INDEPENDENT PUBLICATION BY Raconteur RACONTEUR.NET | #0928 | 11/12/2024

CIO 2025

02 IS THE CIO ROLE TOO IMPORTANT TO LOSE?
06 HOW TO MEASURE THE ROI OF AI DEPLOYMENT
15 A TWO-STEP SOLUTION TO AI IMPLEMENTATION

CIO 2025

Distributed in
THE TIMES

Published in association with
TECH SHOW LONDON

Contributors

Jon Axworthy
A journalist, specialising in health, technology, science and the future, with work published in *T3*, Wareable and The Ambient.

Tamlin Magee
Senior technology writer at Raconteur. He's interested in the impact of new technologies on people and society.

MaryLou Costa
A business writer specialising in the future of work, sustainability and technology, with work featured in *The Guardian* and *The Observer*.

Megan Tatum
An award-winning freelance journalist based in Malaysia. She covers business, tech and health for a range of business and consumer publications.

Raconteur

Special projects editor
Ian Deering

Sub-editor
Gerrard Cowan

Commercial content editors
Laura Bithell
Larnie Hur

Commercial content executive
Jessica Lynn

Commercial production managers
Alex Datcu
Ellen Newsome

Production executive
Sabrina Severino

Design and illustration
Kellie Jerrard
James Lampard
Samuele Motta

Design director
Tim Whitlock

Certified B Corporation

Although this publication is funded through advertising and sponsorship, all editorial is without bias and sponsored features are clearly labelled. For an upcoming schedule, partnership inquiries or feedback, please call +44 (0)20 3877 3800 or email info@raconteur.net.
Raconteur is a leading business media organisation and the 2022 PPA Business Media Brand of the Year. Our articles cover a wide range of topics, including technology, leadership, sustainability, workplace, marketing, supply chain and finance. Raconteur special reports are published exclusively in *The Times* and *The Sunday Times* as well as online at raconteur.net, where you can also find our wider journalism and sign up for our newsletters.
The information contained in this publication has been obtained from sources the Proprietors believe to be correct. However, no legal liability can be accepted for any errors. No part of this publication may be reproduced without the prior consent of the Publisher. © Raconteur Media

@raconteur raconteur-media @raconteur.stories

raconteur.net /cio-2025-dec-24

LEADERSHIP

CIO roles face the axe, but at what cost?

Facing pressure for greater efficiency, some firms have decided to remove their tech leaders. But as businesses digitalise their operations, are the risks too high?

Megan Tatum

It's been a challenging few weeks for Asda. In November, the supermarket confirmed it would make sweeping staff cuts at its Leeds and Leicester head offices, with nearly 500 employees facing redundancy.

When it emerged the retailer's chief information security officer (CISO) and head of security operations were among that number, concerned staff reportedly quizzed Asda execs on whether disbanding the senior tech team would leave the company vulnerable to a customer data breach.

Asda insists it will not. In a statement, a spokesman said: "We have a dedicated function that works hard to ensure that our internal systems and the data we hold remain secure in the face of cybersecurity challenges faced by all businesses."

But the worried response from Asda's workforce begs the question: in an age of digital pre-eminence, are security and information chiefs the riskiest roles to lose in a restructure? Without a dedicated leader at the helm, could organisations' data and digital assets be exposed to new threats? And if such a move is unavoidable, how can firms minimise the danger?

There's no doubt that the responsibilities typically assigned to a CIO and CISO – information security, technology and IT deployment – are seen as business-critical by C-suite leaders. According to new research by Accenture, more than 40% of C-suite job postings in the UK in the past year have been data-related. One in four FTSE 100 board-level executives now say they're proficient in technology, up 12% over the past three years.

"With almost every modern company using the cloud in some form and stakeholders radically changing how they consume services, the need for high-level IT is essential for most business operations," says Andrew Smith, CISO at Kyocera Document Solutions UK, a global manufacturer of high-tech ceramics, electronic components, solar cells and office equipment.

The CIO and CISO are indispensable when they're the sole strategic leads for these business functions. "Depending on how they are led, controlled and implemented, complex digitisation and IT projects can make or break a business," he says, making it dangerous to dispose of the CIO role.

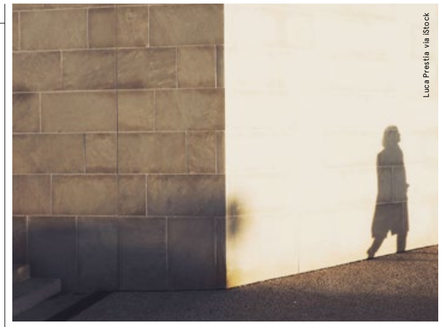
Luca Prestia via iStock

David Morimanno is director of identity and access management technologies at IT consultancy Xalient. He outlines the potential consequences of ditching roles such as CIO or CISO.

"Without that leadership, projects can stall, operational efficiencies may suffer and critical systems can become vulnerable to cyber threats," he warns. "Removing this role without ensuring a capable replacement leaves the organisation exposed to breaches, data theft and regulatory non-compliance, issues that can carry both financial and reputational consequences."

Any gaps in oversight undoubtedly put organisations at risk. But this doesn't necessarily mean the CIO role is exempt from the corporate chopping block, according to Sachin Shah, management consultant at Bain & Company.

Although IT is more important than ever before, Shah believes this has diminished the scope of the CIO, rather than elevating it. "The technology operating model is changing," he says. "In the past, you had one person who would manage everything from applications to infrastructure to networks and telecoms to end-user computing."

However, responsibility for tech, IT and information security is now distributed and segmented more widely across the workforce, rather than being concentrated in a single C-suite role.

Some companies are distributing these functions across multiple senior roles, rather than one, with the appointment of a chief data officer, chief digital officer and so forth, as well as a CIO. Others have introduced 'business-product owners' – a leader with some technology literacy who oversees the deployment of a specific tech-intensive product, with the team responding directly to the product owner rather than a CIO, Shah explains.

He adds that some organisations are also reviewing capability sourcing strategies to outsource more areas of IT infrastructure to third-party providers. This is particularly prevalent in managed security services.

This segmentation and delegation of a CIO's areas of responsibility has arguably made it easier to part ways with them, with less risk. However, there are still some fundamental steps any business must take before handing their CIO a redundancy notice, notes Smith.

The senior leadership team must first ask themselves some fundamental security questions. For example, how will the business risk profile be controlled and managed without a dedicated C-suite member to focus on it?

Companies must also be careful about relying too heavily on outsourcing companies without proper scrutiny, Smith adds. "Do you fully trust the outsourcing company? What are their credentials? Where will your data be stored? Any outsourced IT company adds an extra layer of risk for potential cybercriminals, as an additional stakeholder now has access to your data."

You almost always need someone who understands the business and its needs back to front, he continues. "While third-party providers often claim to offer a seamless service, you need someone on the inside who works with the business's systems every day."

But what if the decision is unavoidable? In such cases, Morimanno says "robust planning, clear communication and a strong commitment to maintaining digital leadership are essential to navigating the transition successfully."

Companies should first reallocate leadership responsibilities, ensuring any replacement "has both the technical expertise and the strategic vision", he says. He also recommends that organisations "establish a digital advisory board or cross-functional leadership team to maintain oversight of critical initiatives" and conduct a comprehensive audit of the firm's digital assets, infrastructure and existing cybersecurity measures. "Identify any gaps or vulnerabilities that might arise from the leadership change," Morimanno advises.

Where possible, retain members of the CIO's team, expanding their roles to maintain continuity if it is appropriate. And finally, "be transparent with employees, stakeholders and customers about the restructuring process. Explain how the company plans to effect its digital strategy and safeguard its infrastructure despite the change in leadership," Morimanno concludes.

Business leaders who skip any of these steps are likely to find themselves facing tough questions – as Asda has learnt.

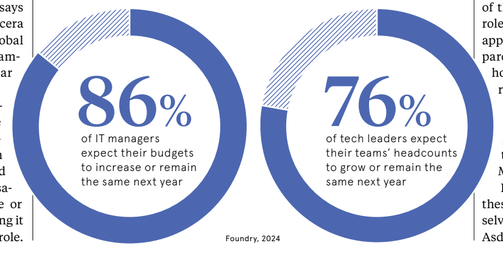

Foundry, 2024

INSIGHT

'Addressing issues in emerging tech requires more than internal discussions'

Simon Press, senior portfolio director, tech shows at CloserStill Media, discusses how knowledge-sharing events can help leaders to cope with the pace of change in business tech

Consistent innovation has become essential for modern businesses. As companies embrace digitalisation, they face increasingly complex challenges. They must innovate rapidly, remain agile and integrate sustainability into their core values.

As a result, organisations are changing their perspectives on tech adoption. They are focusing more on how new technologies can generate sustained value.

That doesn't mean firms have slowed their investments in advanced technologies such as AI. But the power of AI lies in its ability to address real-world problems, not in its novelty. According to a study by Deloitte, organisations are seeking improved efficiency, increased productivity and cost reduction from their AI investments. Two in five (42%) say they have actually realised these benefits.

Demands for infrastructure are also growing alongside compute and storage requirements. As leaders increasingly focus on improving efficiency and scalability, they are recognising that the physical infrastructure supporting AI adoption is just as important as the algorithms that guide the technology.

The importance of resilience in business has never been more evident. Recent events such as global supply chain disruptions, the transition to hybrid working and rising geopolitical tensions have tested organisations' ability to adapt and recover. Firms are adopting multi-cloud strategies to improve flexibility and responsiveness.

However, resilience is not solely a technical challenge, it is also a cultural one. Building resilient systems goes beyond enhancing cybersecurity, it requires a culture of adaptability and collaboration.

Partnerships among service providers and hyperscalers highlight the trend towards grater cooperation and the value of shared knowledge to fortify resilience at scale.

Moreover, sustainability can no longer be ignored by organisations. It is now a central business priority that technology leaders are seeking to incorporate into their strategies to decrease environmental impact without damaging performance.

As the demand for AI solutions grows, business leaders are beginning to explore new avenues to power AI operations while minimising their carbon footprints.

Businesses now understand that sustainability goes beyond compliance – it can also enhance competitiveness through cost savings and brand trust. But achieving meaningful progress requires collaboration to integrate sustainability into operational strategies.

The increased use of data and AI also raises questions of safety, accountability and trust. Tech leaders are now responsible for handling these complexities and promoting a culture of transparency.

This widening remit is shifting the role of leadership. It is not enough for leaders to employ solutions, they must also be the caretakers of responsibility, establishing practices that support the business and society.

Addressing these issues requires more than internal discussions. Events such as Tech Show London, as well as its new Cloud and AI Infrastructure show, provide platforms for people to deepen connections in the industry.

A study by Kearney reveals that 45% of businesses point to a lack of technical skills as a key barrier to the adoption of GenAI in their organisations. Attending these kinds of events or alternative knowledge-sharing spaces also offers a chance to upskill teams, enabling employees to stay ahead of technological shifts.

By attending events where industry leaders share their experiences and employees engage with upskilling opportunities, organisations can better position themselves to meet the challenges of the future with confidence.

Simon Press
Senior portfolio director, tech shows
CloserStill Media

Leaving legacy tech behind can unburden your business

While legacy-tech migration seems daunting, the real risk lies in inaction. The key is embracing change as a continuous, incremental process to unlock value across the business

Tech migration projects can be costly, daunting and fraught with risk – and when things go wrong it's generally the IT team that gets the blame. So it's hardly surprising that CIOs might look at their company's ageing but functional tech infrastructure and ask whether migration is really worth it.

The answer, however, is a resounding yes. While migration and modernisation projects are often challenging, inaction is the bigger threat to long-term business success.

Legacy technology can burden companies with spiralling costs from third-party vendors and hobble their ability to move quickly in a competitive market. Other issues include technical debt and skills challenges, as people who understand how to work with older systems retire or leave the company.

"Every day, each application becomes older and harder to maintain. The more effort it takes to migrate, the greater your risk ," says Thilo Rockmann, CEO of LzLabs, which transforms existing IT by placing it in a modern computing environment.

Many IT leaders are aware that legacy technologies are holding back businesses. Indeed, a recent survey by ISG and LzLabs found that 95% of tech leaders are concerned about the implications of not modernising mainframe applications and data.

But the complexity of incumbent systems, or lost source-code issues, often means migration plans are repeatedly kicked into the long grass.

Respondents to the survey also highlighted cultural resistance and regulatory compliance as significant barriers to migration. These issues can lead to a state of paralysis. CIOs and other C-suite leaders may acknowledge the need to update legacy systems, but the necessary support and resources for the journey – or even agreement on the direction of travel – never materialise.

"There are multiple barriers and varying interests that ultimately paralyse the entire organisation," says Rockmann. "Everyone wants to have a voice, get involved and exert influence, but this desire for involvement may not lead to swift action on what the business actually needs."

The cycle of change

Overcoming this paralysis is the first step towards achieving greater agility, winning and retaining more customers, embracing modern applications and exploiting the full power of data. And, it may not be as challenging as it seems.

One assumption that can prevent the migration of systems and applications to new environments is that it must be achieved in a single huge leap. Instead, migration is a continual process of change and evolution.

"The key issue in our industry is that moving beyond legacy technology is not a one-time project," says Rockmann. "It's a continuous cycle, as what is new today becomes outdated tomorrow."

Viewed this way, legacy technology is not bad per se. It may simply be that the organisation's fundamental beliefs or business strategies have evolved, and systems and applications that were once cutting-edge are now holding it back.

"A new system isn't necessarily better than the old one; it simply meets the needs of the current environment more effectively at this time," explains Rockmann. It is important therefore to plan for future changes to the organisation's tech needs.

But trying to anticipate all of these changes isn't the best way to approach migration. "If you plan too far in advance, your goals may appear as insurmountable obstacles, potentially stifling innovation," Rockmann explains.

It's also important to understand that a migration project should unlock the value embedded in the company's existing application portfolios – not discard the foundations.

Rockmann points to a vehicle manufacturer that LzLabs is currently working with. "They have an application that manages all the logistics for bringing parts to the assembly line, which they've built themselves. No matter what they plan to build tomorrow – motorcycles, tractors and so on – or whether they want to implement AI, that application will still have value."

An incremental approach

LzLab's approach to technology migration and modernisation focuses on four core principles: preserving what needs to be preserved, changing only what needs changing, maintaining interoperability and using open-source technologies.

These principles underpin its Software Defined Mainframe® (SDM), which uses binary rehosting – a means of migrating legacy mainframe applications to the cloud without the need to rewrite or recompile the application.

> Moving beyond legacy technology is not a one-time project. It's a continuous cycle, as what is new today becomes outdated tomorrow

It's a low-risk, low-cost way of moving valuable applications and data so they once again serve the business's needs, enhancing rather than hindering innovation. Using binary-compatible interfaces also enables an iterative and incremental approach to migration.

For instance, a global automotive manufacturer recently transitioned its business-critical processes and applications step by step from its legacy mainframe to the SDM.

This was necessary owing to mainframe overload caused by resource-hungry applications, sluggish application responses to sales and customer requests and a shrinking skills base. But the goal was always to offload and functionally complement the mainframe rather than completely replace it.

"You've got to make sure that you preserve what needs to be preserved and only change what needs to be changed, focusing on a difference to the business," says Rockmann.

"Nobody would say, 'All Londoners need to move out of London so we can completely redo the whole tube system, and once you return, it's all going to be new and shiny'," he says. "There's constant change, constant construction. IT is not so different in this sense."

Some of the key results from the automotive manufacturer's migration project include enhanced service levels for customers, increased scalability of capacities and a significant reduction in mainframe-operating costs.

Thanks to careful planning and proper execution, the company's business-critical applications and data are now fit for the future rather than stuck in the past.

For more information on overcoming the legacy tech burden, visit lzlabs.com

LzLabs®

INTERVIEW

'Trust is somewhat transitive. But a lot of the time you have to put the work in everywhere'

As AWS shifts its OpenSearch search engine to the Linux Foundation, **David Nalley**, the tech giant's director of open source, reflects on the possible implications

Tamlin Magee

Big tech now funds and maintains much of the open-source software community, but conflicts still occur when one company is accused of amassing too much power. The Open Source Summit in Vienna in September pointed to a possible way forward: handing control to non-profit foundations.

That's according to David Nalley, director of open source at Amazon Web Services (AWS). At the summit, AWS handed its OpenSearch 'fork' of Elasticsearch – the back-end analytics engine – to the Linux Foundation, a non-profit that promotes and governs open-source projects.

The move is significant in the highly intricate world of open-source software partnerships. In software, a fork occurs when a project diverges from its original codebase, which in turn changes its governance. The OpenSearch transfer means that a project which had been largely controlled by AWS – even though it is open source – is now vendor-neutral.

Elasticsearch disagreements bubbled to the surface in 2021 when the original founders of the codebase, Elastic, shifted the software's license from Apache 2.0 to something called Server Side Public License. This meant that the project was no longer truly open source. The decision was motivated by Elastic's mounting irritation over a perceived capture and subsequent monetisation of Elasticsearch by AWS for its Amazon Elasticsearch Service – supposedly without giving much back to the codebase or maintenance.

Whatever the motivations for the move, AWS felt strongly enough to take action. "Having that codebase be open source was important to us and to our customers so we took the extraordinary step of creating the [OpenSearch] fork," says Nalley.

The fork was a success, quickly shooting into the top 50 database engines. "We've been acting, in the intervening three years, as the steward for the project," Nalley says.

But many have reservations about any single firm having such strong ties to particular projects. Nalley heard from AWS customers that the presence of one vendor so close to OpenSearch had impacted the health of the project, hence the decision to transfer the project to the Linux Foundation.

"Customers and partners wanted vendor-neutral, independent governance," Nalley explains. "We're hearing from a lot of customers who perceive extra risk when a single vendor dominates or controls an open-source project. Several of our customers told us that they have patches for OpenSearch, but their company has a policy against contributing that code unless it's at a vendor-neutral place."

Open-source software forks are increasingly receiving backing from hyperscalers and major tech companies, driven by licensing changes.

Amazon, Oracle and Microsoft have all backed Valkey, the open-source alternative to Redis, a large data store. Valkey is also hosted by the Linux Foundation and is already outpacing the original codebase, by some accounts.

Some detractors suggest these open-source forks are a consolidation of big tech's power on the open-source ecosystem or are simply economically driven decisions to avoid paying licences.

But Nalley says the AWS/Elasticsearch or Valkey model – where a fork is hosted by a foundation – can help organisations to reduce their risk profiles when consuming open-source projects, especially as it relates to single-vendor control. "This is going to increasingly become a factor for companies who are consuming open-source software," says Nalley, noting that AWS takes this into account before using an open-source project.

> Folks say there's decreased trust when a single vendor can arbitrarily make decisions about an open-source project, whether that's the technical direction or the lisencing or whether to continue a project at all

"Folks have been saying there's decreased trust when a single vendor can arbitrarily make decisions about the future of an open-source project, whether that's the technical direction or the licensing or whether to continue working in the project at all," he adds. "That will drive a lot of attention to open-source foundations. Whether it will mean a lot more software moves to foundations, I don't know."

One thing's for sure: tech giants will remain involved in open-source software at all stages, whether they hand control to a foundation or not. The Open Source Contributor Index ranks commercial entities by their total contributions to open source projects. A quick peek reveals a lot of activity from huge players such as AWS, Google, Microsoft, Intel, Huawei, IBM and Nvidia.

This is nothing new. Many of these organisations have a long history of involvement with open source. But the broader open-source community maintains a healthy scepticism towards corporations on the periphery, despite their significant contributions. So how can businesses win their trust amid demands for vendor-neutral governance?

"A lot of it comes down to proving over time that you're making the investments necessary to help sustain open source," says Nalley. There are no shortcuts.

"The Cloud Native Computing Foundation talks a lot about 'chopping wood and carrying water'," Nalley says, referring to the Zen Buddhist proverb about everyday tasks remaining the same whether you have found enlightenment or not. In this case, the wood and water are writing code and fixing bugs. He points to projects such as PostgreSQL, an open-source relational database, where AWS is the top reviewer of code.

If businesses wish to really demonstrate their commitment to open source, Nalley suggests, they should consider contributing to ecosystems where they don't currently have any products. For AWS, one such project is developing the emergent programming language Rust. The company even has a full team working solely on the project.

"We're doing that because, just like everyone else, we need a more performant, more stable Rust programming language, and a set of tools like the compiler and standard library that are easily consumable and will work for folks," he says.

Showing commitment through contributions and maintenance, especially with no obvious dog in the fight, can be valuable for companies seeking to win and maintain trust.

But Nalley warns that trust is "somewhat transitive, meaning you can earn it in one place and maybe you have enough reputation that it carries over in other places. But a lot of the time you've got to put in the work everywhere."

INVESTMENT

CIOs embrace AI, but many struggle to measure its impact

AI holds huge promise, but its impact can be difficult to assess. We asked CIOs how they calculate the technology's return on investment

MaryLou Costa

The transformative potential of AI has been well-documented, with applications in everything from customer service to data analysis. But many CIOs are grappling with a surprisingly complex question: is it worth the money?

Nearly nine in 10 (87%) organisations are actively developing GenAI initiatives, but only 35% have a clearly defined vision for how they will create business value from GenAI, according to Bain Research.

And there are different views on what constitutes success. Consensus on how to measure the return on AI investments is rare, according to a survey of nearly 600 CIOs and heads of IT by Gong, a sales platform.

So where are major firms focusing their AI investments – and how do they measure the impact? The survey found that 55% focus on productivity, but a similar share look at efficiency and revenue (53% each), and 46% focus on employee satisfaction.

AI has transformed insurer Axa's business, bringing benefits everywhere from customer service to risk assessment and fraud detection. By introducing a corporate version of GPT to its call centres, call resolution time has been slashed from five minutes to five seconds, as agents can swiftly retrieve policy document references to customer questions, says Axa's UK and Ireland CIO, Natasha Davydova.

AI-enabled pricing platforms have made pricing more efficient, helping the company's underwriters complete their customer risk assessments and pricing proposals in hours rather than weeks. AI-enabled IT observability tools, meanwhile, have been implemented to detect and prevent IT incidents, reducing the total number of incidents and pushing down the time to fix.

A trial of Microsoft Copilot, which helps teams summarise and draft documents, is being extended, with its value based on productivity increases, error reduction and employee satisfaction, all of which have changed for the better, according to Davydova.

This comprehensive tech stack doesn't come cheap. Davydova declined to provide details of Axa's spending, saying the company's AI budget is confidential. However, ChatGPT at enterprise level is quoted at $30 (around £24) per user per month. The price for Microsoft's Copilot Pro, meanwhile, is currently published at £19 per user per month. In a 150,000 strong global business such as Axa, this would amount to £43.2m and £34.2m respectively.

Still, that's not a huge outlay for a giant like Axa. The company's total tech spend in 2023 of all platforms, not just AI, was reported by Global Data to be $2.2bn (£1.74bn).

> **Most CIOs are trying to simplify their operations, but the issue with AI is, what do you remove off the back of it?**

And the changes are already making a difference to Axa's bottom line, Davydova says.

Indeed, the company has already recorded profit growth of 5% this year, according to industry media reports, with life and health insurance premiums up 7%, as Axa's half year 2024 results confirm.

So how does it measure the value of the technology? Davydova believes the best use of AI is in customer-facing areas, which "help enhance the growth of our revenues and profitability, because customers choose to stay with a high-quality supplier of insurance services".

She adds: "By analysing KPIs from operational efficiency and enhanced customer experience to improved risk management, cost savings and employee productivity, the evidence indicates that AI contributes positively to the company's goals and bottom line."

Jean-Philippe Avelange is CIO at Expereo, a business connectivity firm. He says all conversations around AI initiatives begin with a question: "What's our starting point that we want AI to help with?"

Avelange says Expereo decided in 2023 to upgrade its Salesforce platform to the AI-enhanced Agentforce offering, which includes features like real-time AI-powered guidance in customer interactions. The company would not share financials, but a total package at the published price of $500 (£394) per user per month, would amount to $2.4m (£1.9m) for the 400 Expereo employees using the platform.

Like Davydova at Axa, Avelange focuses AI deployment on customer service, with ROI metrics concentrated here. He also correlates AI-platform rollouts to productivity gains and any resulting increases in employee satisfaction.

Avelange outlines key focus areas for Expereo. "How many emails are we sending per customer-service agent? How long does it take an agent to handle a case summary? How much time is spent on a customer update? We then assess the cost for that specific AI use case, start prototyping and commence frequent rollouts to gather quick feedback," says Avelange.

Implementing new technology comes with a financial price, but there's also an environmental cost that CIOs shouldn't overlook, notes Louise Bunting, CIO at Carbon Net Neutral Technology Solutions, a corporate carbon-measurement and management company.

For example, it took 1,287MW/h of electricity to train the large language model (LLM) GPT-3, according to the Association of Data Scientists; that's roughly equivalent to the usage of an average American household over 120 years. Moreover, Gartner has predicted that by 2030, AI could consume 3.5% of the world's electricity, while each GPT query requires roughly half a litre of water to cool its servers.

This all adds to an organisation's carbon footprint, Bunting warns, which CIOs must consider when assessing the ROI of AI. "Most CIOs are trying to simplify their operations, but the issue with AI is, what do you remove off the back of it? You're adding tech, but not taking anything away. If you've got a carbon target, you're adding something that is probably the most power-hungry system, consuming up to four times more than a standard technology stack. That is a big problem from an environmental perspective," says Bunting.

Bunting recommends a particular line of questioning when considering whether AI will be worth the cost. "Is it actually adding value? Or could you do what you need to with tech that you've already got? Is it actually going to save us any money, when we could do the same thing through the automation and digitisation of processes, without AI?"

> **If the value of AI is that I can respond to my customers faster, does that warrant having a whole team governing its use?**

Avelange says that upgrading Expereo's existing Salesforce tool to Agentforce has helped to minimise the costs that come with bringing in new AI products, such as network usage spikes, plus the need for extra bandwidth and security layers.

"It's not a simple boxed product that you buy and you're ready to go," he says. "Working directly in our existing Salesforce platform alleviated this risk and ensured that we could keep costs contained while maintaining ROI."

Yet another overlooked cost of AI, Bunting adds, is governance. That includes the need for specialist staff members to create frameworks and processes for how AI should be used and monitored in a business, as well as additional legal support to clean up the mess if AI gets it wrong, which is still a reality. An AI governance director can earn a salary of up to £74,000, according to Glassdoor.

"If the value of AI is that I can respond to my customers in five seconds, does that really warrant having a whole team governing its use?" asks Bunting.

As Avelange notes, quick wins aren't everything. "The risk of any organisation making short-term considerations about ROI on AI initiatives is that they could potentially miss out on any long-term gains and benefits," he says.

He's convinced that AI "will profoundly transform the way any company operates" and that it "is not a debate of knowing whether AI is worth it or not". However, he's realistic about its complexities.

"It is a matter of survival for enterprises to adopt AI, while remaining very conscious of the costs and hype around it," says Avelange.

MEASURING UP AI

Share of tech leaders who focus on the following when assessing the effectiveness of AI investments

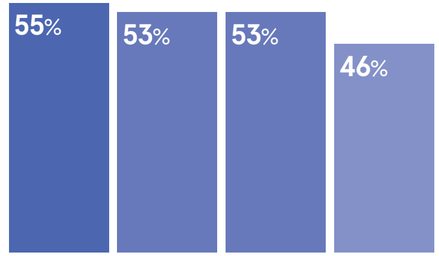

Gong, 2024

THE KEY TO DATA INTEGRITY

Business leaders are facing mounting pressure to implement emerging technologies such as AI. But the effectiveness of digital tools depends heavily on the data that fuels them. Factors including inconsistent data formats, as well as the sheer volume of raw data, are hampering firms' efforts to gain insights from the data they hold. A robust data-governance strategy can help to ensure data integrity and establish trust in an organisation's data.

71%

of IT professionals say their organisations have a data-governance programme

DATA EXPECTATIONS VERSUS REALITY

IT professionals worldwide

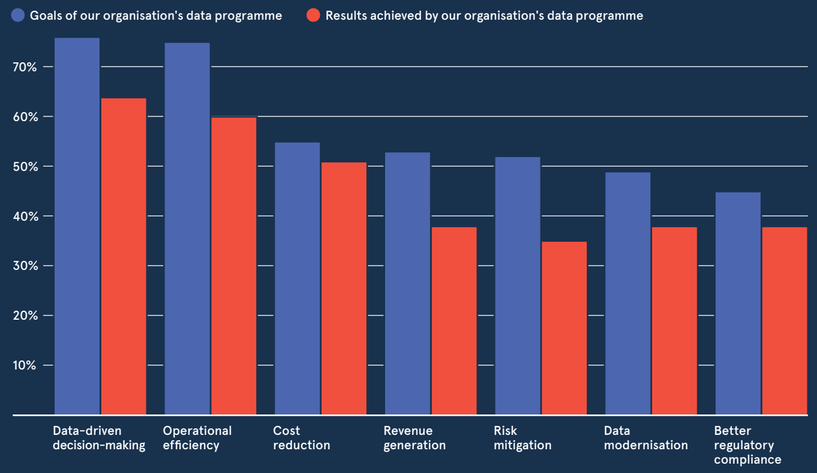

TWO IN THREE IT PROFESSIONALS DON'T FULLY TRUST THEIR FIRMS' DATA

IT professionals worldwide

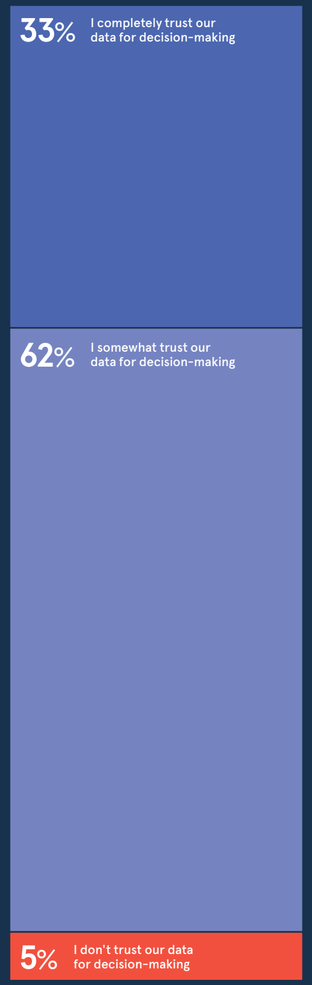

DATA QUALITY IS ALSO A HURDLE TO DATA INTEGRATION

IT professionals ranking the following as challenges to their organisations' data-integration projects

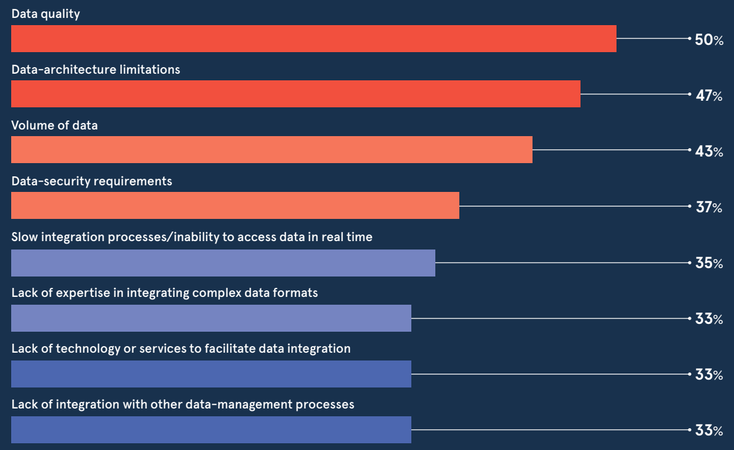

FIRMS CAN IMPROVE DATA QUALITY WITH BETTER DATA GOVERNANCE PROCESSES

IT professionals who say their organisations' data-governance programme has added value in particular ways

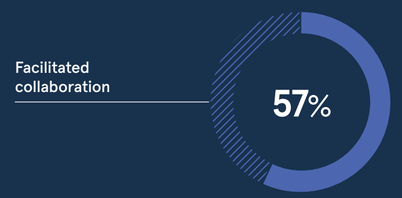

DATA QUALITY IS ONE OF THE MAIN CHALLENGES TO BETTER DATA INTEGRITY

IT professionals ranking the following among their organisations' top-three data-integrity challenges

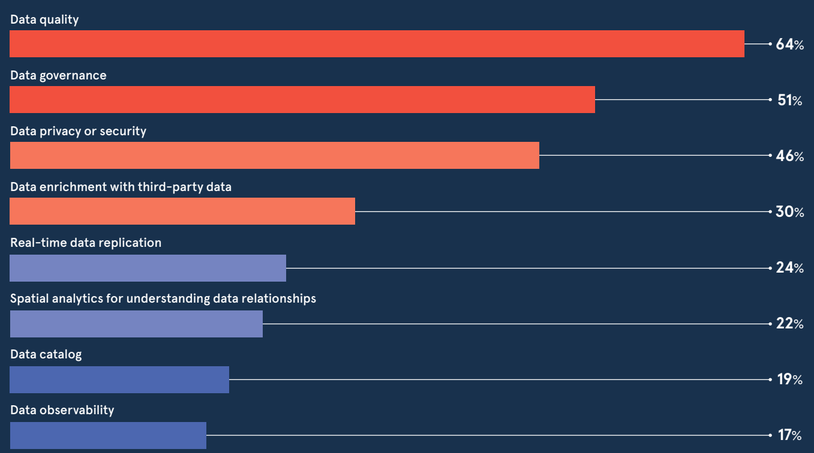

WHAT'S STANDING IN THE WAY OF HIGHER-QUALITY DATA?

IT professionals ranking the following as challenges for achieving higher-quality data

Precisely, Drexel University, 2024

Cool running: reimagining the sustainable data centre

As data centres strain under AI's growing power demands, innovative cooling solutions and sustainable architectures point the way towards a greener computing future

Data centres are the backbone of modern computing, housing the world's servers and data processing, but this comes at a cost. Data centres consume huge amounts of electricity. They account for roughly 3% of global energy use, and that figure is expected to rise to 8% by 2030.

Globally, this amounts to an estimated 200 terawatt hours (TWh) annually – more than the total energy consumption of some entire countries. Much of this is due to the demands of AI.

A single ChatGPT query, for example, requires 2.9 watt-hours of electricity, compared with 0.3 watt-hours for an average Google search. Goldman Sachs estimates that AI will soon add another 200TWh to global data centre power consumption, doubling the current energy demand.

Reducing the energy consumption of data centres is therefore vital for achieving national net-zero targets and reducing business costs.

Supermicro is a market leader in designing and delivering the components for data centres, offering servers designed to use less power while exceeding standard performance levels.

In the relentless pursuit of data centre efficiency, every decision – from server selection to cooling infrastructure – must balance peak performance with environmental stewardship, driving both operational excellence and sustainability goals.

Two of the larger data-centre energy demands come from computers' central processing units (CPU), which execute instructions from computer programs and process data.

While data centres rely on central processing units (CPUs) for general computing, graphics processing units (GPUs) dominate modern AI power demands. Each GPU can consume over 1000 watts – more than double a CPU's maximum draw of 500 watts.

The impact is multiplied by the architecture of AI-optimised systems, which typically contain eight GPUs for every two CPUs. This means a single AI server can have a GPU power footprint that's more than eight times larger than its CPU requirements. When multiplied across thousands of servers in large AI data centres, GPU power consumption becomes the dominant factor in both energy use and cooling demands.

Consider the evolution from traditional data centres to today's AI computing facilities – while CPU requirements have certainly grown, it's the massive deployment of power-hungry GPUs that are driving the unprecedented surge in energy consumption and thermal management challenges.

Graphics processing units (GPU) are the computing hardware designed to render high-quality images and videos efficiently. Originally used for 3D games, they have evolved to become the engine of AI due to their ability to perform many parallel operations.

Therefore, it's not hard to see how the requirements for both CPUs and GPUs have skyrocketed in recent years.

The main energy demand, however, doesn't simply come from plugging more and more servers into the mains. It comes from keeping them cool.

As computational demands soar, traditional air cooling methods – relying on fans, heatsinks and HVAC systems – are reaching their limits in data centres. The challenge was made clear when even tech giants Google and Oracle experienced server downtime during Europe's 2023 heatwave, highlighting the growing thermal management crisis facing the industry.

"Air cooling, while effective in the past, is reaching its physical limits", explains Upadhyayula. "There's only so much air can do to dissipate the heat generated by modern hardware. To continue using air cooling with today's high-power components, you would need larger fans and increased system space to effectively circulate and expel hot air," he explains.

While increasing the dimensions of systems could help, it conflicts with data-centre goals of optimising rack space and density, he says. "These opposing factors make liquid cooling more attractive since it can maintain or even reduce system size while achieving higher thermal efficiency."

Over the past decade, Supermicro has pioneered liquid cooling technology, revolutionising data-centre thermal management. Through an innovative system where coolant circulates to specialised cold plates and heat is efficiently extracted from critical components without direct contact.

This addresses the mounting thermal challenges posed by AI workloads. The technology has garnered enthusiasm from industry experts thanks to its ability to deliver superior cooling performance while maintaining compact system footprints – a crucial advantage in space-constrained environments.

The superior efficiency of liquid cooling systems is undeniable, but there are barriers to adoption, including resistance from data-centre operators wary of the unfamiliar technology and infrastructure requirements.

> There's only so much air cooling can do to dissipate the heat generated by modern hardware

Still, liquid cooling systems permit greater computing density and significantly reduce an organisation's carbon footprint.

For these reasons, Upadhyayula expects to see "more customers gravitating towards liquid cooling to maintain their current data-centre footprints while achieving higher system performance".

Achieving this requires a bespoke partnership approach rather than an off-the-shelf sales transaction.

"We invite customers to our engineering test facilities, where we work collaboratively to identify and resolve their challenges. By prioritising customer-specific problem-solving over rigid metrics, we aim to deliver solutions that truly address their operational needs."

Supermicro maintains complete control over its cooling solutions through comprehensive in-house design and assembly operations. This vertically integrated approach ensures exacting quality standards and enables rapid innovation through direct oversight of every stage from concept to completion.

"Our customer-centric approach involves understanding workloads and tailoring solutions to their specific needs", explains Upadhyayula. "This includes determining whether air cooling, liquid cooling, or another strategy is the most energy-efficient and effective option. Each system is designed with certified components, ensuring reliability and performance."

Reshaping sustainable computing, Supermicro's resource-saving architecture eliminates the need for complete system overhauls by allowing targeted component upgrades, thereby maintaining cutting-edge performance while dramatically reducing electronic waste and operational costs.

The solution can help CIOs gain the flexibility to modernise systems piece by piece, while CFOs benefit from reduced capital spending and lower operating costs, ultimately maximising return on infrastructure investments.

The data centres of the near future will prioritise enhancing efficiency, with liquid cooling and resource-saving architectures playing pivotal roles.

Innovations in cooling technologies can help to move AI from a drain to an enabler - AI-powered management tools can optimise cooling, workload distribution and energy usage in real-time.

Automation will further enhance operational efficiency, reducing human intervention and improving reliability, allowing for higher density in racks and minimising the space requirements for data centres while improving performance.

While the exact form of future data centres will depend on these advances, the overarching goal remains clear: to deliver higher performance with lower environmental impact, paving the way for a sustainable and greener digital future.

In the quest for perfect energy efficiency, the industry pursues the goal of a 1.0 power usage effectiveness (PUE) rating, representing zero energy waste.

PUE measures how efficiently a data centre uses energy by dividing total facility energy consumption by IT equipment energy consumption - the lower the rating, the better the efficiency.

While the current industry average stands at 1.35 PUE, with 1.09 once viewed as near-optimal, Upadhyayula says achieving higher energy efficiency remains a significant challenge.

"It's akin to railroad tracks appearing to converge in the distance. Although they never truly meet, they guide us toward a shared direction, optimising performance and efficiency simultaneously," he says.

Supermicro will keep moving the data centre industry in the right direction towards a greener, cleaner future.

For more information about how NVIDIA's fully accelerated computing platform has provided leaps in AI training and inference, visit nvidia.com

SUPERMICRO

NVIDIA

AS AI ADOPTION ACCELERATES, MANY ORGANISATIONS STILL LACK THE GREEN DIGITAL INFRASTRUCTURE NEEDED TO MAKE INNOVATION TRULY SUSTAINABLE

Percentage of leaders expressing AI-related environmental, social and governance (ESG) concerns

54% see energy consumption and increased green house emissions as a key risk

36% see anti-competitive practices as a key risk

26% see pollution and e-waste as a key risk

18% see increased water consumption as a key risk

Capital Group 2024

The data-centre scorecard: four pillars of sustainability success

In today's digital landscape, CIOs and infrastructure leaders face the dual challenge of delivering exceptional performance while meeting rigorous environmental standards.

Success in modern data centre management requires a balanced approach that optimises both operational efficiency and sustainability. Here's how to evaluate your strategy across four essential pillars.

Intelligent power management

The cornerstone of sustainable data centre operations lies in smart power utilisation. Advanced power management systems now enable real-time adaptation to workload fluctuations, ensuring energy is deployed precisely where and when needed. This dynamic approach prevents unnecessary component strain while significantly reducing energy waste.

Success in this pillar means demonstrating measurable reductions in power consumption without compromising performance. Look for systems that can provide detailed analytics on power usage effectiveness (PUE) and automatically adjust to varying demands. The ability to handle processing spikes efficiently while maintaining optimal performance during quieter periods is crucial for both sustainability and operational excellence.

> Success is about creating resilient, efficient and future-proof infrastructure

Efficient cooling systems

Cooling efficiency represents a critical metric in sustainable data centre operations. With the increasing density of computing resources, particularly in AI and high-performance workloads, traditional cooling methods may no longer suffice. Success in this area requires implementing advanced cooling technologies to match specific needs.

A cooling strategy should be evaluated on its ability to maintain optimal operating temperatures while minimising energy consumption. Liquid cooling solutions, for instance, can offer superior heat dissipation for high-density configurations, while optimised airflow designs might suffice for lower-demand applications. The key is selecting solutions that scale with the organisation's specific needs while maintaining efficiency.

Renewable-energy integration

The transition to renewable energy sources is a defining characteristic of future-ready data centres. Success in this pillar involves more than just purchasing renewable-energy credits – it requires a comprehensive strategy for integrating sustainable power sources into existing operations.

Measure success through the proportion of operations powered by renewables and the reduction in carbon emissions. Consider both centralised and distributed approaches: larger facilities might benefit from direct access to hydroelectric or solar power, while edge locations could leverage local renewable resources. This hybrid approach ensures sustainable power delivery across the entire infrastructure estate.

Advanced thermal design

The fourth pillar focuses on sophisticated thermal management through intentional design. Rather than treating cooling as an afterthought, successful data centres integrate thermal considerations from the ground up. This proactive approach encompasses everything from component placement to airflow optimisation.

Measure success through metrics such as thermal efficiency, component longevity and cooling system performance. Look for designs that minimise hot spots, optimise air or liquid cooling pathways and reduce the overall energy required for thermal management. The most effective solutions will demonstrate improved component life spans while maintaining or enhancing performance capabilities.

To effectively measure success across these pillars, establish clear metrics and regular monitoring procedures. Key performance indicators should include:

- Energy-efficiency ratios
- Carbon-footprint measurements
- Component performance and longevity statistics
- Cooling-system effectiveness
- Renewable-energy utilisation rates

Regularly assessing these metrics helps to identify areas for improvement and validates the effectiveness of sustainable initiatives. The most successful strategies will show continuous improvement across all four pillars while maintaining or enhancing operational performance.

By evaluating your data-centre strategy against these pillars, you can ensure that your infrastructure not only meets current sustainability requirements but is also prepared for future challenges.

Remember that success in sustainable data-centre operations isn't just about meeting environmental targets – it's about creating resilient, efficient and future-proof infrastructure that delivers both business and environmental value.

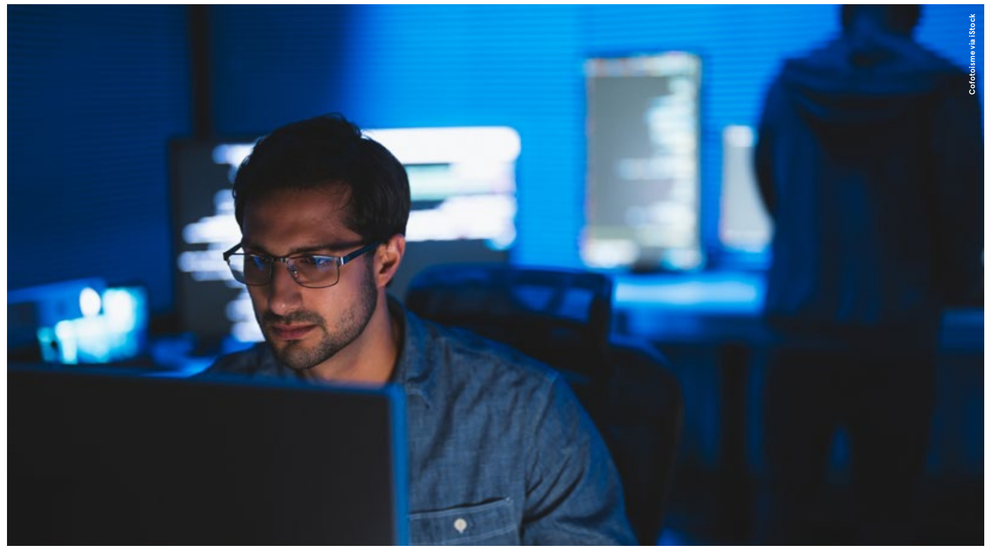
Cofotoisme/via iStock

CYBERSECURITY

Practice makes perfect: how to run a ransomware simulation

Ransomware attacks can be devastating. Regular simulations can help firms weather the storm, but careful preparation is essential to reap the rewards

Tamlin Magee

Ransomware is the gift that keeps on giving for cyber attackers, forcing businesses to pay a sizeable sum or risk losing access to their business-critical data. Ransomware simulations can help firms prepare for the worst – but they must be handled with care.

Most companies can expect to be targeted by a ransomware attack at some point. Six in 10 organisations (59%) suffered ransomware attempts in 2024, according to *The State of Ransomware 2024* report from Sophos. Firms with more than $5bn (£3.85bn) in annual revenue were hit the hardest, with 67% affected by ransomware in some capacity.

The financial and reputational impact on businesses can be devastating. Security personnel have even suffered PTSD-like symptoms after dealing with the fallout from such incidents.

Well-kept data back-ups, fail-safe systems and robust perimeter defences are all essentials for weathering the ransomware storm. Equally important is keeping a cool head in a crisis, showing confidence in decision-making and calmly working to resolve the situation.

That means ensuring all stakeholders are prepared, not just security teams. But what kinds of training exercises can help? And how can security leaders ensure that training is effective?

Cybersecurity training has evolved over the years. Ambushing employees with fake cyber attacks was once a popular method: for example, sending spoof phishing emails to test whether employees would click on a dodgy link. However, these techniques are increasingly being eschewed for more open and transparent training.

One method is to create simulated attacks, engaging so-called red teams, often brought in from outside the organisation, to blind test the cyber defences of blue teams, who comprise the cyber defenders in the organisation. The red team plays the role of the attacker, simulating the kinds of malware or ransomware used by cybercriminals. The blue team must surface these threats, protect against them and repair any damage.

> **If you have a drill without telling someone it's a drill, it can actually be just as disruptive as a real attack**

These simulated attacks take place on portioned-off parts of a network using synthetic junk data, so as not to put the business at risk.

The aim is to expose gaps in cyber defences. For instance, to test whether the organisation's defenders can detect threats, differentiate between false alerts and real threats and coordinate a successful response to an incident.

Organisations with the right expertise can create standoffs between internal red and blue teams. But many will need to partner with third-party vendors, which offer red-team services.

Lorenzo Grillo, head of Alvarez & Marsal's Europe and Middle East global cyber risk services, says such exercises are a "great opportunity to test a company's preparedness, detection and response processes and technologies in a way that mimics real world conditions". That's because they assess the entire control environment to simulate how skilled and motivated cyber threat actors would target an organisation.

However, surprise attacks can put unnecessary pressure on staff and risk making them feel as though they're under constant scrutiny, he adds. This can lead to trust issues between stakeholders.

Alan Woodward, professor of cybersecurity at the University of Surrey, says blind testing is akin to letting off smoke canisters in the office during a fire drill. Such an approach can burn out or panic staff and lead to poorer productivity.

"If you have a drill without telling someone it's a drill, it can actually be just as disruptive as a real attack," he notes.

Red teaming can help to iron out some kinks or expose certain vulnerabilities, but if leaders don't trust teams to perform in real crisis conditions, that "says more about your recruitment processes than anything," Woodward adds.

Instead, he recommends an open and transparent approach to cyber drills. He suggests regular tabletop exercises – typically a 90-minute role-playing session that sets out cyber scenarios for teams and leaders to work through. The National Cyber Crime Centre and the US government's CISA website provide some useful examples.

These games usually involve a facilitator to run the exercise, inform participants of what's happening and instruct them to make decisions. The aim is to create a plausible, realistic scenario and test responses, allowing participants to audit their current decision-making processes, technical defences and continuity plans.

Participants in these exercises can include IT teams, cyber experts and C-suite leaders. Crucially, organisations must include representatives from different segments of the 'command structure' – strategists, tacticians and operational employees.

Properly defined and managed tabletop exercises can help test a company's ability to respond to cyber crises, but businesses must learn to balance this approach with red-team exercises, Grillo says. "Red teaming can expose gaps and enhance defensive skills, while tabletop exercises offer room for safe practice and learning without constant stress."

During a real-world ransomware attack, an organisation's leadership must make tough business decisions. It will be up to them to decide whether to pay the ransom, issue a statement and find a way to ensure business continuity while restoring systems. It is essential that the C-suite is briefed properly and in a language they understand.

"Executives don't need to know how to perform log analysis or reverse engineer malware," says Dan Potter, senior director of operational resilience at Immersive Labs, a cybersecurity vendor.

In the past, adds Potter, security teams have succeeded in terrifying leadership with details of digital disasters and briefings on advanced persistent threat groups. They have been less effective at engaging the business. The facilitator of any tabletop exercises must prioritise inclusivity and encourage participants to speak a common language.

The goal should be continuous learning, says Potter. "One big exercise a year with the same 20 executives is not sufficient. It's not providing the regular cadence or the validation of processes that organisations need."

Given the busy schedules of C-suite executives, it may be difficult to find time for multiple tabletop exercises. This puts the onus on security leaders to keep teams sharp. Potter suggests frequent, small-scale exercises for first-line responders, including hands-on labs, technical skill development training or small-team simulations.

Security teams can then use these activities to brief senior executives on their progress. This can open conversations with leaders about concerns and priorities; it also helps to avoid exercise fatigue among leadership. Ongoing exercises will equip cyber teams with the data to inform leaders of their progress or areas where there's room for improvement, ultimately instilling confidence in the team.

Successful training exercises are built on a security culture that's rooted in collaboration and improvement, rather than shame or ridicule. Employees should understand the need for rehearsals and be clear that exercises are not about catching people out, criticising teams or blaming systems, says Jason Nurse, reader in cybersecurity at the University of Kent. The goal is to work out where there's room for improvement.

Tech leaders should carefully consider the targets, timing and nature of ransomware attacks in their simulations. As well as ensuring the exercises don't unfairly target certain groups, leaders must consider the state of the business before they implement a test exercise.

"For instance, is it the last day of the financial year?" asks Nurse. "Or is the simulation due on the day new software will be installed across the business? While there are certainly advantages to running simulations at these times – and ransomware groups themselves may find these ideal target times – they may cause significant additional stress for employees."

Finally, any business setting up a simulation should consider whether the content is appropriate. There have been instances where organisations have conducted attack simulations that were in poor taste and didn't consider the employee or customer context.

"We've seen attack simulations offering bonuses or alerting to Ebola outbreaks," Nurse explains. "There's a balance to be maintained in achieving and testing security processes without compromising employee morale."

Running cyber attack role-playing sessions might sound like corporate Dungeons and Dragons, but the benefits can be significant. By discussing actions needed to address these imaginary attacks, organisations can identify weak points in their security systems and skills gaps. When ransomware can lead to the destruction of businesses, running simulations can make all the difference.

ROOT CAUSES OF RANSOMWARE ATTACKS

Share of ransomware attacks with the following root causes in 2024, as identified by tech teams

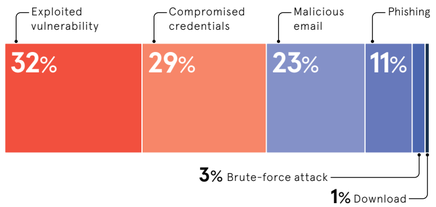

Sophos, 2024

From costly to cutting-edge: a new era of security analytics for CIOs and CISOs

Traditional security-information and event-management systems are no match for today's data complexity and cyber threats. Innovative analytics platforms offer a powerful solution with enhanced visibility, risk prioritisation and cost optimisation

Chief information officers (CIOs) and chief information security officers (CISOs) face tough decisions every day. They understand that harnessing and interpreting data insights are key to any effective cybersecurity strategy.

However, the task has become increasingly complex owing to the sheer volume and diversity of disparate data. Traditional security information and event management (SIEM) tools struggle to keep up, often demanding significant costs for increased data ingestion while relying on operationally laborious threat-detection capabilities.

"Some legacy systems have been in this space for years and have become a core component of security operations. They're very much embedded within processes that the teams are currently operating on," explains Randeep Gill, senior security strategist at Gurucul. "But they were not designed to cope with the realities of today's cyber-threat landscape, IT complexity, evolving regulations and data-sovereignty requirements. Nor were they built to handle the sheer volume of data organisations now face."

Security leaders are forced to make difficult choices about data prioritisation, resulting in either blind spots or unsustainable costs. "It's the lesser of two evils – you either pay a premium or accept a greater level of vulnerability," says Gill. This data security dilemma is not new. However, it is becoming harder to justify a decision to stick with legacy systems.

Indeed, research from 2021 found that half of security professionals were dissatisfied with their SIEM solutions, with 40% citing excessive costs and more having concerns over scalability and data management. This issue is compounded by the rapid growth in data generation – it is estimated that 90% of the world's data was generated in the last two years.

Legacy SIEM providers have attempted to keep up with the demands of modern organisations. But these often result in a patchwork of technology acquisitions or partnerships, which serve only as a band-aid to the problem and remain difficult to use and costly to run.

But what if organisations could reduce risks and costs simultaneously? Modern security-analytics platforms are doing precisely that, in a paradigm shift that addresses the limitations of traditional SIEM solutions and establishes the future of security operations.

Out with the old, in with the new

A new generation of security-analytics platforms address the challenges of the data dilemma facing security leaders. These big-data platforms leverage advanced machine learning (ML) models, artificial intelligence and automation to effectively and affordably gain complete visibility to detect and respond to real threats. They accomplish this in two ways.

The first is native data-pipeline management. This enables teams to accommodate large volumes of data from various sources, preparing it for analytics while also ensuring complete control over data residency. These modules filter non-critical data and direct them to low-cost storage, resulting in cost savings while allowing federated search from within the platform. They also enrich and normalise critical data for analytics readiness.

The second method is advanced analytics. This reduces false positives while streamlining investigation and response efforts by leveraging advanced behaviour-focused ML models. By centralising all relevant data, these ML models put anomalies into context to prioritise and escalate the most risky user and entity behaviour.

But the benefits go beyond cost and risk reduction. They also address the growing complexity and lack of visibility that has plagued many organisations.

Many security teams have invested in a patchwork of tools over the years to address the increasingly sophisticated threat landscape. But this only increases complexity and creates data silos that hinder visibility across the IT environment. This produces a deluge of incomplete alerts, resulting in false positives, which require manual and cumbersome investigations across various tools to validate.

"Getting real value out of these tools has traditionally been a huge challenge and security teams have been disappointed by false promises," says Phil Close, VP of Europe at Gurucul. "To reap the true value of these legacy systems, you have to spend an inordinate amount of time managing, maintaining and navigating across platforms. It's time these teams don't have."

SECURITY OPERATIONS CENTER (SOC) CRITICAL NEEDS ASSESSMENT

Gurucul's autumn 2023 survey of 204 cybersecurity professionals across the US, EMEA and APAC highlights the need for detection clarity while balancing operational efficiency and cost. Respondents were decision-makers or influencers in organisations with 1,000 or more employees and internal SOCs

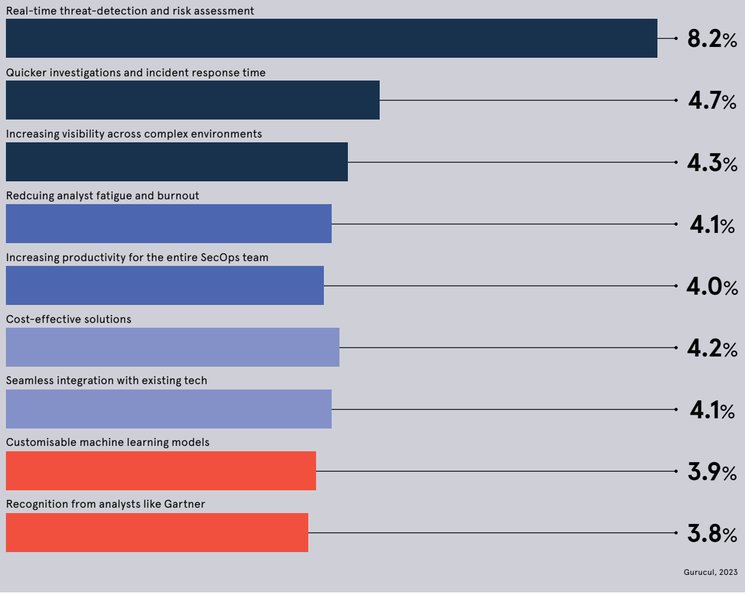

Gurucul, 2023

> Getting real value out of these tools has traditionally been a huge challenge and security teams have been disappointed by false promises

When approaching a security operations modernisation initiative, beyond selecting the right platform, security leaders must ask some critical questions: what does our current cybersecurity framework look like? How are we assessing risk in the context of our environment? How are we measuring the effectiveness of our tools and security posture? Where are our blind spots and what data collection is necessary to illuminate them? Even with the best tools in the world, unless leaders understand the risk in the organisation, they're going to fall at the first hurdle.

"The security-operations centre needs a single source of truth, where all insights should reside. Your analytics shouldn't be running from disparate components within your organisation," says Close.

Eliminate the risk of doing nothing

It's important for security leaders to ensure their SIEM is capable of handling risks in a way that is cost-effective and manageable without sacrificing security. Ultimately, those security leaders must consider the cost of inaction.

That's because many times leaders will know the technology isn't fulfilling their organisation's needs, but they are reluctant to act out of fear that any new investment might be too risky or will just add more costs and complexity to their operations. But reducing cost no longer means increasing risk. The next generation of SIEM solutions can address an organisation's commercial and operational costs without compromising security.

Gurucul's Reveal security analytics platform is designed for agility, flexibility and scalability. Powered by advanced ML and AI, Reveal delivers high-fidelity threat detection and risk prioritisation in real time, cutting investigation times by 50% and eliminating false positives.

Moreover, Reveal offers substantial SIEM cost savings, typically exceeding 40% in reduced data costs compared with traditional SIEM. In essence, next-generation platforms such as Gurucul Reveal can remove those barriers to action.

As organisations navigate the evolving threat landscape and grapple with the limitations of their existing security tools, the time has come to embrace a new era of security analytics. By leveraging the capabilities of modern platforms, CIOs and CISOs can reduce costs, mitigate risks, minimise complexity while maximising analyst output and gain the comprehensive visibility they need to protect their organisations effectively.

For more information please visit gurucul.com

Gurucul

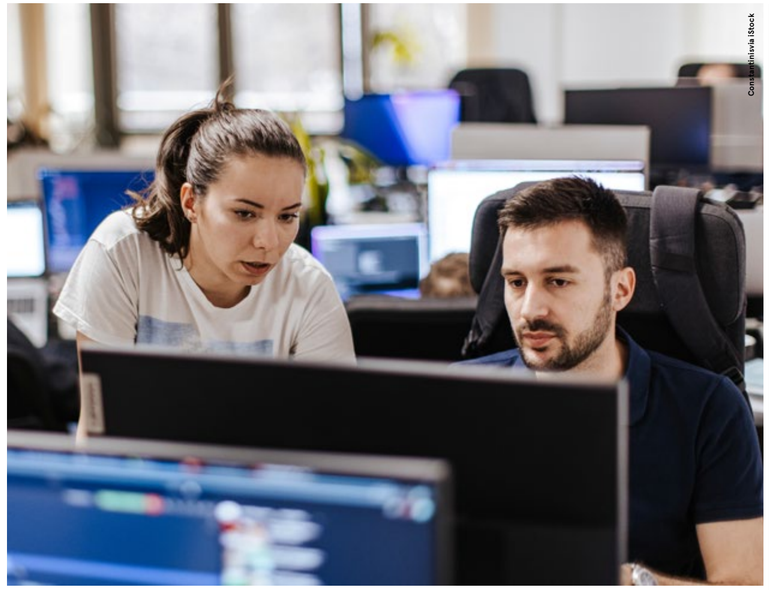

Constantinovia iStock

ARTIFICIAL INTELLIGENCE

As CIOs grapple with GenAI, MIT offers a two-step solution

Organisations are keen to reap the benefits of AI, but many struggle to implement the technology. Success means focusing on tools and solutions, according to MIT

Jon Axworthy

Does your CIO look tired? Perhaps GenAI worries are keeping them up at night.

They certainly face a tough task. CIOs must figure out how to get the most out of the fast-evolving technology and generate business value. The pressure to deliver outcomes has resulted in a lot of trial and experimentation, but a road map for success hasn't been easy to find, especially as use cases vary wildly depending on an organisation's position on its AI journey.

But new research from MIT's Center for Information Systems Research suggests a process that could enable CIOs to implement AI into workflows quickly and safely.

The research was inspired by questions from CIOs and their peers on why they aren't getting the same value from GenAI as they have from data and analytics technologies in the past. Based on a series of virtual roundtable discussions with data and technology executives, it identifies a need to separate the technology into two distinct parts – tools and solutions – before deploying them in a two-step strategy.

AI tools "are designed to be broadly applicable", according to Dr Nick van der Meulen, who co-authored the research. They could include conversational systems, such as ChatGPT, Claude or Gemini, as well as digital assistants embedded in existing productivity software.

"An employee will use a GenAI tool to summarise a document, brainstorm ideas, rewrite an email or analyse financial results," says Van der Meulen. "As one executive in our study put it, they allow for 'productivity shaves'."

Crucially, the report reveals that AI tools also help employees get comfortable with using AI and are important mechanisms for building data democracy in an organisation.

However, it also emphasises that CIOs must understand some basic principles of usage with this first step, most importantly putting in place certain guardrails and backing it up with workforce training.

"Unvetted GenAI tools, in the form of 'bring your own AI', can bring significant risks for an organisation, including data loss, intellectual property leakage, copyright violation and security breaches," explains Van der Meulen. "The guardrails should outline which tools are acceptable and any conditions that may apply. For example, a company may permit GenAI use when prompts draw on publicly available information but disallow it if prompts require company data."

The MIT research also notes that employees shouldn't be left to explore tools independently. There must be company-wide training to teach them how to effectively and responsibly instruct and interrogate GenAI tools so they can get the most out of them.

With these guidelines in place, CIOs can be assured that tools are being used safely. This will also help foster a self-perpetuating understanding of AI best practices across the organisation. As more staff use the tools correctly, best practices will become the norm.

Once a sound knowledge base has been established, CIOs can further build AI architecture and expand its horizons with the introduction of GenAI solutions, which help groups of employees to transform workflows and create value.

For example, Van der Meulen says the research team has "heard from a number of call centres that use LLMs to transcribe calls as they happen and process the content and tone of conversations. This is then used to coach agents in real time to either recommend empathetic responses to frustrated customers or propose upselling opportunities for satisfied ones."

The key to success is to pursue both tools and solutions but use different strategies that dovetail to create a virtuous cycle.

"GenAI tools can serve as a form of grassroots innovation," says Van der Meulen. "Employees can discover promising use cases that can later evolve into more formalised, scalable and lucrative GenAI solutions."

Organisations at different stages of the AI journey must adopt different strategies. The report recommends that the best starting point for GenAI implementation is the targeted adoption of just a few tools from trusted vendors, accompanied by close oversight.

Those further along in their journey should shift their focus to developing GenAI tools into solutions that contribute to strategic business objectives.

For instance, NN Group, an international financial services company, created a ChatGPT 'playground', where employees can use various GenAI tools to test their ideas and figure out ways to make their work more efficient.

"The playground is available to all employees. With a few ground rules in place and by making it easy to use, there is no need for employees to use unsupported tools outside of the playground," explains Tjerrie Smit, NN Group's chief analytics officer. "Launching the playground has been a game-changer for us. It provides a secure and compliant environment where our employees can safely experiment with GenAI. This proactive approach not only encourages innovation but also ensures that we can scale successful ideas into impactful AI applications across the organisation."

> GenAI tools can serve as a form of grassroots innovation. Employees can discover promising use cases that can later evolve

One of the main takeaways from the research is that businesses can choose their approach: buying, boosting or building an AI solution.

Buying means using vendor-provided solutions where the vendor manages the model and operations. Boosting enhances vendor-provided models by incorporating proprietary data through techniques like fine-tuning or retrieval augmented generation (RAG), which customise pre-existing GenAI models with more relevant information from company sources. Building is the most resource-intensive approach, where organisations take full ownership of developing, running and maintaining the model.

"Buy or boost GenAI solutions when you need to move fast and gain competitive parity," advises Van der Meulen. "But build when you need a differentiated GenAI solution that is hard to imitate and provides a competitive advantage."

CIOs must remain vigilant when it comes to business alignment, so that GenAI is never siloed and left in the hands of a few select technologists, as this will starve it of the oxygen of innovation.

As the MIT research suggests, the surest way to accelerate AI's value to an organisation and ensure it is safely embedded is to make it more accessible to employees.

ADOPTION ACROSS THE ORGANISATION

AI adoption rates, in production or developing, across selected use cases

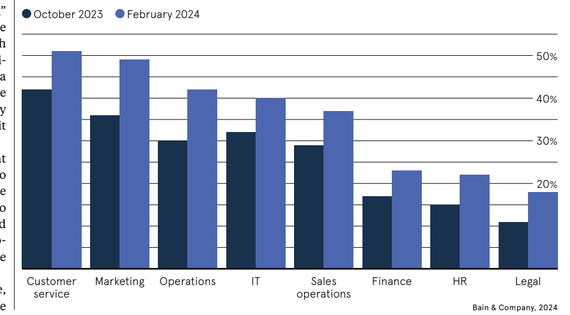

Bain & Company, 2024

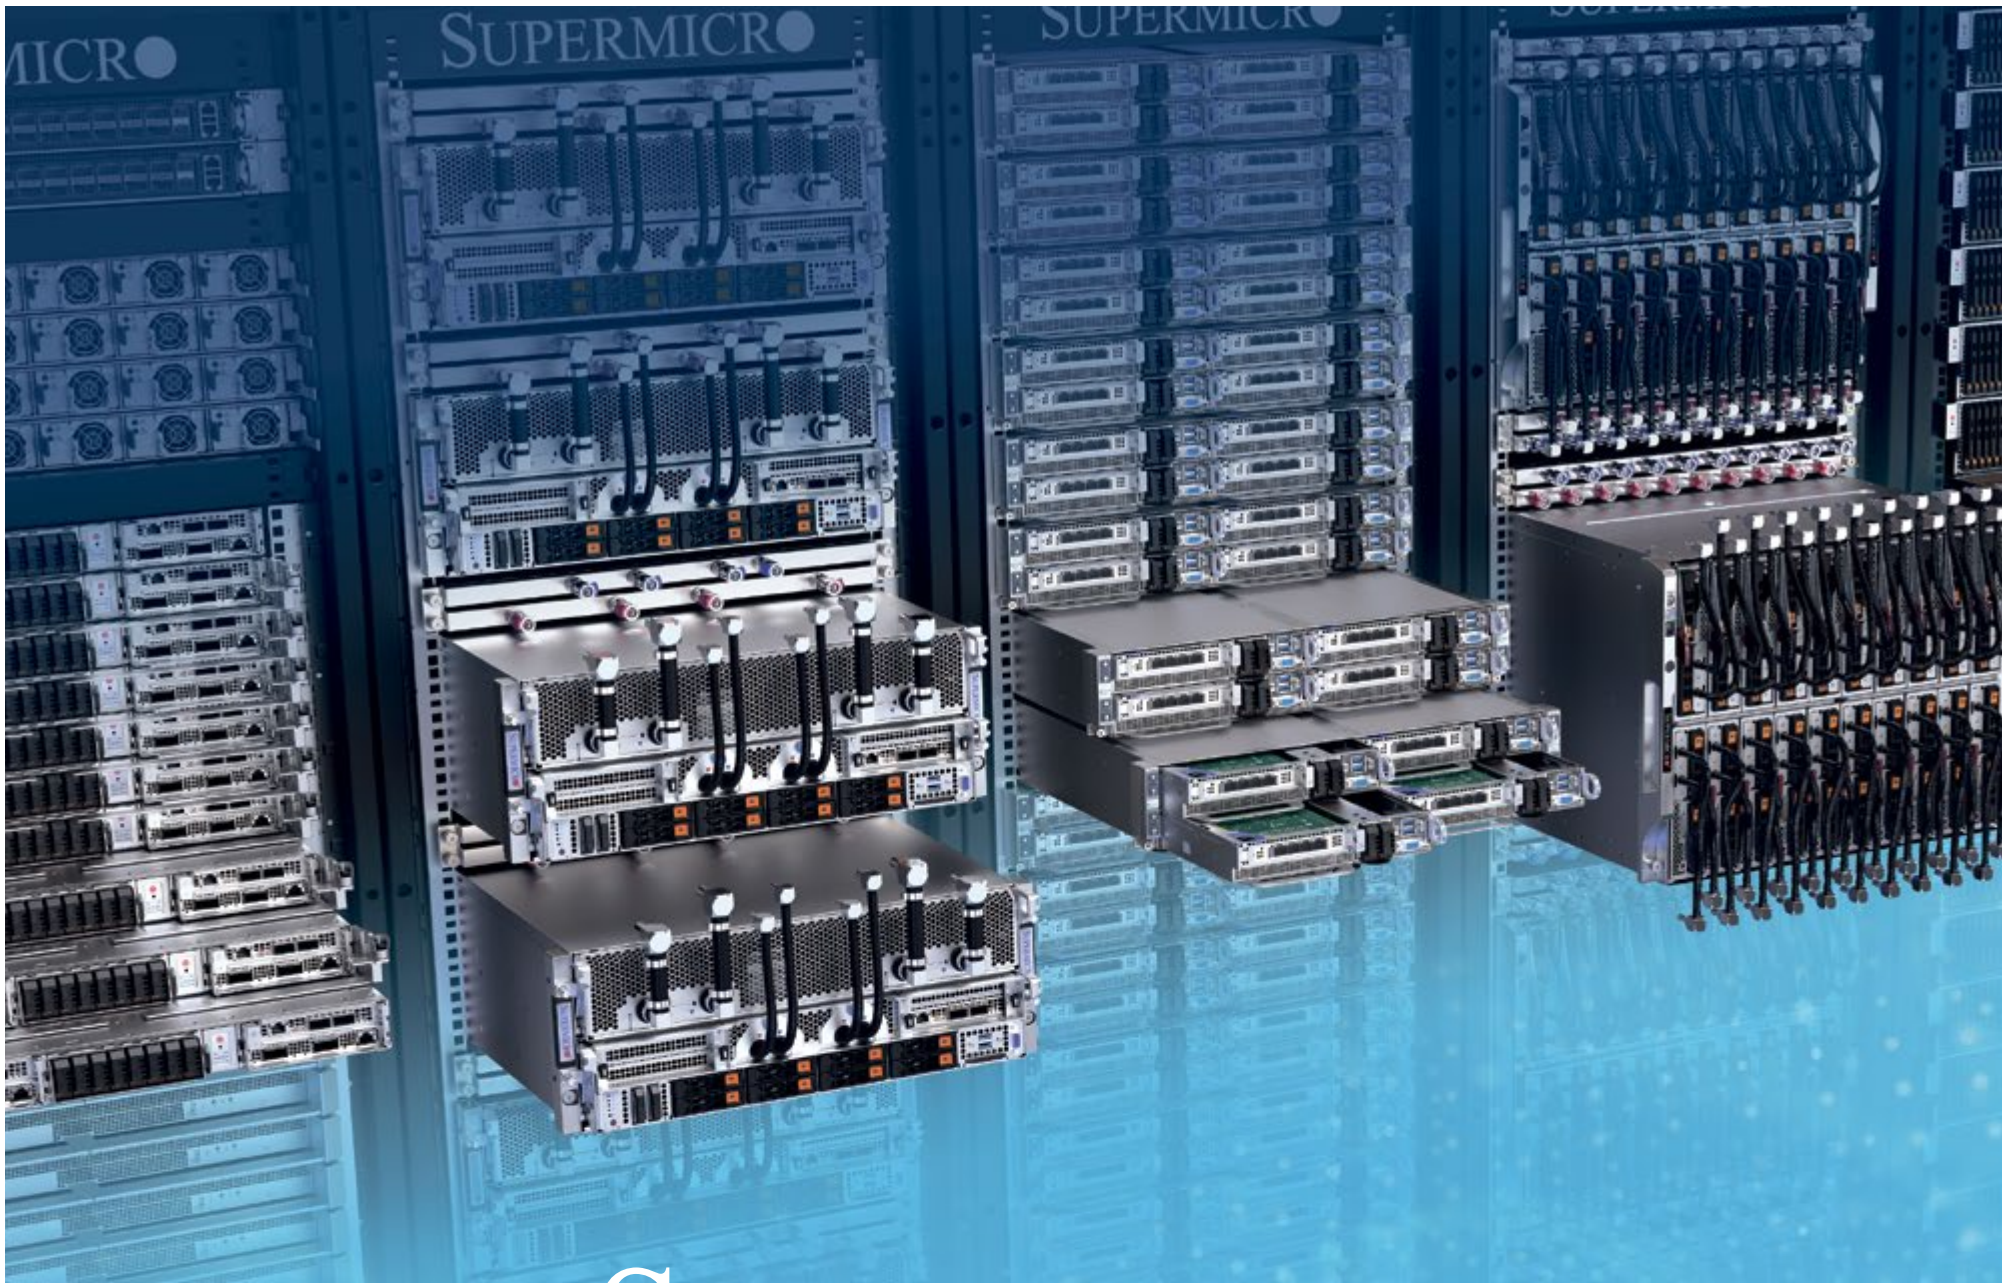

SUPERMICRO

Rack-Scale Liquid Cooling

Total Solutions for AI and HPC

SUPERMICRO | NVIDIA

Learn More at www.supermicro.com