happening online, with data coming in directly to Fan Support, and house it all in one place. By unifying these disparate data sources, ESPN is able to map out all of the data to have one unified voice of the fan.

Real-time view

Customers today have high expectations. They anticipate that businesses will meet them where they are and when they want. As such, ESPN fans respond to experiences that are timely and tailored to their specific needs, and reject those that aren't. To meet these customer demands, ESPN deployed real-time data analytics to understand issues fans were experiencing as they happened, such as lagging, freezing or lost service.

Automating processes to identify key broadcast issues

Automation is key when it comes to customer care generally. For ESPN, they needed to not only see customer service issues in real-time, they also wanted to be able to deflect these issues from their contact centre and promote fan self-service. By harnessing BI and analytics, ESPN is able to identify keywords regarding the product, the related issues, the frequency of that issue, and the associated sentiment. All of which are then presented in a concise summary that allows leaders to act immediately. For example, sharing responses through chatbots, FAQs alerts and live site announcements to let fans know the issue is being addressed, which ultimately relieves strain on Fan Support teams.

Utilise AI-powered alerts to notify problems

Through AI, ESPN is able to track how a fan is consuming sport from various TV providers to understand what the fan in the digital seat is viewing. Using thresholds, ESPN receives AI-powered notifications and alerts via mobile when a number of issues related to a specific TV provider tips over the set limit. This provides information in the moment, so that customer service leads can collaborate with the appropriate tech/product department to get the issue resolved as quickly as possible.

> Customers today have high expectations. They anticipate that businesses will meet them where they are and when they want. As such, ESPN fans respond to experiences that are timely and tailored to their specific needs, and reject those that aren't

The ultimate result

With the help of BI and analytics tools, ESPN can now make data-enabled decisions, and truly understand how its broadcasts are being received by fans. By mobilising data and AI, ESPN customer service leads can spend less time digging around for the problem and more time focusing on how to improve the overall virtual experience for its viewers. By tapping into customer sentiment, ESPN has created a unique business advantage. The impact of all of these efforts together, has meant that ESPN's customer satisfaction is up by 9%, and customer self-service by 200% year over year.

Advanced customer analytics, at incredible speed and precision, isn't just a possibility now, thanks to an array of converging technologies - it's an imperative for companies hoping to connect tightly with their audience. Indeed, we are seeing this reflected in business behaviour. According to research conducted by Harvard Business Review Analytic Services, 60% of enterprise business leaders say customer analytics is extremely important. Companies who utilise real-time customer data are those that achieve customer retention and loyalty.

For other companies looking to rationalise and consolidate their customer data the best piece of advice I can give is to start by looking at a solution that can bring together customer data from multiple touchpoints, while effectively promoting quality data governance. Without a joined-up data solution, businesses are at an increased risk of a garbled customer view that offers no real value. With a solid data governance protocol, data is clean and free of errors which will grant the integrity needed to attain the modern consumers expectations.

Delivering beyond dashboards – understanding the power of data-driven insights

What do you think of when you hear the term "business intelligence" (BI)? Most, but not all, will immediately turn their thoughts to reports, dashboards, graphs and charts. If you spend just a few minutes walking through a BI trade show that is exactly what you'll see on the screens of the vendors exhibiting at the event.

BY PEDRO ARELLANO, VP OF PRODUCT MARKETING AT LOOKER.

WITH THE MAJORITY of organisations still of this mind-set, the concept of reporting has become inseparable from the understanding of BI. However, this means many companies today are unable to unlock the most important aspect of data – the insights.

An outdated vision of BI

There is no doubt that data continues to flood in at unprecedented levels, leading in turn to a rise in the amount of unstructured data that companies are working with. This, combined with changes in technological regulations such as GDPR, can make it difficult for businesses to process and correctly analyse all information available to them. Insights are instrumental to the success of modern organisations – by using modern solutions, businesses are able to present the most up-to-date information to the right people.

Tunnel vision BI

If the purpose of BI is to produce a report or a dashboard, that means additional steps are required before that data delivers meaningful value to an organisation. Someone has to take that dashboard, draw some conclusions from it, and take a corresponding action that has an impact – presumably a positive one – on the business. Without this action, insights from BI are useless.

These days, companies cannot risk being stuck using this outdated view of BI. However, more recently we are seeing a shift in approach; BI is no longer only being used to create and support better internal decisions, but instead the technology is embedded in operational processes.

This means that in order for companies to fully harness the power of the information at hand, we need BI solutions that deliver more than just reports and dashboards.

The real benefits of BI, beyond reporting

Organisations have the power to upskill and educate themselves on how to best use BI within business. Forrester crafted a new term for these systems that help companies move beyond reports and dashboards and can close the loop between data and action. They're called "systems of insight" and Forrester describes them as an evolution of BI and analytics that have the potential to "harness digital insights – new, actionable knowledge that can be implemented in software – and apply these insights consistently to turn data into action."

A modern approach can enable companies to deliver data-driven experiences that seamlessly integrate data into business workflows. We must recognise that the future of BI involves more than feeding charts and graphs to a data analyst.

Data-driven insights and customer experience
The future of business is data driven, and new solutions that allow organisations to experience data in new ways, by infusing it into everyday functions, are driving better employee and customer experience.

These "data-driven experiences" can be tailored to specific operational workflows, like automatically presenting a discount offer to a customer that's likely to churn, automatically adjusting bids for under- or over-performing online ads, or using natural language to ask about inventory levels in Slack and ordering additional units based on the answer.

If we drive our vision based on the understanding that business professionals don't live in a BI tool, and that a modern data and analytics solution should reach us through the applications we already use, there is an enormous opportunity to transform the way companies use data.

It is important as we move forward in technological advancements to view data as more than just something we analyse, and instead see it as an integral and actionable component of a business process, helping organizations become vastly more efficient, productive, and successful.

Fuelling a synchronised supply chain with cross-company collaboration

For twenty years or more, organisations have concentrated on a cascaded decision-making process, with each specific function focused on optimising its own particular view of the supply chain, often to the detriment of other departments or units within the business. Departmental decisions taken would often impact the whole organisation in ways that were poorly understood or not fully apparent at the time.

BY CLAIRE MILNER, SENIOR VP EMEA BUSINESS, KINAXIS.

IN THIS AGE of digital transformation, organisations are increasingly seeking out new approaches that allow them to remove the time and latency between different processes and decision points, enabling people across the wider business ecosystem to collaborate on an always-on, synchronised supply chain.

The birth of supply chain digitisation
Siloed decision-making has always been at the heart of organisational misalignment. As an example, individual cost-related decisions can have a direct effect on an organisation's overall production volume and distribution, but without visibility over the whole system, decision-makers can unintentionally cause challenges for the whole business – with the consequences only becoming apparent over time. In an increasingly digitised world, it's essential for these disparate functions to align. If the teams selling the product are not in line with those producing or suppling or the product, the resulting misalignment can create significant organisational issues and tensions.

With growing consumer demand for speed and personalisation, alongside factors such as trade volatility, tariffs and regulatory changes, it's clear that legacy supply chain systems are no longer fit for purpose. Customers are increasingly expecting flexible delivery options, personalised products and services and increased visibility, but to keep up with these trends, companies need to evolve their supply chains to match these expectations while maintaining cost-effectiveness. As such, the age of digitisation has sparked an agile transformation for supply chains – with a marked growth in the agility of supply-related decision-making, as well as an improved ability to mitigate risk and capitalise on digital opportunities.

Digital supply chain tools such as platform-based models can help create interconnected and transparent supply chain ecosystems – keeping solutions, people, data and machines connected digitally to facilitate easier management and create value throughout the supply chain. Research from McKinsey shows companies that focus heavily on digitising their supply chains can expect to boost their annual growth of earnings by 3.2 percent – the largest increase from digitising any business area – there is a clear financial benefit to this approach.

Shifting the mindset for engagement at all levels

Though it's vital to get the buy-in of supply chain professionals, shifting the mindset of the wider business through effective change management is central to success. Overall, transforming an organisational mentality is about taking every employee on a journey; communicating the individual value of the changes for them, as well as what it means for the business going forwards. Great technology is only as good as the people using it. If employees aren't convinced that new solutions will benefit them or their roles, siloed thinking and use of traditional tools such as Excel will continue, even if there are better technologies available.

Digital transformation starts with attitude transformation, and it starts at the top. Driving supply chain transformation with board-level endorsement is critical to getting the rest of the company to collaborate and change their mindset. Without that, supply chain managers will encounter resistance to change, which in turn will restrict the opportunity to achieve transformational success.

Alongside executive advocation, aligning KPIs and metrics across different functions can help every team see how their actions affect each other and the overall supply chain. However, this visibility can only be achieved if everyone in the organisation is held accountable to a single vision and corresponding KPIs. Equipping employees with training opportunities and new data and technology tools that enable them to gather insights and then share them across the organisation can also help facilitate adherence to a shared mission.

Transforming the supply chain

With outdated technologies, such as email and Excel, still present in supply chain planning, agility is fundamental to future success. By planning what to produce and ship, how much, when and where, companies are empowered to be as efficient with their resources as possible. And by utilising the latest technologies such as machine learning and artificial intelligence alongside current processes, existing practices can be easily adapted to help businesses achieve digital transformation; reducing the time and minimising the latency between different processes and decision points, and enabling every member of the supply chain to collaborate in synchronicity.

This collaboration capability allows easy understanding of a supply chain's overall performance, allowing users to synchronise all aspects of their supply chain, and as a result, see and interpret the impact of a change across the whole chain. The result is dynamic supply and demand balancing. However, the effectiveness of supply chain collaboration relies upon two factors: the level to which it integrates internal and external operations, and the level to which the efforts are aligned to the supply chain settings in terms of geography, demand, and product characteristics. As such, success relies on an organisation's specific circumstances, and identifying what they need to do to fully benefit from their collaborative efforts.

Connecting the dots with concurrent planning

Legacy supply chain planning is based on a culture of assessing individual chain links in isolation, rather than looking at them as a fully connected chain. However, this is fraught with risk: if one link breaks, the chain will likely be affected – but the other links will be unaware. Concurrent planning offers a real benefit to business, with unparalleled responsiveness, the ability to quickly react to new opportunities and mitigate new problems.

By working to continually balance the end-to-end supply chain, concurrent planning can align supply, demand, capacity and inventory, as well as synchronising sales and operations planning and integrated business planning. This facilitates the ability for everyone to share information at the same time; and with only one data set across the ecosystem, the whole chain is in sync at any given time, with immediate access to the people, information and results they need. Most importantly, this approach means that all businesses across the supply chain can work to resolve issues immediately – intercepting and dealing with problems as they arise, rather than having to manage the fallout afterwards – and allowing for smarter decisions to be made based on insights from the whole supply chain. With disruptions large and small prevalent on a day-to-day basis across the globe and across every supply chain, such as those caused by COVID-19, this connectedness, agility and speed has never been more important.

Looking to the future

As organisations increasingly embrace the need for digital transformation, changes in supply chain management are only set to continue. Tools such as machine learning are allowing businesses to glean insights from an ever-growing number of data sources, with newer and smarter ways to capitalise on the opportunities they provide. By incorporating data in a way that fast-tracks supply chain transformation and facilitates faster, more intelligent decision-making, organisations will future-proof themselves in the market, giving themselves a competitive edge against new and powerful disruptive forces."

Five key reasons to move closer to the edge

Today, the challenge facing CIOs is how to take advantage of disruptive technologies, resilient data centres and edge computing sites in order to develop an infrastructure strategy that encompasses faster technology refreshes and greater digital services than we have experienced before.

BY SEAN HILLIAR, CO-FOUNDER AND DATA CENTRE MANAGER, IP HOUSE.

THE 2020 STATE of the CIO Executive Summary by IDG reports 89% of IT leaders ‘believe the CIO increasingly needs to rely on trusted advisors to help navigate emerging technologies, processes and methodologies.1’

The 2019 Harvey Nash & KPMG’s CIO Survey, entitled, ‘A Changing Perspective’ also highlights that only 26% of technology leaders in 2019 felt “very well” prepared for cyber-attacks, a decline of 3% on figures 5 years ago2. It’s clear that with the help of a secure and reliable colocation provider or data centre partner, you can alleviate many of these risks.

At Gartner IT Symposium/Xpo exploring the top industry trends for 2020, Research Vice President Brian Burke stated that ‘edge computing will come to be a dominant factor across virtually all industries and use cases as the edge is empowered with increasingly more sophisticated and specialized compute resources and more data storage.’ As data increases exponentially, the importance of mitigating any outage risks becomes more prevalent.

However, there is a significant opportunity for forward thinking CIOs implementing best-in-class strategies to increase revenue. As Gartner also forecasted that digitally trustworthy companies will generate 20% more online profit compared to those which are not3.

London remains one of the most strategic international technology hubs. Forbes elucidates on research provided by Forrester predicting that the edge cloud service market will grow by at least 50% this year4. Independent colocation facilities such as IP House form an important conduit of near-edge services for CIOs, leading their companies towards the next phase of edge transformation.

We believe there are five key factors that CIOs should consider when looking for a colocation provider to help them get closer to users at the edge: -

#1. Location

Location is often the most important criteria when selecting a data centre to host your critical IT infrastructure. Cost vs QoS, accessibility, network latency and factors such as scalability and capacity are just a some of the areas that may directly impact CIOs ability to deliver digital services in the future.

IP House is in the heart of Europe’s data hub, E14 - London. For years connectivity providers have built their networks and services around the home of rich connectivity links in this prominent district. As the termination point and gateway to transatlantic links, it has long been the choice for service providers who understand the technologies and infrastructure underpinning the global network and routing of traffic across the world.

#2. Customer Service & Technical Advisory

At IP House we respond to our customer’s needs day or night without exception. Our customers receive a response to queries or requests within 1 hour of initial contact being received. IP House provides a variety of communication protocols with alternative back-up options.

Our technical advisory team work closely with our clients to determine their data centre strategy, both for their immediate requirements and predicted future scope. Our business ethos is founded on customer service and satisfaction, as such our client testimonials and online reviews reflect our dedicated approach. We understand the importance of rapid response and that your services are both critical to you and those that are dependent on the uptime of your equipment.

#3. Uptime | Reliability

Data centres must achieve unprecedented availability to meet the demands of business requirements. IP House is an EN50600 Accredited, designed and built to Tier III standards, with N+1 configurations throughout all critical systems, including Power, Cooling and Networks. Our facility provides 99.98% uptime and is constructed by industry-leading vendors and their ultra-reliable systems. With multiple redundancy layers and critical backup infrastructure, we take away the worries of staying online and maintaining uptime.

IT monitoring helps to detect and diagnose problems before they cause major issues, reducing costly downtime in the process. Our data centre is highly connected, and we have full visibility of all critical systems through a single pane solution, using Schneider Electrics, AI-enabled EcoStruxureTM IT.

#4 Security

IP House's high security level is defined by our purpose-built design and layout, which benefits from the omission of windows to reduce the potential risk of intrusion. Access into the data centre is restricted to authorised personnel only.

Our advanced security system comprises; 24-hour surveillance, remote monitoring, anti-intrusion measures, remote monitoring, multiple man-traps and a multi-factor access procedure with control points at every ingress and egress point. The data centre is protected by extensive CCTV throughout with ANPR tracking, intelligent motion detection and also incorporates Grade III intruder protection with perimeter detection with anti-climb deterrent.

#5. Connectivity

The data centre has been established for over 20 years and has a rich fibre network running throughout the facility. Availability of the world's leading Tier I carriers ensures you can connect to the farthest reaches on the globe, whilst ultra-low latency connectivity routes to prominent Exchanges and Peering points, together with Direct Cloud, connect to all major Service providers are available from multiple providers. Network diversity can also be configured to further protect your connectivity.

Our flexible network backbone delivers high capacities in greater densities. Pre-terminated HT RapidNet solutions provide a combination of fibre and copper connections within a single U-Space, giving our customers more room for their equipment and cabling management. This future proofed solution can also deliver 100GB speeds through elite quality fibre and termination modules.

Today, IP House provides a safe, secure, resilient and reliable colocation facility within the heart of London, right on the edge of London's financial district.
The facility is purpose built to house companies' critical infrastructure, incorporating all the necessary redundancies so that CIOs can rest assure that their business will remain online and be accessible at all times.We have also partnered with an esteemed group of industry leading vendors to provide ultra-resilient, highly reliable service offerings for our customers.

To help CIO's further in the decision making process, IP House has created a complementary Ebook discussing the Top 15 Points to Consider before choosing a colocation data centre in which to securely house your IT and network assets. Download it here today.

Further reading

# 1 2020 State of the CIO Executive Summary – CIO from IDG
# 2 CIO Survey 2019: A Changing Perspective – Harvey Nash / KPMG
#3 Top 10 Strategic Technology Trends for 2020 – Gartner Inc., 21st October 2019
#4 Predictions 2020: Edge Computing Makes The Leap – Forbes, 2nd Dec 2019

Operationalising AI:

Whimsical dream to grounded reality

AI is constantly evolving and, with every development, its value to businesses grows. As demonstrated by Accenture's global study[1] with C-suite executives, the implementation of AI is becoming synonymous with growth and profitability. In fact, 84 per cent declared AI as imperative to their growth ambitions, while 74 per cent went as far as saying they would go out of business if they failed to implement the technology over the next 5 years.

BY DR. MADHAV DURBHA, GROUP VICE PRESIDENT INDUSTRY STRATEGY AT LLAMASOFT.

HOWEVER, expectations of AI might be outgrowing its reality. A recent Gartner report[2] has suggested "through 2020, 80 per cent of AI projects will remain alchemy, run by wizards whose talents will not scale in the organisation." In our experience, the most successful organisations have taken a very pragmatic approach to AI with use cases that solve nagging problems which could not be solved with traditional approaches. This pragmatic approach accepts that AI is a collection of many different types of advanced algorithms, and not focused on Deep Learning alone which is generating a lot of attention.

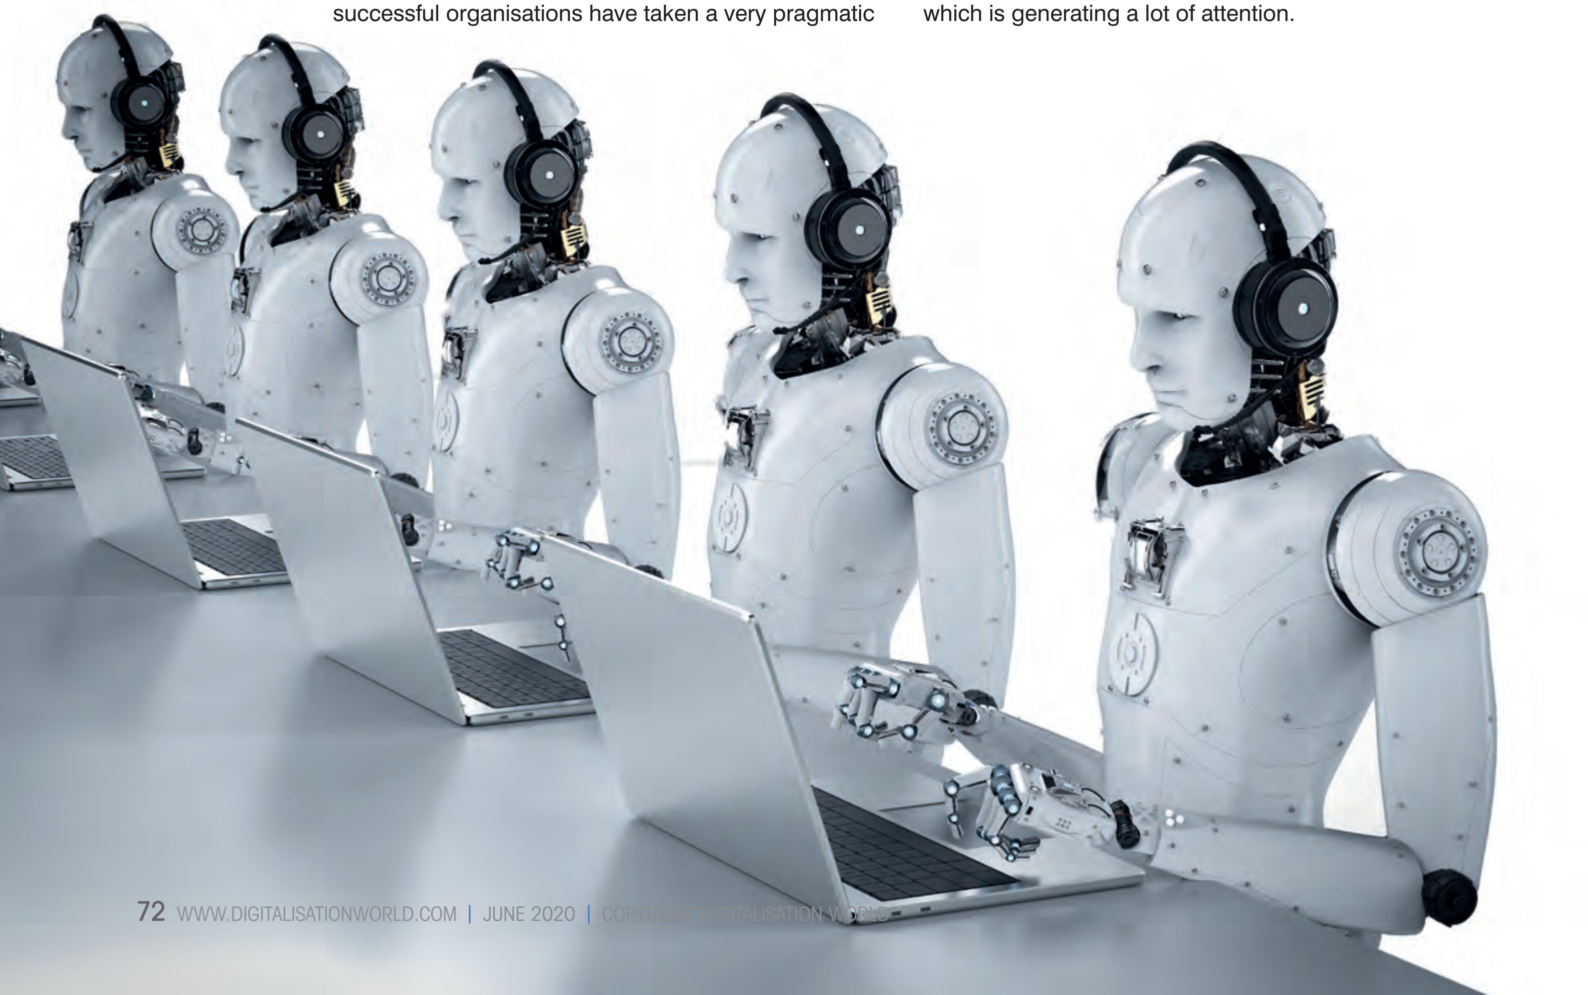

Let's look at some usecases where AI application moves beyond one-off insights and can be operationalised – packaged into an app and used enterprise wide, getting embedded in the daily decisioning process.

Preparing for unpredictable order volumes

First, let's look at how companies manage unpredictable order volumes. A major consumer packaged goods (CPG) customer of ours was facing the challenge of very spiky, unpredictable ordering behavior by a major ecommerce retailer. This was constantly resulting in elevated chargeback penalties due to stockouts, or excess stock, increasing inventory risk. The traditional statistical forecasting approaches it was using were failing to deliver the needed predictability to manage this risk.

However, the retailer in this context provides a wealth of information beyond traditional time series data, such as the product searches, customer page views and in-stock availability. While much of this data is not typically housed in an ERP system, which is the foundation for traditional planning systems, using AI and blending the data from traditional and non traditional sources, an app was built and deployed to the planners for the ecommerce channel.

This provided them with visibility into the next order quantity with probability factored in. The app also elevated the "cut" risk (order not filled in full), while making prescriptive recommendations such as stock transfer orders and production changes to ensure product availability.

Performing chargeback analysis and determining rootcauses

Our second CPG customer example looks at the struggle to meet ever shrinking retail customer delivery windows which causes increasing chargeback penalties from strategic customers. Upon close evaluation, much of the data the company had access to was not being effectively used. This included data on when the order was placed, manufacturing status, inventory status, when the order was changed, when it was picked and delivery information.

With such an abundance of data, this turned out to be a great case for AI. Instead of manually piecing together information to understand a single failure or group of failures and ascertain the root cause, AI helps identify patterns in the data. This is the modern take on the classic 5-Why approach. AI goes through the first 3 or 4 Why's automatically. It pinpoints the full sequence of events that led to a failure.

Once it finds these patterns, it is easier to fix the root cause and predict future failures. But, just running this analysis once is not enough. Hence an AI powered app that is refreshed with the latest data was deployed to make the process repeatable.

Embarking on your AI journey

While the power of AI can be transformative, it is important for organisations to filter the hype from reality, and keep in perspective that AI is the means to the end rather than the end in itself. Given the rapidly evolving nature of AI, senior leaders should consider engaging external partners with a proven track record in kickstarting initiatives.

Companies should start with a burning problem that can be solved by applying data intelligently. Trying to perfect the quality of data and AI models in the first go can be the enemy of speed to value. Instead, an agile approach of experimenting, learning, and adapting on the go, starting with the data already available (with all its limitations) is a practical approach. Quantify and communicate the benefits throughout the organisation so excitement can be built around AI for more use cases, and to stimulate additional ideas. By taking a pragmatic approach to deploying sustainable AI, organisations can unlock tremendous amounts of value that is simply not possible with traditional approaches.

Further reading

1 Accenture Research 2019: AI: BUILT TO SCALE: From experimental to exponential

2 "Predicts 2019: Analytics and BI strategy" by Gareth Herschel et al, Dec 2018, Gartner ID: G00372971

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

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

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

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