smartCIO

EMEA Edition | Brought to you by Workday | Vol. 10

The Trust Equation: Responsible AI in EMEA

Also in this issue:

Five ways tech platforms can protect private data

Build employee trust – and enthusiasm – with responsible AI

Scan for previous issues:

Editor’s comment

Welcome to the latest issue of smartCIO magazine, where we provide you with insights on the latest technology trends, opinions and regional perspectives from IT leaders across EMEA.

If you followed events from the World Economic Forum in Davos, then you’ll know that trust was the central narrative at the annual gathering of the world’s political and business leaders. And as AI continues to dominate technology, media and business agendas, C-suite leaders must wrestle with the challenges of trust and responsible AI.

In this issue, we explore why the success of AI hinges on the ability of AI leaders to win the trust of the workforce and deploy the technology in a responsible way. We’ll also look at the broader role of trust across IT and how organisations are facing this challenge head on.

IT commentator Martin Veitch has seen his fair share of seismic technology shifts, and he gives his thoughts on how IT leaders can win the trust battle.

We’ll also delve deeper into a new global research study from Workday that charts the existence of a trust gap between executives and their employees – and what business leaders should be focused on if they are to close the gap.

Meanwhile, Anja Fordon unpacks some of the reasons why AI excites and alienates in equal measures, as she looks beyond the fear and hype of this technology phenomenon.

In a feature with Workday customer Veolia, we look at how to build trust in a time of data insecurity.

And we’ll also hear from Workday partners Accenture and KPMG, who offer tips and advice on how to build employee trust – and enthusiasm – with responsible AI and by addressing the IT skills gap.

We hope you enjoy this issue and find it useful on your own AI journey.

Angelique De Vries - Schipperijn
President, EMEA, Workday

Contents at a glance

Join the conversation in three ways

SUBMIT story ideas, **CONTRIBUTE** your own articles for publication or **SUBSCRIBE** to be a part of the smartCIO community and receive the quarterly digital issue and info on local events.

One address does it all!

Email us at **smartcioemea@workday.com**

In the AI age, leaders need to build trust... but how?

With great uncertainty surrounding AI, security and governance teams are operating in a sea of change. But one thing's for certain: trust has never been more important.

By **Martin Veitch**,
Industry Commentator

Trust, my trusty dictionary tells me, is "belief in the reliability, truth or ability of someone or something". That's a decent starting point but we all know that it can be hard to achieve. As children, it is drummed into us that "trust is earned, not given". In adulthood, the importance of transparency, honesty and accountability are repeatedly proven. Sadly, however, we also become increasingly familiar with just how tenuous trust can be, how easily breached and lost. We all know the importance of trust and nowhere is trust more valued than in the process of change, the uncomfortable and unfamiliar terrain that must be negotiated to get to a better place. And with AI, companies are undergoing the mother of all changes.

What does trust mean?

The fact is that trust is a complex matter. Paul Thagard in Psychology Today has called it "a complex neural process [that is] rarely absolute, but … restricted to particular situations … a binding of current experiences, memories and concepts". In other words, we aren't like Professor Pangloss in Voltaire's novel Candide, blindly trusting everybody. Instead, we build up our sense of who we trust and in what circumstances and to what degree, based on a lifetime's body of direct and indirect experiences. And once that trust is exposed as being fool worthy, reinstating it is very hard.

Few of us will build our own LLMs or feel the need to create bespoke core applications outside those that create direct competitive differentiation. That means selecting vendors and other partners that have strong records on data stewardship.

Paul J. Zak, writing in the 'Harvard Business Review' has gone further, measuring the production of the chemical oxytocin to show that trust can lead to dramatic positive shifts in stress levels, energy, productivity and engagement. Even without that sort of hard evidence, most of us will instinctively agree that trust is a good thing that leads to positive outcomes but, particularly in the wake of The Great Resignation and quiet quitting, all leaders need to invest in building trust and understand why it is a better route than blunt retention measures such as imposing golden handcuffs. This is especially the case today for CIOs where AI and machine learning (ML) will surely act as super-catalysts that will lead to massive changes in how we work.

Change is tough without trust

Change is one of the toughest aspects of business and you may have heard a version of the old wisdom that while change is tough, not changing will end up being far tougher. Change requires a well-planned strategy of course, with lots of due diligence to show that the decision to enter a new market, geography or business model is correct. But much of the challenge in change lies in soft skills, the ability to lead, persuade and build consensus. And underpinning all of these is trust.

Leaders need to win over employees, partners and customers. AI puts trust front and centre again because it works by collating large data sets so it has never been so important to lay out clearly what data is being collected, how it is being attained and how it will be used, together with transparency as to risk factors and unknowns. For AI not to stand for 'Angst Inducing', we need to implement safeguarding policies now and make sure everyone can access and understand them.

To what extent do we trust AI today? Not much, which is perhaps no surprise when we think about generative AI issues such as hallucination – AI making stuff up, effectively – or LLMs sucking data from unknown sources on the public internet and not always respecting copyright. New research conducted by FT Longitude for Workday points to a clear gap between leaders and staff:

- **Mind the gap:** 70% of leaders welcome AI and 65% are confident their organisations will deploy it in a trustworthy manner but for employees the numbers are just 46% and 51% respectively.
- **Manage, don't dictate:** Four out of five employees say they have observed no collaborative interactions with their employers over AI and have received no usage guidelines.
- **Educate and pledge:** One in four employees are not confident their organisations will place their interests above those of the organisation. Also, 69% of leaders predict a future scenario where AI reduces manual labour to a significant degree – but just 38% of employees agree.

This research is aligned with a growing body of evidence as to the importance of trust in AI. A 2023 report by KPMG with the University of Queensland, Australia, based on 17,000 global participants, found strikingly negative reactions with 61% saying they feel ambivalent or wary towards AI. However, looked at more closely, the data yielded more granular results with, for example, far more unwillingness to trust AI in HR compared to AI in medical diagnosis. Also, it's worth noting that prejudice against AI can't be purely ascribed to negative Luddite views, as 85% of respondents said they believed AI will result in a range of benefits.

We also know that AI is one of the biggest and fastest-moving booms ever in technology history, making it difficult to predict the next twists and turns in how it operates, how it is applied and how it is policed. So, if your organisation is betting on AI for significant business change, it will be as well to explore the concerns of your people and partners then work out how you assuage them.

As with every major technology change then, we know there are obstacles to be addressed. Think of cloud, SaaS, blockchain or e-commerce: there are tricky stages we go through before we feel comfortable. Many of us will feel that we are in a Wild West scenario and miles away from what Gartner calls the 'Plateau of Productivity.' So, how do we get there?

Building trust, one brick at a time

When we think about practical steps, a few stand out.

Use a trusted technology platform

Few of us will build our own LLMs or feel the need to create bespoke core applications outside those that create direct competitive differentiation. That means selecting vendors and other partners that have strong records on data stewardship and rigorous measures to safeguard AI. Ask your suppliers what they are doing in AI governance and security, then ask them to prove it. Request access to customers who are peers. Seek evidence of architectures that prevent bias through techniques such as 'dynamic grounding' to capture only the most reliable and up-to-date information from LLMs. Look for strong access and retrieval controls and data-masking capabilities to protect sources. Insist on tight retention policies and the ability to identify and block toxic content. Beyond this, seek out vendors that are heavily involved in setting standards and putting in place safeguards for AI.

70% of leaders welcome AI and 65% are confident their organisations will deploy it in a trustworthy manner.

Draw the line between automation and augmentation

One elephant in the AI room is that there are widespread fears that it will lead to a massive disruption of white-collar working as certain human jobs become thought of as no longer necessary because machines can do them better. It's critical that leaders explain that AI is about augmentation and replacing human chores with machine intelligence, liberating human beings to deliver what they are good at: empathy, creativity, collaboration and problem solving.

Sunlight is the best detergent, so keep people in the loop and give them a voice. Accenture is an example of a company that has put its cards on the table and clearly said that it does not plan job cuts, but it does expect massive productivity improvements through AI. That sort of clear messaging will go a long way to reassuring staff who may be feeling vulnerable or exposed. And if your organisation is not willing to say that and is eyeing the opportunity purely as a vehicle for mass job cuts and cost cutting, then maybe it's time to consider your employment options.

Walk the walk

The politician who says there will be no tax rises and then hikes them anyway immediately loses faith. The business that champions ethical conduct and then exploits its workers or collaborates with unethical partners also burns through trust. Actions speak louder than words, so promises on AI need to be followed up on.

AI is a CIO opportunity

In his book 'The Open Organization', former Red Hat CEO Jim Whitehurst commended the notion of democratising decision-making across the enterprise and even outside it. This included the idea that leaders should make it clear when they don't know something rather than pretending to be omniscient.

Most modern CEOs should be aware of the implications of AI for their organisations, but they will urgently need the assistance of CIOs and other specialists to understand technical, legal and ethical complexities. By building trust and by getting ready now for the probability of massive technologically enabled disruption, smart leaders will be able to ride what promises to be one of the great business waves of our times.

As one CIO told me, "Trust is a two-way street and it is hard to negotiate. Nobody can put their finger on exactly what are the skeins and fibres that bond us to have faith in each other. But we know that once it's gone, it's gone, so the message has to be to handle with care."

All leaders need to invest in building trust and understand why it is a better route than blunt retention measures such as imposing golden handcuffs.

Expert perspectives on the future of AI in business and society

By **Anja Fordon**,
EMEA Staff Writer

Check out what the experts had to say about the future of AI at last year's Workday Rising EMEA. Expect insights from Dr. Tomas Chamorro-Premuzic and many more.

In an era where AI and machine learning (ML) are increasingly becoming integral to our daily lives, understanding their implications is more critical than ever. At Workday Rising EMEA, experts from various fields gathered to discuss the transformative potential of AI and how to bridge the trust gap that currently exists in this domain. Here's what they had to say:

Understanding AI: Beyond the hype and fear

"AI is a defining technology of our times," said Psychologist, Author and Entrepreneur, Dr. Tomas Chamorro-Premuzic, emphasising its permanence and influence. However, there's a need to move past extreme views in mainstream media and adopt a nuanced understanding of AI's capabilities and limitations. This balanced perspective is vital for workers and businesses to effectively leverage AI. Here are some of the highlights of the video:

AI in the workforce: Amplification and mitigation: Kathy Pham, Vice President AI & ML at Workday, highlights the historical context of managing work and finances, stressing the importance of understanding these processes in an analogue world before integrating AI. "The speed of technology and its ability to amplify our work is high," she noted, suggesting that AI and ML can both enhance and challenge traditional business practices.

“

AI is a defining technology of our times

Dr. Tomas Chamorro-Premuzic,
Author and Professor of Business Psychology

Reskilling for the AI era: The role of finance and HR leaders in transitioning workers from mundane tasks to more creative activities is crucial, according to Dr. Tomas Chamorro-Premuzic. However, this shift isn't automatic. It requires active management and incentivisation. Upskilling is essential for leveraging AI's next phase.

The trust gap and human-machine interaction: Filip Gilbert, Global Workday GTM & HR Technology Lead at Accenture, identified a trust gap concerning AI, centred around personal impact and employability. Building trust involves demonstrating continued investment in employees and making AI's role in augmenting activities understandable and predictable. "Human-machine interaction [...] will actually increase that trust," he asserted. Interestingly, Dr. Tomas Chamorro-Premuzic pointed out that distrust towards AI often reflects existing organisational distrust. The disparity in AI perception between employees and leaders indicates a misalignment in strategic business goals, an issue that requires leadership intervention. New research by Workday shows a trust gap has formed at all levels of the workforce but especially among employees. You can read the whole report to learn how to close the trust gap.

Regulation, policy and AI implementation: Pham and Chandler Morse, Vice President Corporate Affairs at Workday, discussed the opportunities for organisations to collaborate with governments on AI regulation and policy. Such collaboration can lead to better user experiences and products. Morse also mentioned the importance of responsible, transparent AI implementation, highlighting consensus among industry leaders on best practices for bias mitigation, privacy and explainability.

These insights from industry leaders provide a valuable roadmap for navigating the complexities of AI in the workplace and society. By addressing the trust gap, fostering transparent and ethical AI use and prioritising skill development, we can harness AI's transformative power responsibly and effectively.

“

The speed of technology and its ability to amplify our work is high

Kathy Pham, Vice President AI & ML at Workday

Three insights to help CIOs navigate evolving AI regulations

The age of AI is causing huge disruption to the way the world does business – and so we can expect a raft of new regulations to define how to use AI responsibly. The question is, how can CIOs prepare and respond to those incoming regulations? Here are three key insights we think you need to know.

By **Steve Dunne**,
EMEA Staff Writer

As global leaders shape AI policies, CIOs should keep transparency, data privacy and vendor accountability top of mind.

AI is evolving at breakneck speeds but as is often the case when it comes to technology, policy is struggling to keep up.

Still, the next couple of years promise to be big ones for AI regulations as global leaders create policies aimed at governing next-gen AI applications, including large language models (LLMs) like ChatGPT. European Union policymakers, for one, have agreed on the basics of the AI Act, a sweeping set of laws meant to capture the potential of the technology as well as buffer against its risks.

The exact shape that AI regulations take is still evolving and will likely need constant updating. After all, ChatGPT is still in its infancy. A few of the key questions being considered are:

- In a global economy, how will different countries enable and limit the use of AI?
- How can AI leverage data while ensuring sensitive information remains secure and protected?
- What practices can best mitigate bias in AI applications and outputs?
- What documentation will be required to prove that AI has been developed responsibly?

As government agencies and NGOs continue to grapple with these crucial questions, CIOs find themselves in the hot seat. While forging ahead amidst regulatory uncertainty comes with risks, delaying the development and deployment of AI applications could have long-term consequences on profitability and growth.

Risks aside, 60% of businesses are adopting AI and machine learning (ML) in some way, according to the C-Suite Global AI Indicator Report by Workday. Furthermore, the research found that IT leaders will most likely be the ones expected to make a company's AI deployment a success. To stay ahead of the curve, CIOs must identify how the business can benefit from AI, define clear use cases, and introduce governance policies that will enable the responsible use of AI innovation.

"If you're regulating AI, you can't go about it by regulating the technology. Technology evolves. So you have to regulate the users and look at the context," said Thomas Boué, Director General, Policy at Business Software Alliance (BSA). "In high-risk uses of AI, the idea is not to prevent them from happening, but to put the safeguards in place to ensure that AI can be used, developed and deployed for the benefits of society."

Here are three insights CIOs can use to guide their AI practices as both the technology, and the regulations surrounding it, evolve.

Transparency is king

No CIO wants to invest in innovation just to have regulatory changes block their organisation's path forward. But the AI opportunity is too promising to pass up. To make meaningful progress in an unpredictable market, CIOs must ask for – and enable – transparency across the enterprise and the ecosystem.

This starts with clearly communicating where and how AI will be used, as well as what the organisation aims to achieve. With so many unknowns framing the AI conversation, radical transparency helps set expectations about what applications will be able to achieve, alleviate stakeholder fears and demonstrate accountability.

CIOs can promote transparency by pulling back the curtain for internal and external stakeholders, including regulators. Outlining data handling practices and privacy measures in detail can help organisations prove that data has been used ethically and transparently. Providing insight into which algorithms were chosen and why can also showcase how bias is being considered and addressed.

While exactly what kind of documentation will be required by regulatory bodies is still being determined, "there's a tremendous consensus on transparency," said Chandler Morse, Vice President, Public Policy at Workday. "For example, people believe that, if you use AI in HR, there should be full transparency on what's happening, how it's happening, what data is being collected and what inferences are being made."

When companies communicate openly about AI, employees are also more willing to ask questions about how and where to use new applications. This decreases the risk that teams will use AI inappropriately – and increases the likelihood that they will innovate confidently and responsibly.

Explainability reduces risk exposure

While compliance guidelines are being developed, explainability should be a CIO's North Star. In the context of AI, explainability is concerned specifically with decision-making. Transparency provides visibility into how AI is developed and deployed, but explainability focuses on how the system thinks – and the logic it uses to come to conclusions.

Regardless of whether governments require regulatory pre-approval or self-assessment of AI – the European Union has chosen self-assessment, putting the burden on software developers and AI providers, said Jens-Henrik Jeppesen, Senior Director, Public Policy at Workday – CIOs will need to easily communicate the internal workings of these applications. For example, businesses may be asked to prove that no copyrighted material was used to train their AI, even if a third-party developed the model it is based on.

Employee management offers a case in point. When AI is used to inform hiring, promotion or termination decisions, CIOs will be asked hard questions about how the AI was trained, how bias was addressed and how private data was protected during implementation. That's critical considering that some countries are considering legislation that makes companies using high-risk AI models – such as those meant for healthcare or education – more responsible for any damage that results from that use.

It's not that companies need governments to tell them how to build the technology – they don't. Rather, they need safeguards in place to assure their customers that these products and applications are safe.

> "If you're regulating AI, you can't go about it by regulating the technology. Technology evolves. So you have to regulate the users and look at the context.
>
> **Thomas Boué,** Director General, Policy, Business Software Alliance (BSA)

> “
> People believe that, if you use AI in HR, there should be full transparency on what's happening, how it's happening, what data is being collected and what inferences are being made.
>
> **Chandler Morse,** Vice President, Public Policy, Workday

General purpose AI models, known as foundation models, can be fine-tuned to complete a variety of tasks. Companies often purchase these general models from vendors – but once a company incorporates a foundation model into its products or operations, its leaders will be responsible for assuring regulators that the technology is in compliance with new rules. This means CIOs must ensure they are delivered with comprehensive documentation, including background on model architecture, feature engineering, testing procedures, and security measures.

“So companies ought to have very close conversations with their vendors to make sure that they have the emerging regulations firmly in hand, and that they have governance programs that are aligning with those emerging regulatory requirements,” Jeppesen said.

CIOs must also track where internal teams use the foundation model, how it has been integrated into the company’s products and operations and what additional data was used to fine-tune the application. As adoption increases, ensuring explainability will be an enterprise-wide assignment. “AI is no longer an off-the-shelf thing that you install on the system and it just works,” said Boué. “It's something that is negotiated, that is discussed and that changes all the time.”

Safeguards encourage innovation

Businesses rarely clamour for more regulation but, when it comes to AI, most industry players agree that better guardrails are needed.

It’s not that companies need governments to tell them how to build the technology – they don’t. Rather, they need safeguards in place to assure their customers that these products and applications are safe. “There is a level of trust that comes with regulatory surety,” Morse said.

To build confidence while regulations are being developed, CIOs should examine what AI leaders are doing in this space. For example, early adopters that are rolling out AI for HR are focused on protecting the fundamental rights of individual employees, applicants and candidates each step of the way.

Taking a proactive approach also prepares global companies for the inevitable variation in AI legislation around the world. Defining key terms and aligning with the core principles of responsible AI, including transparency, explainability, discrimination and bias mitigation, and privacy protection, can help companies go further faster – while also enabling collaboration across jurisdictions.

“There is a commonality of purpose, which is to have an interoperable environment for innovation and deployment of these technologies,” Jeppesen said. “Most countries have broadly similar objectives – to have all the benefits of this technology, while ensuring that it is safe, trustworthy and can be used with confidence.”

Five ways tech platforms can protect private data

Data is the lifeblood of many organisations, but when unprotected, it can cause more harm than good. As data privacy gains more attention, here are the tools and practices companies can utilise to boost data privacy and security.

By **Patrick Evenden,**
EMEA Staff Writer

As private data becomes ever-more valuable, it's also getting harder to protect. Hackers are using AI to sneak into IT systems faster and more effectively – and companies must strengthen their security to keep up. At the same time, regulations are evolving rapidly, often raising the compliance bar.

In this environment, business leaders say security and privacy are the greatest risks to leveraging AI and machine learning (ML) in their organisations, according to the C-Suite Global AI Indicator Report by Workday. To manage these risks while also tapping into the benefits of AI, CIOs need technology to do more on all fronts.

In addition to making data easier to organise and analyse, technology platforms can help organisations stay ahead of cyber threats and privacy rules. With the right cloud-based tools, IT teams can build privacy into systems from the start, rather than continually playing catch up as new risks and regulations emerge.

Here are five ways tech platforms can help organisations navigate new privacy needs in tumultuous times.

1. Make transparency automatic

Collecting user data has countless benefits – but it also comes with serious risks. Not only do most jurisdictions have strict rules about how private data can be collected and used – rules that will only intensify with increased use of AI and machine learning – they also require companies to clearly outline for users what they plan to do with it.

“In recent years, some companies have received multi-million-dollar fines as a result of failing to meet the requirements around transparency and provision of information,” said Patricia O'Gara, Senior Principal, Data & Privacy Engineering at Workday.

Full transparency is necessary for end users to make informed decisions about the type of permissions they want to grant a given organisation. But it can be difficult to deliver the legal information people need in an accessible way. The right technology platform can help companies present privacy information persistently, either on a homepage, in a footer or within a central dashboard that users visit frequently. However they're presented, providing clear links to privacy notices that can be updated as needed helps ensure users can access the required information with the click of a button.

Just 34% of respondents in a global privacy survey said they have conducted data mapping and understand their organisation's data practices.

2. Put users in control

The idea that individuals should retain control over their data is at the heart of most privacy laws. While the exact requirements continue to evolve, proactive companies can stay ahead of the game by letting users decide what data can be used for what ends.

For example, a company might want to track user metrics on their website using analytics software. This can help personalise marketing, career-focused messaging or inform future offerings, but users must opt-in to sharing this type of data via cookie banners.

However, that's only one piece of the puzzle. People should also have control over how their data is being stored and processed – which requires companies to give users a peek behind the curtain. For example, if an employee requests access to their personal data, IT teams need to be able to quickly create a report that shows what information the company is tracking, who can access it and how it is leveraged to inform decision-making. Many organisations, however, have room for improvement – just 34% of respondents in a global privacy survey said they have conducted data mapping and understand their organisation's data practices.

Admin guides and fact sheets can help companies clearly communicate how personal data is used by ML models, giving people the context they need to make an informed choice about what they will allow.

"It's really about what data is used as input, what the output of the machine learning capability is, how we are doing bias evaluation and how our machine learning model is trained," said Sabine Hagege, Director, HCM Product Strategy at Workday. "People need a lot of information to understand how the data is processed."

3. Get granular with consent

In many situations, users will be comfortable sharing some personal information for specific purposes. Companies are then responsible for making sure that data is only used in approved ways. And if a company works with consumers or employees in multiple jurisdictions, it must ensure that data isn't shared with or pulled from regions with different privacy laws.

How can CIOs navigate all the moving parts? It starts with the proper configuration. Technology platforms that offer a localisation framework give IT teams the power to determine what type of information can be tapped for different people based on who they are, what role they play and where they're located.

"It's best when you can configure for each purpose you collect data for. So is it for diversity and inclusion or statistics and metrics?" said Hagege. "Then, on a country-by-country basis, use the consent response to configure your other processes and control how that data is used."

> How can CIOs navigate all the moving parts? It starts with the proper configuration.

> "
> It's really about what data is used as input, what the output of the machine learning capability is, how we are doing bias evaluation and how our machine learning model is trained.
>
> **Sabine Hagege,** Director, HCM Product Strategy, Workday

4. Purge data you don't need

Many privacy rules also demand that personal data be deleted when it's no longer needed. Consent should be given for a specific purpose and timeframe – and companies must permanently erase, or purge, that information afterwards.

To stay in line with expectations and regulations, each company needs a data purge plan. CIOs should work with their IT teams to determine which data should be purged when – and then schedule mass deletions on a regular basis.

However, that alone is not enough. Companies must also be able to purge an individual's data at will, either because their status has changed or because they've requested it. For example, a CIO might want the data of every terminated employee purged immediately after they leave the company. Or a job candidate might request their data to be purged if they aren't hired.

IT should make it easy for people to get their data deleted – but it's important to remember that "purging is irreversible," said Hagege. "So it's very important that you implement some controls and make sure that whoever has access to purge is fully aware that it can't be undone."

5. Keep private data confidential

Getting consent to collect and use private information doesn't make it any less private. CIOs must keep this in mind when determining who can view what data – and take the steps needed to keep sensitive information confidential.

Certain types of data, such as birth dates, government identification numbers and health information, are valuable on the black market. Because these data types are a prime target for theft or exploitation, they must be handled very carefully each step of the way.

For example, when IT teams are implementing new platform features or functionality, they can scramble data to block testers from viewing private data. Data scrambling uses real personal data to create realistic but fake datasets for testers to use. This enables the consistent and rigorous Quality Assurance needed to deploy new technologies while also limiting exposure.

Data masking also helps companies keep private data confidential. For instance, while an individual's manager needs to see their salary, every person on the HR team doesn't. Serving up private data on a strict need-to-know basis can protect employees' privacy and security, while also helping companies navigate a variety of local privacy laws.

"That's context-based security, which is really valuable when you work in multinational organisations," said O'Gara. "It's an entirely flexible model that allows you to be in full control of who has access to specific data."

People should also have control over how their data is being stored and processed – which requires companies to give users a peek behind the curtain.

Build employee trust – and enthusiasm – with responsible AI

AI, especially Generative AI, is only in its infancy, creating varying degrees of anticipation, uncertainty and wariness among employees. To build trust, companies will need to have honest conversations, clear policies and a dedication to training.

By **Emily Teesdale**, Senior Manager, KPMG

Mohammed Bari, Director, Powered HR, KPMG

Generative AI promises to usher in a new era of productivity – and companies that get it right could leave competitors in the dust. But what exactly does that look like? The truth is, no one really knows for sure.

This emerging technology, which uses advanced machine learning (ML) algorithms to generate entirely new content – from text and images to videos and presentations – is still finding its footing in the business world. Yet, business leaders understand the urgency of the situation: 80% of decision-makers agree that AI is required to keep their business competitive, a global Workday study found. If they wait until generative AI has gained traction to get on board, their organisation could quickly fall too far behind.

To stay ahead of the curve, companies are experimenting on many fronts. From automating complex tasks to brainstorming creative solutions, they're looking for new ways to enhance efficiency and accelerate innovation. The opportunities are exciting – but it's not yet clear what this technology will mean for employees.

Overall, CIOs expect increased productivity, increased collaboration and increased revenue and profits to be the top benefits that come from integrating AI and ML within the IT function, according to a report from Workday on AI in IT. However, only one-third of employees say they have a good understanding of AI and how it can be used in the workplace, a Forrester Consulting survey found.

Plus, there are several fundamental challenges that organisations must tussle with as they adopt generative AI. Ethical training, responsible use, robust governance and regulatory compliance are just a few of the critical factors CIOs must consider. Without the proper governance, AI can create more problems than it solves.

"ChatGPT and other generative AI tools will very confidently answer your questions, but it's based on the data they have access to. It's not always accurate," said Emily Teesdale, Senior Manager at KPMG. "So, a lot of companies are starting to draft policies and procedures to manage those inherent risks."

How companies tap, train and fine-tune generative AI tools could also have a major impact on both customer and employee trust. IT leaders must demonstrate that AI can be rolled out responsibly – protecting privacy, preserving jobs and producing accurate content – to get people on board.

For their part, employees are interested in learning more. Roughly three in four say they hope their company explores more AI implementation. But organisations need to strike the right balance between innovation and ethics to build enthusiasm about new ways of working. Otherwise, internal resistance to change could hinder meaningful progress.

Here are ways that CIOs can adopt and implement generative AI solutions that will both boost the bottom line and empower employees to drive responsible change.

Employees are interested in learning more. Roughly three in four say they hope their company explores more AI implementation.

Develop (and communicate) a clear AI strategy

Generative AI may seem to work like magic, but successful rollouts don't happen overnight. Getting farther faster with this technology demands a clear vision of what the organisation wants to achieve – whether that's boosting productivity, increasing customer satisfaction or improving the employee experience. From there, teams can start brainstorming different ways to meet those goals.

However, what they can achieve depends on the quality and quantity of the data AI models have access to. While some out-of-the-box solutions come pre-trained on relevant datasets, most models must be fine-tuned with proprietary data to deliver the most meaningful results. That means CIOs must focus on connecting internal data in a responsible way.

CIOs should also ensure the organisation's AI strategy keeps scalability top of mind, thinking through how new solutions will integrate with existing processes. The point is to improve results while also staying nimble, adopting technology that can adapt as both the business and AI applications evolve.

While proactive strategic planning is essential to make generative AI investments as effective as possible, that doesn't mean CIOs need a fully baked plan to get started, said Mohammed Bari, Director, Powered HR, at KPMG.

"You can have a strategy cooking while you're analysing your use cases," he said. "Don't wait, though. Go ahead, get started. Start thinking, start brainstorming and start experimenting."

Dip your toe in with specific use cases targeting pain points

Your generative AI strategy tells teams where they should be headed. Specific use cases show them which path they should take – and that extra direction can make all the difference.

"What we're seeing with AI is it's a bit more use-case-led," said Bari "So, I've got a big problem in recruitment. I've got a big problem in redeploying talent. Well, how can AI solve that?"

Take for instance, a company that receives thousands of CVs every day. It's impossible for an individual employee to sift through all of them – but generative AI can help bubble the best candidates to the top. By focusing on skills – what the business has, what it needs and what different applicants bring to the table – generative AI can quickly find the best fit. With an internal skills marketplace, organisations can also quickly find ideal candidates already in their ranks.

While the details of each use case will vary, focusing on major pain points can help companies get quick wins – and help teams learn how generative AI really works. When people start to apply this technology to their daily work, they'll identify potential uses that make their jobs easier. And as employees become personally invested in AI rollouts, larger productivity gains promise to appear.

> “
>
> ChatGPT and other generative AI tools will very confidently answer your questions, but it's based on the data they have access to. It's not always accurate.
>
> **Emily Teesdale,** Senior Manager, KPMG

Only 16% of leaders believe that employee buy-in is critical for AI success.

Prioritise ethics and governance from the ground up

AI models are trained on massive datasets – but that doesn't mean that data is always accurate. It could be biased, replicating the unconscious bias of its human trainers, or just plain wrong. Plus, there's the risk that training data has been manipulated by malicious actors, as these types of cyberattacks are becoming more valuable and therefore common.

In this environment, CIOs must develop a comprehensive risk management plan that addresses these potential threats to build and buy AI solutions that can be trusted. Establishing guidelines for responsible AI usage, data privacy and transparency shows employees that you care about ethical issues – and puts some of the power in their hands. While CIOs must lead by example, trustworthy AI takes commitment from everyone involved.

To get employees engaged, offer regular training on the company's ethics policy, relevant regulations and how to identify and address ethical issues. Encourage open communication by clearly outlining how employees can escalate concerns and establishing whistleblowing policies that protect employees from retaliation. Address bias before it becomes a problem by increasing diversity on teams developing and using new AI applications. Taking these steps early on will show employees that the company is taking its commitment to responsible AI seriously.

To get employees engaged, offer regular training on the company's ethics policy, relevant regulations and how to identify and address ethical issues.

Take the human impact into account

Generative AI could completely transform the business landscape – and employees aren't quite sure what that will mean for them. For example, will productivity enhancements translate to layoffs? Or will their duties change in a way that outpaces their skills?

These worries are realistic and CIOs need to take them seriously. If teams aren't willing to find new ways to work alongside generative AI, innovation will be stymied. Yet, only 16% of leaders believe that employee buy-in is critical for AI success.

To make employees active participants in AI innovation, take the time to show them where humans add the most value. Highlight the strategic, creative and logical tasks that require a discerning human mind. Then help them explore how generative AI can elevate their own roles by automating the mundane tasks no one wants to do.

Not everyone will be open to change, so find the people who are willing to adapt and invest in them as change management leaders. Skills are teachable, but attitude is not. In an uncertain environment, cultivating enthusiasm, curiosity and empathy will be key to future success, Bari said.

"The things that make us human – that make us good people and good colleagues – are the things that will make us stand out," he said.

About KPMG UK

KPMG LLP, a UK limited liability partnership, operates from 20 offices across the UK with approximately 18,000 partners and staff. It operates in 143 countries and territories with more than 273,000 partners and employees working in member firms around the world.

How companies can thrive with trusted AI

All the exact applications of AI are unknown, but organisations should consider privacy protections, human judgement and simpler systems to guide successful AI use.

By **Anja Fordon**,
EMEA Staff Writer

From writing marketing content to delivering faster, more accurate financial predictions to streamlining supply chains, business leaders see opportunities to win big by implementing AI across the board.

And no one wants to be left behind. In fact, nearly three-quarters of business leaders say they feel pressure to increase AI adoption, according to the 2023 AI IQ report from Workday. But where – and how – AI will deliver the biggest gains is still cloaked in uncertainty.

"What we would have done six months ago is not what we're doing today. What we're going to do in six months, we don't know today," said Shane Luke, Vice President, AI and Machine Learning (ML) at Workday.

To reach the glittering potential on the horizon, CIOs must lead their organisations through a minefield of hidden privacy, security, bias and ethics issues. Yet, the policies, rules and best practices that would normally guide them are still being developed. With so much still up in the air, it’s no surprise that nearly half (49%) of CEOs say their organisation is unprepared to adopt AI and ML, the C-Suite Global AI Indicator Report by Workday found.

“If people don't trust technology, they won’t use it,” said Tom Girdler, Principal, Product Marketing at Workday. “At the same time, if we can build technology that is underpinned by a robust framework, we create an amazing dynamic where trust and AI can really thrive together.”

By increasing confidence in AI, CIOs can amp up adoption and start delivering the exponential value business leaders expect. What will that take? It starts with adhering to three principles of trustworthy AI – and opening the door for teams to experiment responsibly with this transformative tech.

> “
> What we would have done six months ago is not what we're doing today. What we're going to do in six months, we don't know today.
>
> **Shane Luke,** Vice President, AI and Machine Learning (ML), Workday

1. Assess privacy risk – and plan for compliance

Not every AI application comes with the same level of risk. To navigate evolving privacy issues and compliance concerns, IT teams must understand the unique issues each use case presents – and help the business prioritise projects accordingly.

Some risks should be deemed unacceptable from the start, such as the presence of bias or discrimination in AI models that could lead to unfair or discriminatory outcomes concerning race, gender or other protected characteristics. Security vulnerabilities, unauthorised data collection and unreliable predictions are other examples of unacceptable risks that should stop an AI project in its tracks.

Other risks are more manageable, but require close oversight. For example, many AI models learn as they're used, which means new interactions could introduce new types of bias. IT teams must develop guidelines for the individuals providing new inputs – and monitor evolving outputs to ensure they remain fair and accurate.

CIOs must also keep an eye on compliance. While most privacy laws are still catching up with the latest advancements in AI, technology leaders should be prepared for what's to come. Regulations will most certainly vary across borders, but there is broad agreement that a few key factors are essential to the development and deployment of trustworthy AI.

"It's about transparency, technical documentation, recordkeeping, human oversight, accuracy, robustness and cybersecurity," said Jens-Henrik Jeppesen, Senior Director, Public Policy at Workday. "The idea is that technical standards are going to be developed to match each of these regulatory requirements and companies will certify to these standards."

2. Keep humans at the helm

Science fiction writers love to imagine dystopian futures ruled by sentient AI. Of course, technologists know that AI can't think – it can only come to conclusions based on its training data. But this can be dangerous in its own regard.

When unthinking machines make purely data-driven decisions, they often ignore crucial contextual factors. For example, an AI-driven financial model that relies on historical data to make projections may not account for current geopolitical conditions or recent shifts in market sentiment, which could significantly influence business outcomes.

For AI to inform sound business decisions, humans must remain involved each step of the way. From training and testing to implementation and adoption, organisations must use AI to amplify human potential – not the other way around.

"The real question is, how do you put that into practice?" asked Kelly Trindel, Chief Responsible AI Officer at Workday.

It takes cross-disciplinary, open-minded collaboration across different domains, she said. In these early days, CIOs must build the teams and organisational structures needed to develop the guidelines that will promote fairness, accuracy, reliability and robustness as the organisation brings new applications online.

77% of leaders worry that at least some of their data is neither timely or reliable enough to use with AI and ML.

"The people who actually know how this stuff works, they really need to be involved in how you put together your AI governance," Trindel said. "We're seeing it as a developing best practice to have separate lines of reporting for those who develop governance for AI systems and those who are frontline developers of AI systems."

3. Design simpler systems to mitigate bias

Bias in AI can't be completely avoided. Every human has their own opinions – and humans train AI based on what they believe to be true. However, proactively working to mitigate bias from the start can go a long way toward building more ethical and equitable AI systems.

"The design of the system is by far the most important," said Luke. "You can design the system to be very unlikely to produce something that you don't want. So that's the starting point."

Because training data will determine AI outputs, CIOs must ensure all applications are built using trustworthy data that has been examined and validated by diverse human teams. While testing outputs is important to mitigate bias that sneaks into the model, this should be the organisation's safety net – not its first line of defence, Luke said. "It's not about trying to check or police outputs. That's much harder to do and it's never definitive."

For example, large language models like Chat GPT are trained on large, general datasets that allow them to deliver long-form responses in convincing natural language. But these datasets often include poor content, such as misinformation found online. A substantial 77% of leaders worry that at least some of their data is neither timely or reliable enough to use with AI and ML. As an alternative, CIOs and their system designers should consider building applications with a smaller scope, trained to complete very specific tasks.

"They're not as capable at doing very general things, so they're less mesmerising," Luke said. "But they're very capable at the tasks they're supposed to do, while being less capable of doing something you don't want."

"

It's about transparency, technical documentation, recordkeeping, human oversight, accuracy, robustness and cybersecurity.

Jens-Henrik Jeppesen, Senior Director, Public Policy, Workday

To build trust in a time of data insecurity, embrace a proactive strategy

Mastering data privacy requires a commitment to continually mapping risks. The first step: embracing best practices and automated tools that do the heavy lifting.

By **Steve Dunne**,
EMEA Staff Writer

Data is the lifeblood of today's digital economy, but only if it can be trusted. While data fuels innovation and agility, it's also at the centre of growing security threats and regulatory requirements. The ability to strike the right balance between data protection and operational excellence is increasingly tied up with the ability to inspire trust among customers, employees and investors.

Many companies have experienced a year-over-year increase in cyberattacks, including 54% of European organisations, according to a September 2023 IDC report. That's one reason why data security now tops executive agendas across the continent. In fact, 45% of surveyed CEOs in Europe said they will prioritise spending on data security, risk and compliance to support trustworthy data collaboration and sharing. Respondents' top operational security priority? Data privacy and regulatory compliance.

Given all this, what CIOs now urgently seek is confidence that the company's critical business, people and financial data is not only compliant but secure from cyberthreats and internal bad actors. The fastest route to such confidence is now clear. Tech leaders need a proactive data privacy management strategy that leverages best practices and top-of-line technology, and deploys custom automation tools that bolster controls, monitoring and audits.

"Risks are growing along with the complexity and scale of data," said Mark Eaglefield, Head of Digital Products, Veolia UK, a global leader in environmental services that operates in nearly 50 countries. "Without a proactive stance, you can't quantify the organisation's level of risk – until after the damage is done and trust has been lost."

Protecting data – and building trust – by default

When it comes to data security and compliance vulnerabilities at multinational organisations, it can be hard to know where to start. In 2019, Veolia's tech leaders made a key decision that has grounded its data management strategy. The company became a full-platform Workday customer.

A unified tech environment spanning human capital management, finance, payroll, recruiting and other areas established a baseline of well-defined, documented and managed processes around user access and security, Eaglefield said. "That's our solid foundation, which we strengthen with best practices woven into our evolving, proactive data privacy strategy."

Leading among those practices is a steady drumbeat of education and awareness-building efforts around data privacy and user security. Veolia's IT organisation educates a range of stakeholders – such as end users, auditors and IT teams – via training sessions and published policies and procedures, through various communication channels. "It's a continuous cycle – education and awareness is key," Eaglefield said. "We never assume our stakeholders know and understand how we aim to protect data and what the stakes are."

Another best practice Veolia has embraced: forming and leaning on a team of dedicated, in-house security experts. These experts, deeply fluent in the Workday operating environment, collaborate closely with internal data protection teams. They're up-to-speed on current data privacy legislation and attendant regulatory requirements impacting the business. They act as peer reviewers in a way, helping to ensure that the organisation's policies, procedures and controls always reflect the current threat and regulatory landscape, Eaglefield said.

What CIOs now urgently seek is confidence that the company's critical business, people and financial data is not only compliant but secure from cyberthreats and internal bad actors.

45% of surveyed CEOs in Europe said they will prioritise spending on data security, risk and compliance to support trustworthy data collaboration and sharing.

“In terms of user-based security and compliance, these experts are crucial collaborators, allowing us to continually strengthen the design of our particular configuration,” he added.

The right tools for the job

Every business has a unique data environment to protect and related risks to guard against. At Veolia, there was growing awareness among security leaders of vulnerabilities related to proxy access.

The company had a proxy policy for its non-production environment, allowing users granted proxy access to see all the same data the individual being proxied typically sees. While just a small number of trusted individuals were granted such access, the lack of data masking still left Veolia open to potential data protection compliance and breach risks. The solution: Smart Shield, a tool designed by Kainos to enable data masking for specific proxy users in Workday.

“Now we can make sure that a user assigned to, say, a finance proxy group can’t view any compensation data relevant to the individual that they're proxying in as,” Eaglefield said.

At large organisations with thousands, or tens of thousands, of users, manually auditing user configurations relative to segregation of duties and system access levels is impossible. Audit automation must be part of any proactive data privacy solution – which is why many IT leaders rely on 360-degree security monitoring tools.

“The more complex your business is, the harder it is to get a panoramic view of data-related risks,” said Kim Freestone, Product Principal at Kainos. For example, “employees who have inappropriate and unrestricted access to highly sensitive data.”

Veolia, which has 14,000 users, chose to implement Kainos’ Smart Audit tool, which automates data security monitoring, including flagging business processes and data at high risk for fraud and breaches. There’s huge value in having an overarching viewpoint, in terms of segregation of duties and identifying related conflicts, Eaglefield notes. Preventative checks review whether users’ data access level is justified, and a daily digest email offers his internal controls team a helicopter view flagging anomalies, along with detailing current conflicts and what’s under review.

> “
> The more complex your business is, the harder it is to get a panoramic view of data-related risks.
>
> **Kim Freestone,** Product Principal, Kainos

> “
> We never assume our stakeholders know and understand how we aim to protect data and what the stakes are.
>
> **Mark Eaglefield,** Head of Digital Products, Veolia UK

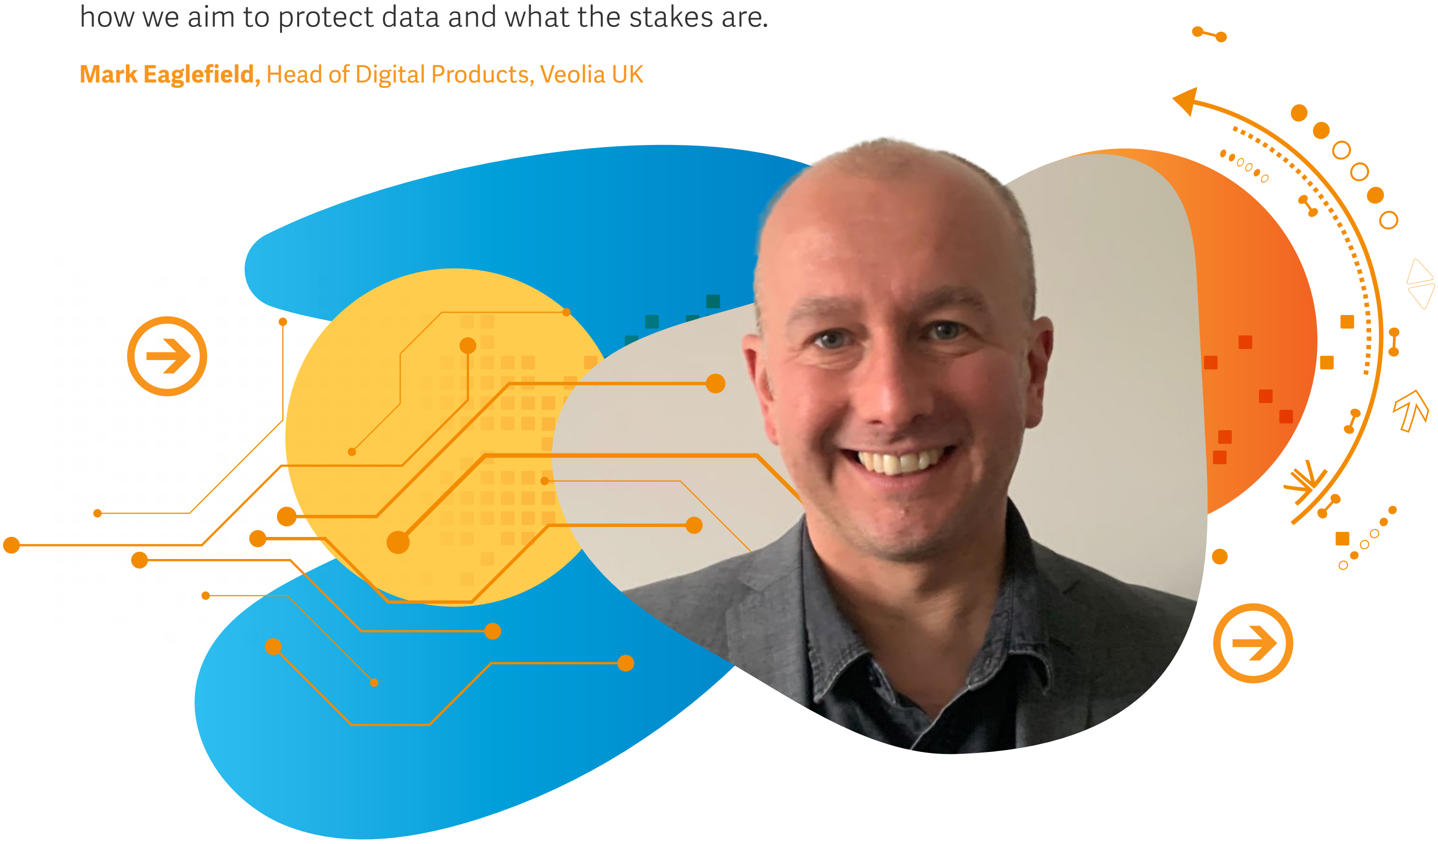

Proactive risk assessment involves more than just putting controls and processes to protect data in privileged access areas, Freestone says. “It’s also about assembling a body of evidence to show auditors, whether internal or external, that you’re serious about mitigating risks.”

Keep an eye on complacency

In an era of rising security threats and emergent AI capabilities, trust is earned in the realm of data privacy and protection. That will become even more true as organisations are challenged to both tap the power of AI while also complying with entirely new regulatory frameworks governing data sets and practices, such as the European Commission’s AI Act.

There is no such thing as absence of risk when it comes to networked data. But with a proactive data management strategy that embraces both best practices and the best of what current technology offers, CIOs can build and then fine-tune a future-ready IT infrastructure. Instead of stressing over unknown risks, leaders can find assurance in a highly configurable security model governing all of the enterprise’s apps and data. The right tools can surface risks and then prompt the right protective actions before it’s too late.

The biggest data misstep is complacency, Eaglefield said. “Don’t wait for a problem to arise – get ahead of it right now.”

Bring Your Own Key delivers data protection peace of mind

Stakeholders – including governments – have increasingly high data security expectations. Retaining full control of root encryption keys can build trust while easing regulatory compliance.

By **Anja Fordon**,
EMEA Staff Writer

Cloud computing has transformed how organisations do business – and how they think about security. The days of designing an IT security strategy around on-premises servers and databases are long gone. The shift to a cloud-first operating environment, along with the growing strategic value of data to businesses, necessitates a data-focused security approach.

But finding the right balance between cloud-based agility and data security is tricky for many organisations. Especially for businesses in highly regulated industries such as financial services, utilities, healthcare and those seeking compliance with more stringent data regulations such as the EU's GDPR. After all, working with a cloud service provider (CSP) typically means entrusting that partner with encryption services to protect users' data. This means no matter how secure the cloud environment, another organisation holds the encryption keys.

For particularly risk-averse organisations looking to strengthen trust with customers and eliminate compliance concerns, a new solution has emerged in recent years: Bring Your Own Key (BYOK).

Why BYOK

The basic value of BYOK is simple – it allows an organisation to encrypt their data in the cloud with their own root encryption key. Then, they can allow or deny access to the underlying data by sharing (or revoking) their root key with a SaaS provider. It's a single point of control for data access.

"With BYOK, you really have control over all of your organisation's data," said Tammo Buss, Workday Technical Lead at EWE, a German utility that delivers energy, telecommunications and IT services. "It's more than just an encryption service – it's really an encryption management service."

EWE had a very specific reason for implementing Workday BYOK: ensuring GDPR compliance. The organisation's data privacy officer and legal team decided that EWE needed full control over its own data in a cloud environment. A specific benefit of BYOK in EWE's highly regulated sector, Buss noted, relates to customer data audits.

"It's great that we have a product in which everything is very transparent, letting us streamline security auditing and the compliance response," Buss said. In December 2023, EWE went live on a variety of Workday solutions including Workday BYOK.

Beyond compliance audits, the big benefit of BYOK is total control over access to your data.

The basic value of BYOK is simple – it allows an organisation to encrypt their data in the cloud with their own root encryption key.

Taking full control

In general, BYOK can be implemented in more than one way. How an organisation manages its root encryption keys depends not only on its risk appetite and ability to manage keys in house, but on the CSP's underlying key management service.

As one example, a CSP may have the capability to use a root encryption key generated outside its system, but it needs the customer to upload the keys to the CSP's servers. This approach could enable a CSP to handle some encryption management duties, such as root key rotation, on behalf of the customer, thereby ceding control of root encryption back to the service provider and away from the customer.

EWE, however, was able to retain total control of all aspects of the root encryption key because Workday BYOK allows the customer to fully own and manage that key outside of Workday. EWE set up a customer-managed key management service in AWS, which EWE's Workday account then interfaced with.

"Through this customer-managed key approach, we were able to encrypt all our Workday data and all our tenants with our own key and have full control over it," Buss said.

AWS has certified that at no point can it access a customer-generated key, whether used in its key management system or in a hybrid security module – the cryptographic processing system that protects digital keys. That's important for proving compliance during System and Organisation Control (SOC) audits, Buss highlighted. "Our legal counsel was very happy with the documentation provided."

Greater responsibility – and real ROI

Workday BYOK can deliver clear value in terms of augmented data security and compliance peace of mind. But organisations considering implementing BYOK capability should think carefully about potential challenges.

Most obviously, BYOK involves assuming greater responsibilities, which may involve additional expenses. The organisation's internal IT team may need to take on managing an AWS account, and if an external partner needs to access the root key, the team would need to handle that process. An IT provider could handle some BYOK duties, but such delegation is at odds with the main purpose of BYOK.

And if the organisation loses the master key? That's a real problem.

"There is no backdoor," said Gautam Roy, Principal Product Manager at Workday, who has supported EWE's Workday account. "If the customer revokes access to the root key, Workday loses access to the data. That's the whole premise of BYOK. We don't do key backups. We need real-time access through the root keys to access data."

> "
>
> The cost of implementation and maintenance was actually quite low, from our perspective. It's easy to use and maintain, and we've significantly reduced data security risks.
>
> **Tammo Buss,** Workday Technical Lead, EWE

> “With BYOK, you really have control over all of your organisation’s data. It’s more than just an encryption service – it’s really an encryption management service.
>
> **Tammo Buss,** Workday Technical Lead, EWE

For EWE, the extra work and responsibility relative to BYOK was worth it. The organisation chose to eliminate GDPR compliance risk through a technical solution, rather than just reduce it through a data privacy contractual agreement. All European companies, and especially those in tightly regulated industries, must carefully address data privacy risks, Buss said, while remaining mindful that regulatory frameworks will keep evolving. With new EU AI regulation now on the way, organisations should prepare to be agile when it comes to data security practices.

And Workday BYOK can help in that respect. “We’re happy to have BYOK in place because when compliance requirements change, we know we’ll have technical control over our data,” Buss said. That should reduce the need to get legal teams together to update data privacy agreements with new contractual parts, he pointed out.

So while BYOK incurs costs, it can lower the cost of future compliance-related activities, while also helping to prevent expensive data security breaches.

“The cost of implementation and maintenance was actually quite low, from our perspective,” Buss said. “It’s easy to use and maintain, and we’ve significantly reduced data security risks.”

Beyond compliance audits, the big benefit of BYOK is total control over access to your data.

Get in touch

SUBMIT story ideas, **CONTRIBUTE** your own articles for publication or **SUBSCRIBE** to be a part of the smartCIO community and receive the quarterly digital issue and info on local events.

One address does it all!

Email us at **smartcioemea@workday.com**

workday.com

20240318-smartcio-volume10-magazine-engb