Out of the Shadows

How enterprise organizations are responding to the rise of shadow AI

OCTOBER 2024

All images in this report were created with the assistance of DALL·E 2

For enterprise organizations, shadow AI is a reality that can no longer be ignored

Over the past two years, the use of AI within enterprise organizations—and generative AI platforms in particular—has grown at an astonishing rate. According to a study conducted by PwC, 73% of US companies with revenues in excess of $500 million are now using AI in at least some areas of their business,[1] while over half have taken steps to incorporate generative AI specifically into their day-to-day operations.

But alongside the rapid adoption of AI has come the emergence of a troubling new phenomenon: ***shadow AI***. From individual employees using their personal ChatGPT accounts through to entire departments going rogue and deploying unsanctioned AI products, security and compliance teams have faced an uphill battle when it comes to ensuring that AI use cases within their organizations fall within the scope of official IT governance.

The proliferation of shadow AI creates undeniable risks for enterprise organizations—including entirely new categories of security vulnerabilities. At the same time, however, many leadership teams are reluctant to take too heavy-handed of an approach to shadow AI, lest they risk stifling organic experimentation and grassroots innovation within their organizations. Indeed, many AI use cases began as examples of shadow AI before eventually being legitimized and turned into accepted new business processes and incorporated into the organization's IT architecture.

Through interviews with enterprise decision-makers across a variety of industries, this report uncovers the key drivers behind shadow AI and the varied ways businesses are attempting to thread the needle of minimizing AI risk while still allowing employees to innovate. By doing so, it provides actionable recommendations for providers of AI solutions, demonstrating how they can work with their customers to bring shadow AI out of the darkness and into the light—transforming a potential risk into a strategic opportunity.

METHODOLOGY

For this report, NRG conducted case study interviews with key decision-makers across 10 US-based enterprise-level organizations (i.e., businesses with a global headcount of at least 500). Interviewees spanned a range of roles—including department heads, CTOs, HR directors, and innovation leads, to name a few—but all of them had personal involvement in setting the policies and standards governing approved AI usage within their organizations. The organizations profiled for this research represent a broad range of sectors, including but not limited to technology, professional services, aviation, real estate, and financial services.

Additionally, where applicable, we have drawn on insights from existing NRG research on public attitudes towards AI. See page 16 for links to previous papers on this topic.

IN THIS PAPER, YOU'LL FIND...

[1] "2024 AI Business Predictions," PwC

Shadow AI, by its nature, is extremely difficult to define

In theory, “shadow AI” is easy to define. As an extension of the broader concept of “shadow IT,” the term refers to **the unsanctioned use of AI by employees or freelancers within an organization**—i.e., the use of AI products or platforms without oversight or approval from an organization’s IT department.

This definition encompasses a broad range of use cases, from the relatively benign—e.g., an employee using their personal ChatGPT account to help draft an internal email—through to those that have substantial security implications—e.g., uploading sensitive customer or client data to a third-party platform.

> **It’d be naive for anyone to say that 100% of the AI use within their organization is explicitly sanctioned by the relevant policies. I like to think we’ve been able to get it to around 95%, but no matter what we do, we have to accept there will always be that 5% of individuals going off and using their own tools for their own needs and use cases."**
>
> CTO, technology

Often, this activity is truly “bottom-up,” the result of individuals choosing to experiment with AI products to find workarounds to existing processes or to increase their personal efficiency. But there are also cases where shadow AI is the result of concerted activity across an entire team or department, with team leads actively encouraging people to use specific AI products as part of day-to-day operations. In larger organizations, this can even include the development and rollout of custom AI tools without the awareness of or authorization by the relevant central compliance and security stakeholders.

EXAMPLES OF SHADOW AI

Relatively low risk

INDIVIDUAL LEVEL	TEAM LEVEL
Drafting or refining internal emails using chatbots or AI-based text prediction and grammar editing tools	Teams using AI for internal document managment purposes, helping members find information more efficiently
Generating reports, presentations, or summaries using AI	Teams using AI to automate project management tasks such as resource allocation or scheduling
Using AI-powered translation services for external comms	Teams using AI-driven collaborative platforms for writing reports or drafting documents
Creating copy or imagery for social media posts	Sales teams building AI models for forecasting sales and inventories
Artists and designers using AI imagery in marketing materials and sales collateral	HR teams using AI to screen resumes
Developers using AI tools to write or debug code	Proposal or sales teams using AI tools to draft proposals or respond to RFPs
Using AI tools to summarize or review contracts or legal documents	Product teams developing AI-based text analysis tools to extract insights from customer queries
Uploading sensitive customer data to AI platforms to perform analysis and data-cleaning	Customer service teams developing their own AI chatbots for handling customer queries

Relatively high risk

While this definition may be easy to grasp in theory, applying it in practice can be exceedingly difficult—impossible, even, in some circumstances. The commercial AI landscape has progressed rapidly over the past two years. The IT policies and governance practices of many large organizations, meanwhile, have struggled to keep pace. The result is a situation in which even those tasked with monitoring and curbing shadow AI often struggle to determine whether or not a specific AI platform or use case falls within the acceptable parameters set by their organization's official policies.

This complexity is compounded by the fact that many large organizations have made a conscious or semi-conscious decision to wilfully turn a blind eye to certain cases of shadow AI out of a desire to avoid stifling innovation or being seen as behind-the-curve when it comes to this nascent technology. In some cases, IT teams have even formalized this arrangement by helping individual departments set up parallel data environments explicitly exempted from their organization's standard governance practices.

The determination as to whether a specific use case for AI ought to be viewed as "shadow AI" or "approved AI," therefore, is often a deeply subjective exercise—one that depends greatly on how individual stakeholders choose to interpret and understand policies and guidelines that are frequently non-specific, terminologically loose, or outdated.

It may be more accurate, therefore, to think of these categories not as two sides of the same coin, but as ends along a continuum. Yes, there are certainly some scenarios in which specific AI use cases have been either clearly authorized or clearly forbidden by the relevant policies. But many—perhaps even the majority—of contemporary AI use cases in enterprise organizations fall somewhere in between these two extremes: a nebulous category of "gray AI" that exists in the liminal space created by the failure of institutional governance to keep pace with a fast-changing technology landscape.

> **I'm positive that people are using shadow AI in our organization. It's impossible to quantify it, but I'd be shocked if it wasn't happening. Everyone uses their own personal phones or laptops for work purposes occasionally anyway; if they want to access an unapproved AI tool, it's trivially easy for them to do so."**
>
> HR director, real estate

THE SPECTRUM OF SHADOW AI

SHADOW AI

AI use cases clearly and explicitly **ruled out** by official IT policies

No ability for IT team to accurately track and audit usage

Relevant compliance and security teams actively trying to push back on usage

GRAY AI

AI use cases neither clearly authorized nor clearly ruled out by official IT policies

May have tacit approval of IT teams or senior stakeholders but without proper documentation, leading to ambiguities and a lack of accountability

May take place in parallel data environments operated by individual teams

APPROVED AI

AI use cases **clearly authorized** by official IT policies

Usage can be accurately monitored and audited, and takes place within well-defined guardrails

Platforms clearly approved by any relevant compliance or cybersecurity teams

As AI platforms become more sophisticated, more employees will feel compelled to engage in shadow AI

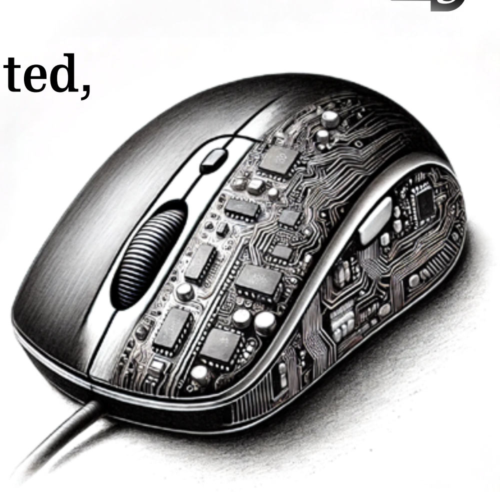

Without exception, the decision-makers interviewed for this report recognized that some degree of unauthorized AI usage within their organizations was inevitable, given the exceptionally low barrier to entry for many popular AI platforms. Even organizations that operate heavily restricted and monitored technology estates have to accept that there's little to nothing they can do to prevent employees from using AI platforms on their personal devices.

Equally, most senior IT and compliance stakeholders recognize that the vast majority of employees who engage in these forms of unauthorized AI usage are not doing so with malicious intent. In many cases, they may be unaware of the potential security risks they are creating by doing so. Particularly for many younger workers, conversational AI platforms like ChatGPT and image generators such as Dall-E are almost as mundane and commonplace as word processors or spreadsheets: hardly the sort of thing one would think might require special approval to make use of in a professional context.

And even where employees are aware that they're breaching their organization's governance policies by uploading sensitive information to unauthorized platforms, they may feel that the risks are small enough to be worth the potential rewards—especially if they're under pressure from their team leaders to hit revenue targets or to cut costs.

Previous NRG research on the implications of AI for creative class professionals found that **many creative workers are deeply concerned about the prospect that AI tools will create a "productivity arms race"**: as more and more of their colleagues embrace these kinds of products, employers and clients will develop heightened expectations for productivity, forcing them to embed AI even deeper into their workflows. It's possible that many employees—in creative positions or otherwise—already feel compelled to turn to AI platforms simply to keep pace with colleagues who are doing so.

> **I don't think that the use of shadow AI is typically motivated by malicious or nefarious intentions. In most cases, it just means that the policies aren't properly understood or communicated; or maybe the policy is just more strict than it really needs to be."**
>
> COO, software

The factors that compel employees to engage in shadow or gray AI can exist at every level of an organization—from the most junior of interns up to the C-suite. Several of the decision-makers interviewed for this report acknowledged that they themselves had made unauthorized use of certain AI products, despite being personally responsible for drafting and implementing policies towards the technology.

Equally, these factors are not limited to any one specific department or vertical. That said, organizations do not have limitless resources; when it comes to shedding light on or discouraging shadow AI, they will find themselves forced to make tactical decisions about where to prioritize those efforts.

To that end, several of the decision-makers NRG spoke to mentioned they were most concerned about shadow AI as practiced by their most tech-literate employees: including developers, engineers, and data scientists, for example. These employees, they reasoned, were far-and-away the most likely to be using shadow AI in ways that could create significant data vulnerabilities. But these employees are also the ones best able to cover their tracks to avoid detection, and the ones best equipped to find workarounds to whatever technical limitations their organization has put in place to deter unauthorized usage of AI products.

> **It is typically the intellectually curious employees who are the most likely to use shadow AI: the people who have a real desire to be innovative or be at the leading edge of their field."**
>
> CEO, technology

REASONS EMPLOYEES MAY USE SHADOW AI

Pressure to innovate

Individual managers may push employees to stay abreast of new technologies and trial new use cases, regardless of wider organizational culture or policies.

Need to keep pace

If colleagues are using AI tools to increase their productive output, individuals may feel they have little choice but to adopt these tools themselves to keep up.

Limited understanding of policies

IT policies governing AI usage may be ambiguous, hard to find, or out-of-date. If changes were implemented recently to define approved AI systems, these may not have been communicated well across the business.

Personal familiarity with AI tools

Employees will have begun to integrate AI tools such as ChatGPT or Microsoft Copilot into their personal lives; they may also have prior experience of using such tools in a professional capacity from previous jobs.

Perceived lack of consequences

If shadow AI usage is widespread, individual employees may assume they're unlikely to face personal repercussions for it. And employees without a background in cybersecurity may not fully appreciate the risks associated with unapproved AI applications.

Frustration with approved tools and processes

Use of shadow AI may reflect underlying dissatisfaction with an organization's existing suite of approved tools.

Drive for efficiency

Teams that are overworked and overstretched may turn to AI tools as an easy way to automate mundane tasks, freeing up capacity to focus on more complex activities.

External collaboration

Teams may default to using AI tools that align with those used by external partners, vendors, or clients, without first checking to see if they have been approved by internal IT policies.

> **"Our biggest offender of shadow AI use is the marketing team. That's because they are in charge of the digital experience and therefore feel they can freely incorporate AI tools into how they manage analytics for it. Sales is also highly prone to using shadow AI; in our industry, there's a consensus that, if someone wins a client, they're entitled to ownership of their data and key information, and can use whatever tools they want to help manage it."**
>
> CIO, financial services

> **"Developers and other employees with technical skill sets are the most likely to use shadow AI—or at least, the most likely to use it at a scale where it creates clear risks. Those employees not only understand the AI landscape, they also have the knowledge necessary to get around whatever failsafes or technical controls you try to set up."**
>
> Senior HR leader, consulting

While the cybersecurity implications of shadow AI are well known, these are far from the only relevant risks

One risk of shadow AI is the potential exposure of enterprise data to the real world. Then there are privacy concerns relating to employee data being made available involuntarily. And on top of that, you have a lack of clear security protocols for some of these applications. Really, the entire ecosystem just opens up an enormous set of hacking concerns."

CIO, consulting

We deal with a lot of third-party customer data on behalf of our advertisers. So for us, the biggest risk of shadow AI is that someone uploads sensitive data into one of these platforms and it ends up getting leaked. I'm also worried about internal communications being exposed, now that everyone's using AI tools to help them summarize and respond to emails."

Chief transformation officer, advertising

As a healthcare provider, we have to abide by incredibly strict rules around patient privacy and the security of medical records. So that's the biggest risk I'm thinking about when it comes to shadow AI."

COO, healthcare

Most organizations are all too aware of the cybersecurity risks associated with shadow AI. There have been plenty of headlines in recent months about hacks[2] and data leakages[3] in prominent AI platforms—highlighting, for enterprise organizations, the importance of ensuring that AI platforms in use within their technology estates have been properly vetted and audited. While the same could be said for any tech platform capable of handling large volumes of potentially sensitive data, this problem is particularly acute in the case of AI. Many LLMs are still, essentially, black boxes; businesses may have little to no idea about how the data their employees upload to popular AI platforms is stored or whether it's being used to train subsequent iterations of the model.[4]

These security concerns are particularly relevant for organizations which operate in heavily regulated industries, such as aerospace or natural resources, and those which regularly deal with highly sensitive customer data, such as financial or healthcare records—as well as those which operate in markets like the EU which enforce particularly stringent data protection standards.

Comparatively less attention, however, has been paid to the broader set of brand and reputational challenges that can result from the proliferation of shadow AI within an organization.

In some cases, reputation and security concerns go hand-in-hand: a significant data breach, of course, can often lead to negative media scrutiny and/or backlash from customers, investors, or other stakeholders. But there are also plenty of other ways in which AI, when not subjected to the proper level of oversight and scrutiny, can damage a brand's standing in the eyes of customers and the wider public.

Several brands, for example, are grappling with the fallout of inaccurate advice provided to customers by AI chatbots; a case earlier this year against Air Canada found that the brand was legally liable for decisions consumers had made on the basis of faulty information given to them by the brand's AI customer service agent.[5]

In other cases, brands—particularly those in the entertainment and media space—have had to deal with backlash from creative professionals online due to their use of AI-generated images, text, videos, or other forms of media. Gaming brand Wizards of the Coast, for example, issued a blanket ban on the use of AI within its titles following controversy among fans in 2023—but then had to apologize to fans again in 2024 after they noticed AI-generated content in advertising images the company had commissioned from a vendor.[6] This case highlights the particular difficulties in eliminating shadow AI for brands that rely heavily on freelancers, agencies, and other third parties.

[2] Kevin Townsend, "Hacker Stole Secrets From OpenAI," Security Week, July 5th, 2024
[3] George Fitzmaurice, "Microsoft Copilot could have serious vulnerabilities after researchers reveal data leak issues in RAG systems," IT Pro, August 19th, 2024
[4] The editorial board, "AI should not be a black box," Financial Times, May 30th, 2024
[5] Leyland Cecco, "Air Canada ordered to pay customer who was misled by airline's chatbot," The Guardian, February 16th, 2024
[6] Oli Welsh, "Wizards of the Coast admits using AI art after banning AI art," Polygon, January 8th, 2024

SECURITY & DATA RISKS

DATA BREACHES

Unauthorized AI tools may not have proper security protocols in place, increasing the likelihood of exposing sensitive corporate data.

COMPLIANCE VIOLATIONS

Use of unsanctioned tools can lead to violations of country or industry-wide data privacy regulations.

LACK OF ENCRYPTION

AI tools may lack encryption, leaving data transmitted or stored in these tools unprotected from interception or theft.

PHISHING AND SOCIAL ENGINEERING ATTACKS

Employees may unknowingly engage with malicious AI tools that pose as legitimate applications.

DATA RESIDENCY ISSUES

AI tools might store data in jurisdictions with different privacy protections, creating risks around international compliance.

LIMITED TRANSPARENCY

AI providers may not provide clear information on how data submitted to their platforms is stored and whether it's used for training LLMs.

INADEQUATE AUTHENTICATION

AI tools may not enforce multi-factor authentication, making them susceptible to unauthorized logins.

INSUFFICIENT BACKUP AND RECOVERY

Shadow AI applications may not be included in the organization's data backup plans, making it impossible to recover lost data in case of failure or attack.

INCREASED ATTACK SURFACE

Unauthorized AI tools broaden the organization's attack surface, providing more entry points for hackers to exploit.

INABILITY TO AUDIT

IT departments may be unable to perform regular security audits on unsanctioned AI tools, making it difficult to track data flows, vulnerabilities, or access logs.

REPUTATIONAL RISKS

INACCURATE AI OUTPUTS

AI-generated errors that go unchecked due to lack of oversight can lead to incorrect business decisions, negatively affecting customer experiences or partnerships.

NEGATIVE MEDIA ATTENTION

High-profile security breaches or errors attributed to AI may lead to media scrutiny and backlash from employees, investors, or other stakeholders.

BRAND DAMAGE

Low quality or inaccurate AI-generated text or imagery featured in sales and marketing collateral may damage the organization's brand equity.

INCONSISTENT CUSTOMER EXPERIENCE

If AI tools are being used to communicate with customers, the organization may lose the ability to provide a consistently positive customer experience.

LABOR BACKLASH

If an organization is seen to be using AI tools in a way that is insensitive to the concerns of labor groups, this could lead to public backlash or organized labor action.

SCRUTINY OF BIAS AND DISCRIMINATION

Use of unapproved AI tools can lead to biased outputs (e.g., discriminatory hiring or lending practices), leading to public backlash and accusations of unethical practices.

These challenges have given rise to two distinct philosophies for managing shadow AI

However, the risks associated with the proliferation of shadow AI are only one part of the equation that leaders within enterprise organizations need to consider when deciding on the appropriate strategy for dealing with shadow AI. Many such leaders are also keenly aware of the risks of being *too* restrictive in their approach to AI. Particularly for businesses in sectors prone to rapid disruption, it's certainly possible that the opportunity costs of being overly cautious may outweigh the security and reputational risks that stem from employees experimenting with AI applications beyond the purview of the relevant governance policies.

The very existence of shadow AI, after all, speaks to a level of grassroots enthusiasm towards the technology among large parts of the workforce—an enthusiasm that many organizations would be keen to harness in service of process and product innovation. As a result, companies find themselves walking a tightrope: **trying to balance the need to innovate and stay ahead of the competition in terms of efficiency and product quality against the desire to avoid serious security incidents and significant reputational blowback**.

For the most part, large organizations appear to have adopted one of two mutually exclusive strategies for solving this dilemma. Businesses that prioritize safety over innovation—whether due to the industry that they operate in or unique cultural factors within their organization—have tended to embrace a RESTRICT BY DEFAULT model. For these types of companies, the messaging to employees is simple: assume that any given AI product is off-limits for professional use, unless it has been specifically authorized for use by the relevant IT stakeholders or you're willing to personally submit it to a lengthy approval and vetting process.

This model benefits from a great deal of relative clarity and simplicity. By limiting the scope for shadow AI and gray AI as much as possible, the "restrict by default" model helps to promote clear lines of accountability; employees, regardless of their level of seniority, don't have the luxury of being able to hide behind policy ambiguity.

There are, however, a number of downsides to this approach. In many cases, IT teams have been tasked by their organization's leadership teams with finding "AI wins." They may justifiably be concerned, therefore, about the political risks of adopting an overly draconian approach to shadow AI that prevents potentially impactful use cases from emerging and taking root.

Moreover, if employees are just going to be able to work around restrictions on AI usage by using their personal devices, then there's a risk that such an approach will not only prove ineffectual, but will also undermine the authority and credibility of a business's IT function.

> As a consultant, I've seen a lot of different approaches to managing shadow AI, and the broader question of shadow IT in general. In some cases, I've even seen organizations actively try to foster it, on the basis that it's better to leverage the ground-up knowledge within a company than try to impose a one-size-fits-all policy solution."
>
> CIO, consulting

> We are taking a very diplomatic approach to the use of external technologies by our employees. On the one hand, we don't want our IT team to lose credibility or authority. But equally, we don't want to stifle innovation and experimentation. So we're being deliberately soft in our enforcement of our own policies. We know this is going to be an educational journey for all of us within the business."
>
> CIO, financial services

> We aren't worried about the use of shadow AI tools, not at all. We actually find that it gives us a boost to our reputation; we want potential employees to look at us as a business that's open to experimentation, one that's using all the modern, sexy, slick applications."
>
> COO, software

> I don't want to squash the innovation that could emerge from employees playing around with AI. So long as people know what the guardrails are, I think they should have as much freedom to experiment as possible."
>
> CIO, aviation

> “We take a very binary approach to AI: we tell employees to assume that they’ll need special dispensation if they want to use any sort of AI platform for work purposes. And that dispensation is only given out in very particular circumstances. We don’t want our policies to be subject to interpretation by employees; we find that a blanket ban is the best way to avoid any confusion or unanswered questions.”
>
> CEO, technology

Conversely, organizations that are more concerned about the risks of not finding the right AI use cases are often more inclined to adopt a more permissive, APPROVE BY DEFAULT model. Within this model, employees may be encouraged to use certain tools or to abide by certain data management principles, but are generally given the freedom to road test new use cases for AI on their own terms. Unless they’ve been specifically told otherwise by the relevant stakeholders, they can assume that any platform or any specific use case is fair game for experimentation.

What this model lacks in clarity, it makes up for in agility; organizations that embrace this approach find themselves able to more rapidly adapt their operating models around the new possibilities opened up by advancements in artificial intelligence. To make this approach work at scale, **IT teams need to function less as police officers, and more as true creative partners**—positioning themselves as the nexus for AI innovation within the organization, providing employees across the business with the tools they need to experiment with AI in meaningful ways, and shining a light on innovations that others ought to know about.

We can see in these two different philosophies two competing visions for how to solve the problem of shadow AI.

Organizations that adopt the RESTRICT BY DEFAULT approach seek to address the problem by stamping it out—making it clear to employees which use cases are off limits, putting measures in place to prevent them from engaging in those use cases, and taking remediative steps when those measures are circumvented.

Organizations with a more permissive approach to the technology, meanwhile, will seek to bring shadow AI out of the darkness and into the sunlight, providing IT and compliance stakeholders with better visibility and “legitimizing” those use cases that warrant it.

TWO COMPETING PHILOSOPHIES FOR DEALING WITH SHADOW AI

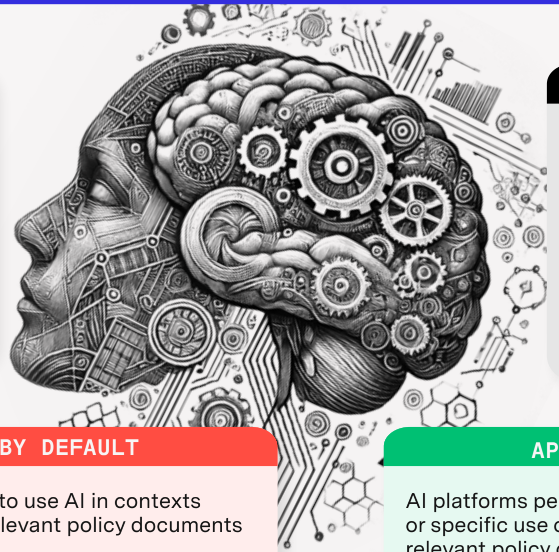

DECISION DRIVERS

- Organization operates in a highly regulated environment (e.g., banking, aviation, or healthcare)
- Organization has broadly conservative culture
- Organization's IT architecture acts as a tightly controlled "walled garden"

RESTRICT BY DEFAULT

- Employees only allowed to use AI in contexts specifically laid out by relevant policy documents
- Potential AI use cases identified by teams subject to extensive—and typically lengthy—screening and approval process
- IT and compliance teams monitor for misuse and crack down on shadow AI

DECISION DRIVERS

- Organization operates in a market prone to rapid disruption (e.g., software) where threat of falling behind the competition is a more pressing danger than the risk of shadow AI
- Organization has an established culture of innovation
- Organization's IT architecture is inherently flexible, with employees working from a range of machines, including their personal devices

APPROVE BY DEFAULT

- AI platforms permitted unless the specific platform or specific use case has been explicitly ruled out by relevant policy documents
- Teams encouraged to experiment with AI and share their learnings across the business in a "test and learn" model
- IT and compliance teams work with stakeholders across the business to facilitate experimentation

If someone wanted to use a separate AI tool other than our preferred ones, they would be allowed to, because our general approach with any business tools, including AI ones, is that it is overall recommended. If someone identifies a product that's a better fit for a given use case, then we would simply go through our normal process of conducting a third-party vendor assessment to get the vendor vetted and approved."

COO, software

Our IT policy is very clear: employees are prohibited from using AI to conduct operations or process business data, unless such AI is the company's own internally-approved AI system or a system a client has requested us to use."

HR director, real estate

Our general philosophy is that employees understand their own jobs better than anyone else, and know best what kinds of applications they need access to. We've had numerous examples in the past where we found employees using tools that weren't officially sanctioned, so we ran them through our normal vendor assessment process to get them vetted and onto our approved list."

COO, software

We don't have any specifically authorized AI tools within our business other than a Microsoft Copilot license. But people are free to use other tools if they find them useful. Some people are always going to be more comfortable with specific applications, so we don't want to force conformity if we don't have to."

CIO, financial services

These macro-philosophies will inevitably inform the specific lower-level controls which organizations put in place in order to monitor and manage shadow AI. In general, these controls can be divided into three distinct categories.

Technology controls physically restrict employees from accessing high-risk platforms, or provide IT teams with direct visibility of how AI tools are being used and misused.

Process controls, meanwhile, set out the actions that various stakeholders need to take in order to properly ascertain whether a specific AI use case is appropriate and to mitigate potential risks.

Finally, there are **cultural controls**: the shared values and principles within an organization that inform how employees approach this technology and how they understand their relationship with the compliance and security stakeholders tasked with responding to the challenges raised by shadow AI.

Organizations with a more restrictive attitude towards AI may be inclined to lean more heavily on technology controls, taking steps to make it as difficult as possible for employees to use AI platforms in unauthorized ways on their work devices.

On the other end of the spectrum, organizations with a more permissive approach to AI will typically spend more time worrying about cultural controls—trying to promote values of accountability, transparency, and innovation that can help employees make informed and intelligent decisions about how and when to use AI platforms.

For this culture-driven approach to work, it's critical that these values are embraced from the top down, with senior leaders educating themselves about AI, assuming executive sponsorship for AI initiatives, and exemplifying the kinds of risk-conscious yet innovative behaviors they want to see normalized across their organizations. It's no coincidence that—according to recent research conducted by NRG and Google—comprehensive C-level sponsorship is one of the critical factors determining whether businesses are able to drive measurable and meaningful value out of their AI investments.[7]

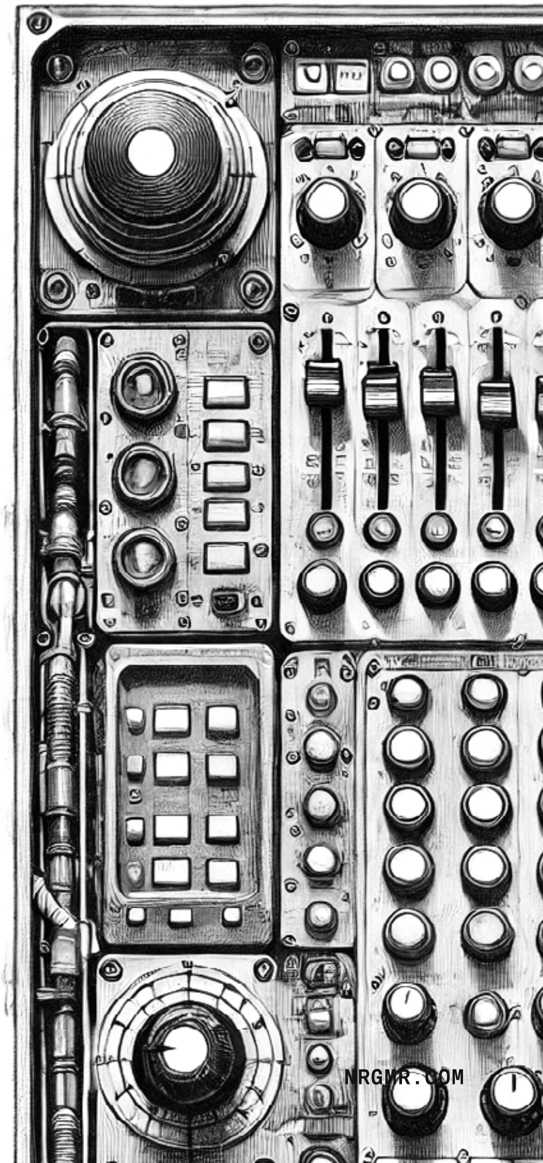

[7] "The ROI of gen AI," Google Cloud, August, 2024

CONTROLS FOR SHADOW AI

> "We have a highly restricted IT architecture, all new applications on employee devices have to be vetted by IT, and usage of online platforms is monitored as well. But that hasn't been enough to stop shadow AI completely; employees will always be able to use their own devices, there's nothing we can really do if they're using their personal ChatGPT account at home to complete tasks."
>
> CEO, technology

TECHNOLOGY CONTROLS

- Use role-based access control (RBAC) to restrict who can install or access certain AI tools
- Create sandbox environments where employees can experiment with AI tools without posing risks to the main network
- Implement software whitelisting to allow only approved applications and AI tools on company devices
- Adopt specialized governance tools that monitor AI use across departments, ensuring tools meet compliance and security standards
- Use data loss prevention (DLP) tools to monitor and prevent unauthorized sharing of sensitive data through AI platforms

> "There are plenty of ad hoc cases that have emerged for AI within our business, cases where people have said 'Hey, this is a tool that I've found to be effective for this task, can I use it?'. So we have a clear governance process set up to handle those requests, with a designated committee who approves or rejects them."
>
> CEO, software

PROCESS CONTROLS

- Create clear and succinct policy documents that are easily accessible to all employees
- Implement a well-defined process for reviewing potential AI use cases, including a risk assessment framework
- Conduct regular IT audits to discover shadow AI and identify non-compliant tools and processes
- Define and rehearse incident response procedures for cases where unauthorized AI tools have been detected, including remediation and communication protocols
- Provide training that helps employees understand the risks associated with unauthorized AI tools and how to properly request and use AI resources

> "To us, the question of enforcing responsible AI use is a cultural one. We want employees to understand that you're still accountable for your own decisions, regardless of what an AI may or may not have told you to do. Equally, we want to make sure all employees are culturally invested in maintaining compliance with GDPR and other relevant standards, regardless of this technology shift."
>
> CIO, aviation

CULTURAL CONTROLS

- Foster a culture of open communication where employees feel comfortable discussing the tools they use with IT and leadership without fear of punishment
- Identify and empower "AI champions" in different departments to act as liaisons between the IT department and individual teams
- Reward teams or individuals who report the use of shadow AI tools, framing it as an opportunity for improvement and learning
- Make authorized AI tools easier to access and create incentives for using them in innovative ways
- Ensure that leadership emphasizes the importance of following AI governance policies and sets the tone for ethical and compliant use of AI tools
- Involve employees and teams across the business in the creation of AI policies and strategies, allowing them to feel ownership and fostering a culture of compliance
- Promote interdepartmental collaboration, with IT acting as a hub for innovation and knowledge-sharing

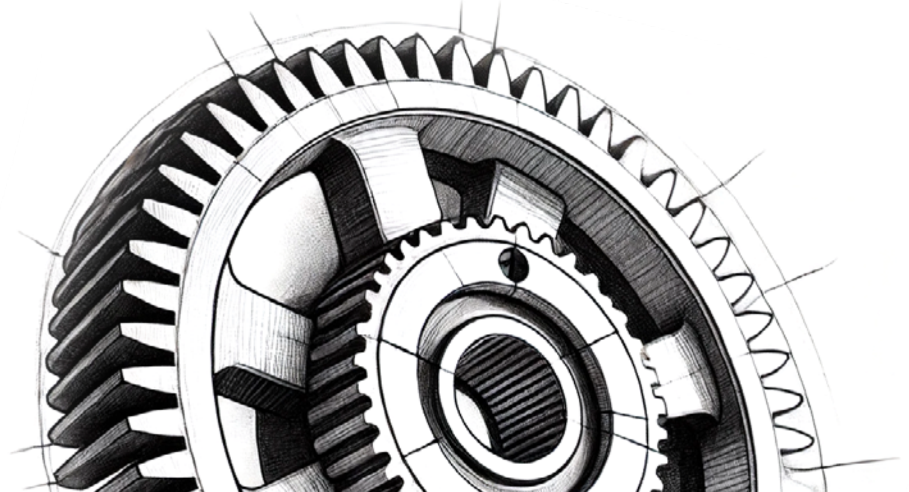

Many organizations are keen to move towards a more permissive model of AI management—but find themselves stifled by factors outside their control

As the AI landscape continues to evolve, the risk/reward calculus that dictates how organizations grapple with the challenge of shadow AI will inevitably shift. Indeed, many organizations that currently practice the RESTRICT BY DEFAULT philosophy would like to move towards a more permissive model that allows them to more effectively harness the potential of this new technology. It's just that the circumstances aren't quite right yet for them to take that leap into the unknown.

This may be the result of internal stakeholder challenges within an organization—but in many cases, it has just as much, if not more, to do with the state of the AI market itself. For one thing, **key decision-makers within large organizations often find themselves constrained by the lack of legal and regulatory clarity around the technology**. Security and compliance leaders want to know: where legal liability sits in the event of something going wrong; the legal status of the training data used in LLMs and other generative AI models; and how AI intersects with established compliance and reporting requirements.[8] Until these concerns are addressed, such stakeholders will inevitably feel uneasy about giving the green light for their organizations to embrace a more open and permissive model of AI management.

Equally, many C-suite leaders and other key stakeholders have concerns about the business models and operational practices underpinning today's most popular AI products. They may, for example, feel uncomfortable authorizing employees to upload sensitive data into platforms whose owners have, in many cases, been less than fully transparent about their own data management practices or the sourcing of training data for their models.[9]

Similarly, they may be reluctant to encourage usage of products where the pathway to profitability remains unclear. In recent months, media coverage of OpenAI's high burn-rate has prompted **speculation about the long-term financial viability of popular generative AI products**; many analysts expect that we will soon see OpenAI and its competitors move more of their models' features behind paywalls, or increase the rates charged to corporate users.[10] Businesses may understandably feel wary about embedding these kinds of technologies into their operating models as long as there's a substantial risk of hefty price rises in the near-to-mid term future.

> “
>
> **The thing that scares me about all these AI companies is: what's the pathway to profitability? You see all these companies burning through capital, you have to assume at some point they're going to find a way to pass the costs onto their customers. I'd love to encourage a more permissive attitude towards AI within our company, but not if it means that we'll end up dependent on technologies that are going to exponentially increase in price within a few years.”**
>
> CEO, software

Until these kinds of concerns have been adequately addressed, many large organizations will feel that they have little choice but to adopt an aggressive approach towards identifying and quashing instances of shadow AI—instead of a more permissive model that seeks to empower employees and harness their enthusiasm for the technology to more productive ends.

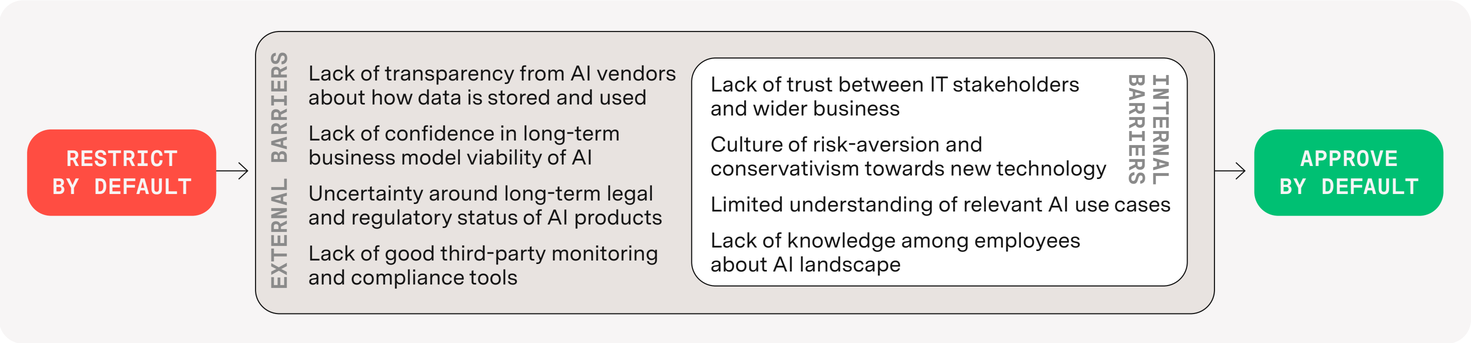
RESTRICT BY DEFAULT
EXTERNAL BARRIERS
Lack of transparency from AI vendors about how data is stored and used
Lack of confidence in long-term business model viability of AI
Uncertainty around long-term legal and regulatory status of AI products
Lack of good third-party monitoring and compliance tools
INTERNAL BARRIERS
Lack of trust between IT stakeholders and wider business
Culture of risk-aversion and conservativism towards new technology
Limited understanding of relevant AI use cases
Lack of knowledge among employees about AI landscape
APPROVE BY DEFAULT

[8] Craig Hale, “Lack of clarity on AI regulation could be holding back businesses,” Tech Radar, June 28th, 2024
[9] Eileen Yu, “Transparency is sorely lacking amid growing AI interest,” ZDNET, May 10th, 2024
[10] Megan Morrone, “OpenAI is looking for new ways to pay its bills,” Axios, September 6th, 2024

> We definitely lean heavily on our vendors for support when it comes to these questions around shadow AI. We've been asking them lots of questions about audit processes, user permissions, capabilities, token length, etc."
>
> CTO, technology

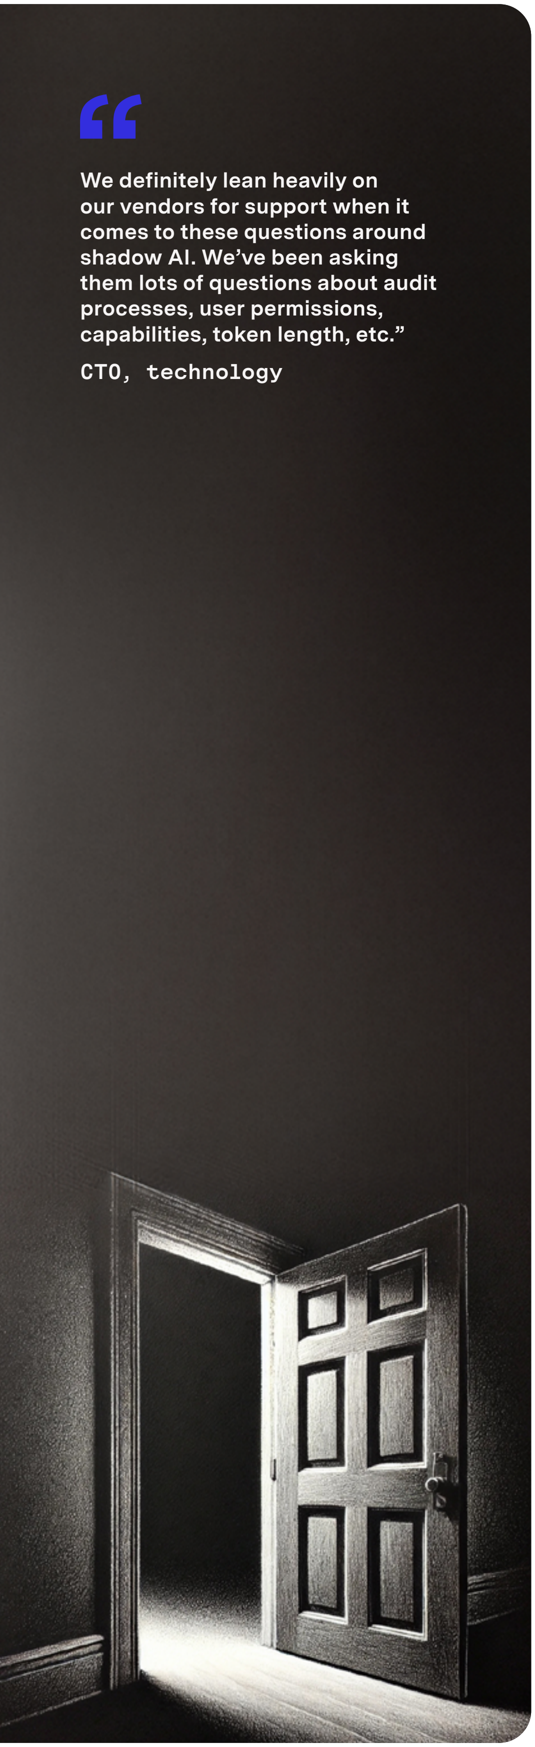

For tech companies, the looming specter of shadow AI offers an opportunity for differentiation

Ultimately, decision-makers within **enterprise organizations are keen to find ways to harness the spirit of innovation that leads to the emergence of shadow AI in the first place**—and to create the conditions under which employees can experiment with the technology without incurring serious security or reputational risks. But in order to do this, they need the companies providing them with AI solutions to meet them halfway and provide a roadmap for addressing some of the persistent uncertainties surrounding the technology.

For tech brands willing to grapple with this issue head-on, there's a real opportunity here for competitive differentiation. Right now, many buyers feel there's a lack of leadership around the issue of shadow AI. Tech companies, in their eyes, are too focused on tactically promoting individual AI use cases, and have not—for the most part—been willing to step back and provide their customers with a strategic vision for how to integrate AI usage across their organization within a consistent governance framework that keeps shadow AI to an absolute minimum.

The problem of shadow AI is not going away anytime soon. Indeed, unauthorized AI usage will only become more prevalent as workers become increasingly familiar with AI products and platforms. It's vital, therefore, that the tech industry gets ahead of the problem now, and that companies developing AI solutions **demonstrate to their customers an appreciation of the risks associated with shadow AI, and show that they're serious about working together as true partners to mitigate those risks**. Tech brands that rise to this challenge will find themselves with a distinct competitive advantage amidst the rapidly-evolving AI landscape.

> A lot of tech companies say they want to help businesses like us solve these problems of shadow AI. But in my experience, unless you're a truly huge organization, they're not really going to give you the resources to do it; they'll probably just point you towards a self-service portal or recommend an online course. I don't feel like these big players in the AI space really have much skin in the game when it comes to shadow AI, to be honest."
>
> CIO, financial services

> I feel that most tech vendors today lack a strategic point of view on how businesses can truly take advantage of their AI tools—how these tools can function at scale to enable business transformation. Mostly, they just seem focused on trying to sell as many licenses as possible, without real consideration for these bigger picture questions."
>
> Chief transformation officer, advertising

KEY RECOMMENDATIONS FOR TECH COMPANIES

How to support customers as they seek to manage the risks of shadow AI
and find an AI governance model that works for their organization

NEAR-TERM

Equip customers with the resources they need to understand the scale of the challenge and make informed decisions

Partner with academics and other relevant experts to produce independent thought leadership on shadow AI

Develop and deliver training programs on responsible AI usage targeted at a broad range of employees

Provide pre-built policy templates that buyers can adapt to their organization's needs

Our suppliers will often tell us that they have the software and the engines to solve a hundred different problems—but the problem with that is that I'm left trying to figure out the individual use cases for myself. It would be much easier for me to promote the use of AI within my organization if tech companies would give more specific guidance about individual use cases and their security and policy implications."

CIO, aviation

Transparency builds credibility. Too many of the big tech companies are focused on relentlessly hyping up AI. I'd trust them more if they were more grounded, and provided a realistic perspective on the limitations and the risks of the technology alongside the benefits."

CEO, technology

While we want to be the ones ultimately making the final decisions about our AI policies and our approach to dealing with shadow AI, we also recognize that the technology partners we work with should be an important voice in the room. Ideally, I'm looking to work with tech brands that understand my business and our sector, and can offer advice that doesn't feel obvious or generic."

CTO, advertising

MEDIUM-TERM

Build and deploy product features that enable organizations to implement their AI governance strategies

Embed security features such as multi-factor authentication (MFA) and end-to-end encryption into AI platforms that could be used to store sensitive customer information

Offer centralized AI dashboards that allow IT stakeholders to implement role-based access control (RBAC) and monitor for unauthorized usage in real-time

Build data governance and compliance features into core AI products, notifying users of governance violations and providing clear audit trails

Allow buyers to create controlled sandbox environments where employees can experiment freely

We've always been very clear about our AI philosophy: we don't want AI tools making decisions in isolation, there should always be that last layer where a human being looks at the recommendations and the wider context and uses all of that information to come to a final decision. So of course, we want to work with partners who share that philosophy and that perspective."

COO, software

When it comes to taking advice on shadow AI, we're always going to be more likely to trust technology partners who have a proven track record in cybersecurity."

CIO, consulting

LONG-TERM

Provide buyers with long-term reassurances around legal, regulatory, and business model uncertainties surrounding AI products

Embrace a greater level of transparency in the development of AI models

Provide buyers with contracts that clearly stipulate legal liability for data breaches and decisions made with the aid of AI

Proactively engage with policymakers to help address regulatory and compliance uncertainties

Offer long-term contracts or alternative cost models to allay buyers' fears of price increases

For more on AI, see NRG's previous reports...

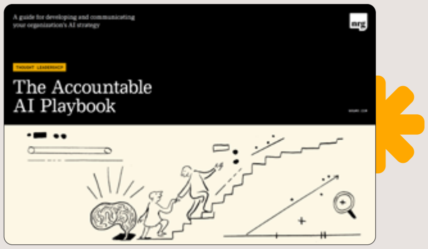
A guide for developing and communicating your organization's AI strategy
THOUGHT LEADERSHIP
The Accountable AI Playbook

The Accountable AI Playbook

Understanding consumers' fears and anxieties around AI, and how businesses can develop messaging that accommodates them

The AI-Powered Human

How AI could transform consumers' daily lives, and the use cases they're most excited about

nrg
The Enterprise AI Journey
Decoding how corporations are buying and deploying AI solutions

The Enterprise AI Journey

Decoding how corporations are buying and deploying AI solutions

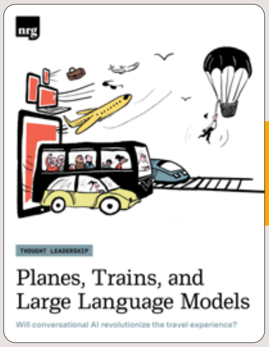
nrg
THOUGHT LEADERSHIP
Planes, Trains, and Large Language Models
Will conversational AI revolutionize the travel experience?

Planes, Trains, and Large Language Models

How AI could revolutionize the travel and hospitality industry

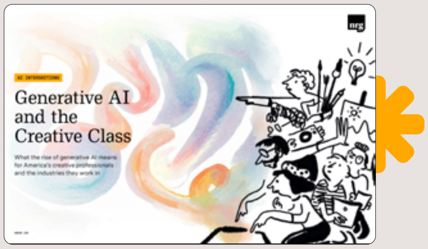
nrg
AI INTERSECTIONS
Generative AI and the Creative Class
What the rise of generative AI means for America's creative professionals and the industries they work in

Generative AI and the Creative Class

What the rise of generative AI means for America's creative professionals

Pressing Play on AI

The impact of AI on the video game industry

National Research Group is a leading global insights and strategy firm at the intersection of content, culture, and technology. The world's most innovative brands turn to us for insights into growth and strategy for any content, anywhere, on any device.

WORDS AND ANALYSIS BY

Fergus Navaratnam-Blair

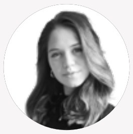
Jasmina Saleh

Rob Barrish

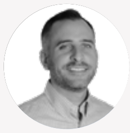
Nick Crofoot

Shannon Crocker

Nicole Clouser

Susan Hoxie

Aaron Williams

Maryalice Postel

Grady Miller

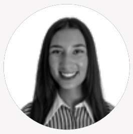
Emily Taylor

DESIGN BY

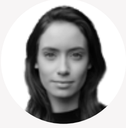
Victoria Lutz

Grace Stees