tines

Securing AI in the Enterprise

A step-by-step guide for IT and security leaders

Contents

Introduction
Eoin Hinchy, co-founder and CEO, Tines

AI has the potential to revolutionize how businesses operate, but its impact on the enterprise has been underwhelming. Years into the latest wave of AI innovation, very few teams have truly transformed their operations through AI.

McKinsey recently reported that just 10% of companies they surveyed had successfully implemented generative AI at scale for any use case.[1] When we explore why AI has fallen short of expectations, the same challenges surface time and time again. AI adoption is blocked by a small, consistent group of issues - misaligned priorities, underwhelming tools, lack of relevant skills, inflexible (or non-existent) AI policies, evolving regulations, and - most significantly - security and privacy risks.

The good news? For proactive IT and security leaders, security and privacy don't have to be obstacles. Tackling these specific security challenges head-on, with a security-first approach to implementing AI, can actually help us unlock the technology's full potential.

Balancing innovation and security can feel like threading a very small needle, and it's easy to see why skepticism and AI hype fatigue persist.

AI undoubtedly introduces complex security challenges, particularly when it comes to safeguarding organizational and customer data. But restricting AI's access to tools and systems isn't a perfect solution either.

AI's impact depends on access to proprietary data and the ability to perform tasks on our team's behalf. While we must tread carefully, we should avoid limiting AI to the extent that it fails to deliver value. When it comes to adopting AI securely, there's no one-size-fits-all framework or solution. Every organization's needs and objectives are different.

At Tines, we're fortunate to work with some of those exceptional teams that have already achieved game-changing results with AI. Their experiences and successes formed the inspiration behind this guide. These organizations have succeeded by investing in solutions with robust privacy and security guardrails, ensuring integration with evolving tech stacks, embracing AI alongside workflow orchestration and automation, and involving security and IT in managing AI priorities and risks across the enterprise.

With this guide, we hope to help IT and security leaders make measurable progress toward secure AI adoption. Whether you're facing a lack of AI governance, struggling to define priorities, or navigating a sea of hyperbolic vendor claims, this guide offers actionable strategies to help you keep moving forward.

You'll find step-by-step guidance to help you overcome blockers, accurately assess your AI needs, and securely deploy solutions that drive real business value. And we're not just sharing our perspective - this content is shaped by insights from top CIOs and CISOs, offering advice rooted in real-world experience.

We really hope this guide brings you closer to realizing AI's value for your organization.

Current state

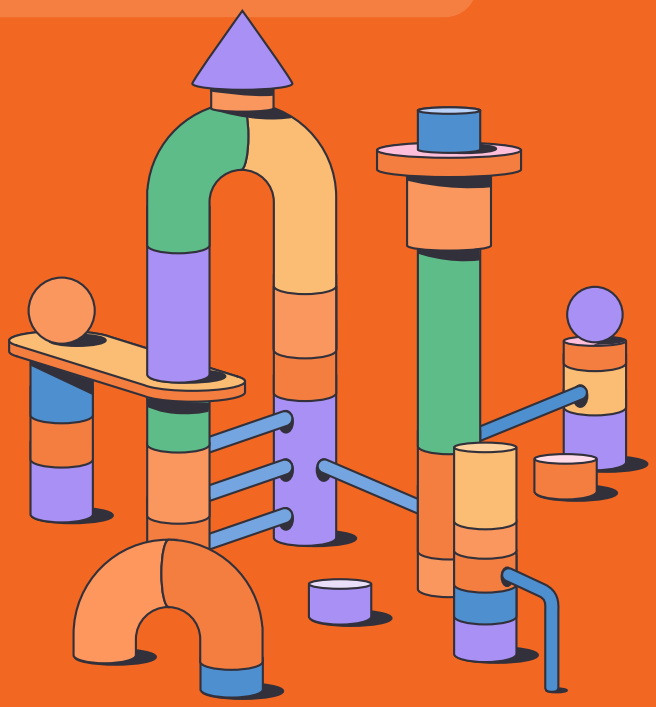

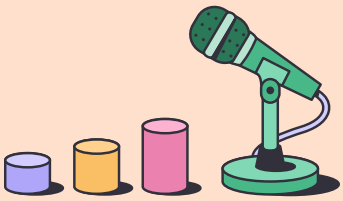

High potential, limited impact

AI's impact on the enterprise has been underwhelming so far due to rigid products that fail to connect data across technology stacks and extensive security and privacy concerns.

Rising demand meets real-world complexity

Pressure to adopt AI from business leaders and employees is growing, but progress is slow thanks to complex challenges like regulatory uncertainty and fragmented decision-making.

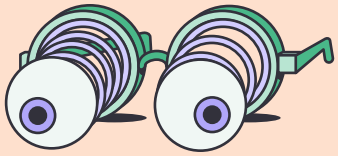

Privacy, security, and regulatory risks hinder adoption

Leaders remain cautious about exposing sensitive data or increasing attack surfaces, while evolving regulations make compliance a challenge.

Shadow AI creates hidden risks

Unauthorized tools and "shadow AI" introduce security blind spots as employees unknowingly share sensitive data in seemingly harmless prompts, which LLMs can collect and store.

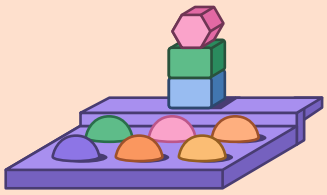

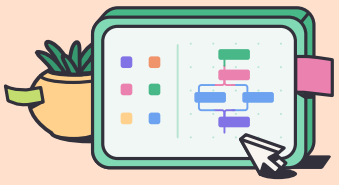

An evolving attack surface

AI enables attackers to move faster, increasing pressure on security teams to respond. Advanced processes are needed to keep pace with emerging threats.

Vendors and media add to the confusion

While AI is portrayed as transformative for attackers, most threats still rely on familiar tactics like phishing. Overhyping AI's role can distract already stretched teams from addressing proven vulnerabilities.

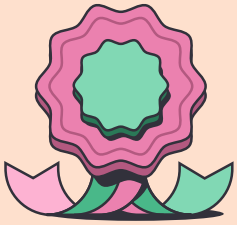

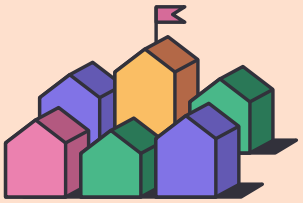

Early wins in AI adoption give us reasons to be hopeful

Despite the hurdles, some organizations are beginning to unlock the potential of AI in targeted ways. Next, we'll take a look at an example from the world of security operations.

samsara

EKCO

Case study
Samsara and Ekco

Security engineers Hela Lucas from Samsara and Kieran Walsh from Ekco share how, in just a few months, they've experienced the benefits of using secure and private native AI features within Tines' workflow orchestration and automation platform.

Enhanced productivity means analysts are empowered to focus on higher-value tasks. "AI helps provide additional context and explanations in a more human-friendly way, enabling us to focus on more complex tasks."

— Kieran Walsh, SOC Engineer, Ekco

Scalability means seamless tool integration enables rapid growth. "AI has already made things a lot quicker. The time to onboard and integrate a new toolset into our analyst's tech stack is now just five minutes."

— Kieran Walsh, SOC Engineer, Ekco

Time savings means significant reductions in time spent on repetitive tasks and reporting. "AI-generated ticket summaries save us so much time - around 15 minutes per case."

— Kieran Walsh, SOC Engineer, Ekco

Faster response times means reduced manual work and improved stakehold communication. "We use AI to translate technical information into readable language, speeding up response times and ensuring stakeholders understand the issues."

— Hela Lucas, Security Operations Engineer, Samsara

Improved accuracy means better spam detection and clear remediation instructions. "AI-powered spam filters have drastically improved the quality of life for our on-call team by reducing false escalations."

— Hela Lucas, Security Operations Engineer, Samsara

Key challenges with secure AI adoption

Recent surveys by Salesforce[2] and Tines[3] highlight the most pressing blockers for IT and security leaders.

Shared challenges, distinct priorities

While CIOs and CISOs share common concerns about security, privacy, and skills gaps in AI adoption, their priorities diverge in key areas. CIOs struggle more with identifying high-value use cases and proving ROI, whereas CISOs are more preoccupied with tackling misaligned priorities and addressing the limitations of inflexible technologies.

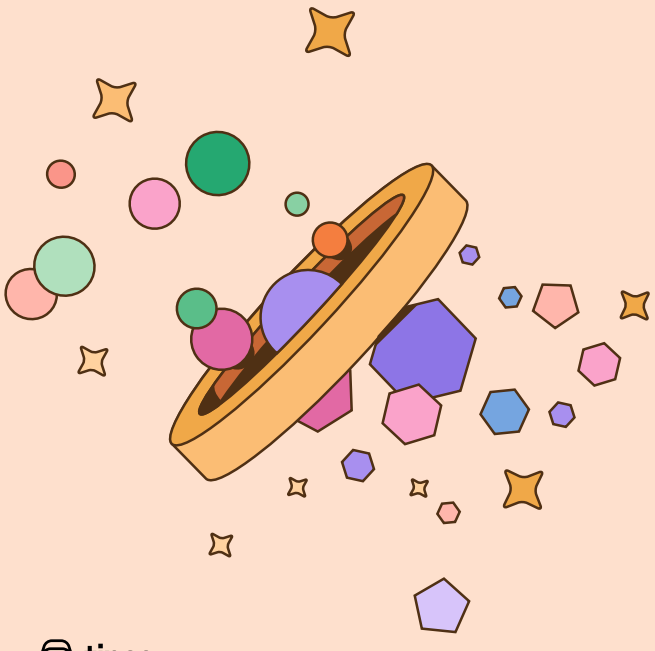

Top five blockers of AI adoption according to CIOs:

- 57% Security and privacy threats
- 52% Lack of trusted data
- 31% Training workforce with required skills
- 31% Inability to identify best use cases
- 29% Insufficient return on investment

Top five blockers of AI adoption according to CISOs:

- 77% Policies and perceptions
- 66% Data privacy concerns
- 60% Insufficient staff and skills
- 51% Misaligned priorities
- 49% Inflexible technologies

Step-by-step guide

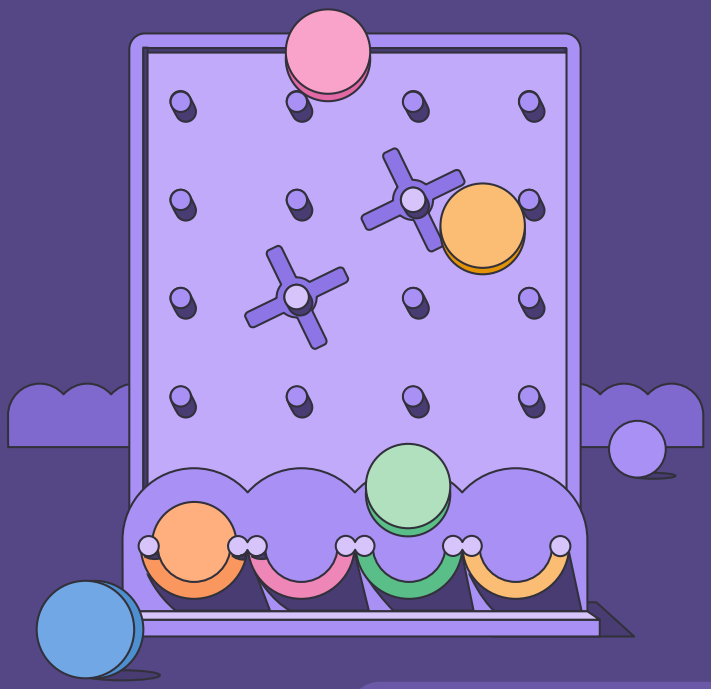

Step 1: Establish who owns AI security

Before security and IT leaders can begin their AI adoption journey, they need to establish which stakeholders will be involved and what policies and processes need to be considered.

They need to decide who has ultimate authority and how decisions will be made. Given that 77% of CISOs cite “policies and perceptions” as a blocker to AI adoption, ensuring alignment across business units through clear charters and protocols is crucial.

Popular AI governance models

Cross-functional committee

This centralized group should include representatives from key functional and business groups across the organization who use AI, including IT, security, legal, HR, among others. Periodic meetings and discussions drive priorities, while clear charters and protocols are key to how a group like this functions.

Business unit ownership

In this model, individual business units are equipped with guidelines for the secure adoption and use of AI tools, with clear distinctions on which tools or use cases require heightened security oversight. Support and guidance from security and IT teams remain critical to ensure compliance and effective implementation.

Security or IT ownership

This is a simpler, centralized approach to a cross-functional committee that may place accountability on a Chief AI Officer but it still requires security leaders to work with key stakeholders to gain support to implement key policies or tools.

CISO perspective: Security leaders must ensure holistic security across the AI lifecycle.

"It's so important that the CISO gets a seat at the table, in terms of creating an AI Center of Excellence, especially when we're thinking about the increased security risks. It's not just about the bad actors knocking on our door, but it is also securing all layers of the LLM, as well as the automation behind it — think holistic aspect of the AI lifecycle, not silos."

Gina Yaccone
Regional and Advisory CISO
Trace3

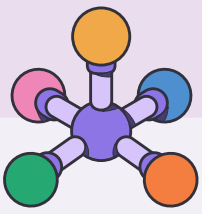

Step 2: Identifying the risks

As AI becomes increasingly integrated into business operations, understanding and mitigating potential risks is crucial to protecting organizational assets, maintaining regulatory compliance, and ensuring quality AI model data.

These risks reflect the challenges faced by technology leaders. As we saw in a previous section, CIOs cite issues such as "security and privacy threats" (57%) and "lack of trusted data" (52%), while CISOs emphasize "data privacy" (66%) and "misaligned priorities" (51%). Aligning leadership perspectives early can help expedite the process of identifying and addressing these risks effectively.

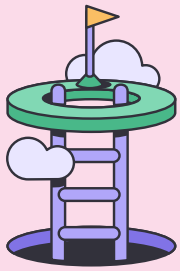

Regulatory and compliance requirements

Organizations must navigate a complex web of emerging AI regulations that vary across industries and geographies. This involves understanding legal frameworks such as:

- Data protection regulations (GDPR, CCPA)
- Industry-specific compliance standards (HIPAA for healthcare, FINRA for financial services)
- Emerging AI-specific legislation like the EU AI Act
- Potential liability issues related to AI decision-making
- International data transfer and localization requirements

Data privacy and security risks

AI systems are fundamentally data-driven, making robust data protection critical. Leaders must protect against:

- Potential unauthorized access to sensitive training data
- Risk of data breaches or model inversion attacks
- Challenges in maintaining data anonymization
- Vulnerabilities in data collection, storage, and processing pipelines
- Potential for inadvertent personal information exposure
- Complex data sharing and third-party vendor risk management

Model transparency and bias concerns

The "black box" nature of many AI models presents significant ethical and operational risks. These risks include:

- Algorithmic bias leading to discriminatory outcomes
- Lack of interpretability in complex machine learning models
- Challenges in explaining AI-driven decisions
- Potential reinforcement of historical prejudices
- Difficulty in conducting meaningful audits
- Risks of unintended consequences in critical decision-making systems

Risk mitigation strategies

Comprehensive risk management requires a holistic and proactive approach. Security and IT teams should:

- Implement rigorous testing and validation processes
- Develop continuous monitoring and evaluation frameworks
- Create incident response and remediation plans
- Invest in ongoing employee training and awareness
- Establish clear escalation protocols
- Conduct regular risk assessments and penetration testing
- Build redundancy and fallback mechanisms

Governance and policies

Effective AI risk management demands clear organizational guidelines. Leaders should consider:

- Developing comprehensive AI usage policies
- Creating cross-functional AI governance committees
- Establishing clear ethical guidelines for AI development and deployment
- Implementing mandatory training programs
- Defining clear accountability and responsibility matrices
- Encouraging a culture of responsible AI innovation
- Creating transparent reporting mechanisms for potential AI-related issues

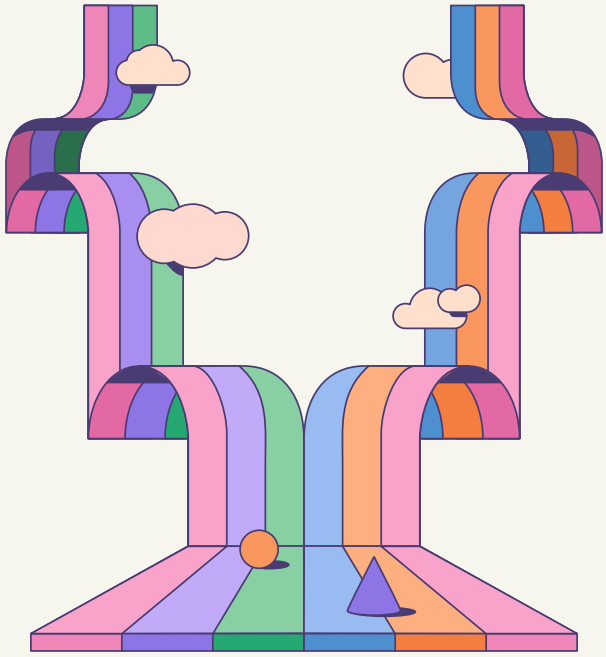

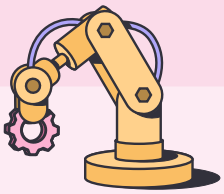

CISO perspective: AI tech selection by business unit balances control with empowerment

"I've come across a couple of organizations that are fascinating in their ability to prescribe what AI models certain business units can use. For example, in HR, you use this model to help support your decision-making, and if you need a new one, you can go through our selection process. So you have a policy that team members will feel informed by, enables them with the right technology for the right use case, and helps them feel secure and confident while they use it."

Matt Hillary
VP, Security, and CISO
Drata

Step 3: Familiarize yourself with AI security best practices

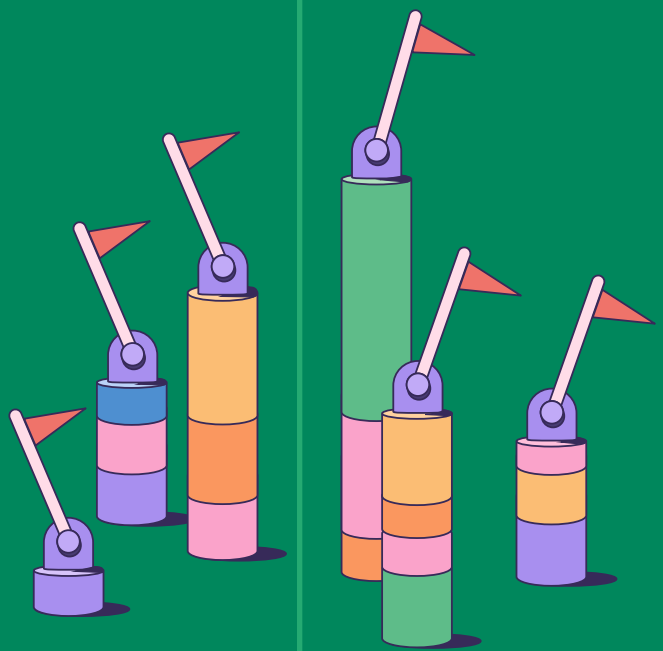

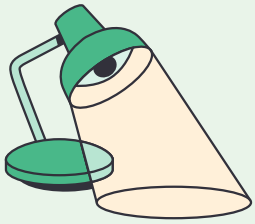

Organizations must develop dependable frameworks that protect both tech assets and organizational integrity while maximizing the transformative potential of AI technologies. Let's explore key best practices for achieving this balance.

Start small and iterate

Begin with controlled, low-risk AI implementations that allow for gradual learning and refinement. Pilot projects enable organizations to identify potential vulnerabilities, assess performance, and develop sophisticated governance mechanisms without exposing critical systems to widespread risk.

Invest in prompt engineering skills

Equip your team with prompt engineering expertise to enhance the reliability and predictability of AI systems. Well-designed prompts can dramatically reduce unintended outputs, mitigate potential security risks, and optimize AI performance across a wide range of applications.

Keep humans in the loop

Maintain critical human oversight in AI decision-making processes. Human judgment provides essential context, drives ethical decision-making, and nuanced understanding that AI cannot independently replicate, serving as a crucial safeguard against potential errors or biased outcomes.

Establish guardrails

Create robust technical and procedural constraints that guide AI behavior within predetermined boundaries. These guardrails include essential measures to prevent AI-generated "hallucinations" or false outputs, ensure compliance with data privacy regulations, and mitigate legal risks associated with AI decisions and actions.

Implement continuous monitoring

Develop real-time monitoring systems that track AI performance, detect anomalies, and enable immediate intervention. Advanced monitoring can identify potential security breaches, performance degradation, or unexpected behavioral patterns.

Invest in robust cybersecurity

Implement advanced cybersecurity measures specifically designed for AI systems, including encryption of training data and model parameters, secure access controls, and regular AI security audits.

Develop comprehensive training programs

Create organization-wide AI literacy and security awareness programs. Educate employees about potential risks, ethical considerations, and best practices for responsible AI interaction and management.

Maintain flexibility and adaptability

Design AI systems and governance frameworks that can quickly adapt to emerging technologies, regulatory changes, and evolving security landscapes. Build architectural flexibility that allows for seamless workflows.

CISO perspective: Hallucinations happen – your team needs a safety net.

"I think it's important to push your vendors on their AI philosophy, especially when you think about AI for security operations. Most of the foundational models, if you read through their documentation, say, 'Hey, be careful, our AI can and will hallucinate.' The problem is that AI won't tell you if it's hallucinating, it'll just do it. So you need to have a good story around falling back to a human or sanity-checking the output of an AI. If your vendor doesn't have a good answer around that, that's going to be a little bit of a red flag."

Matt Muller
Field CISO
Tines

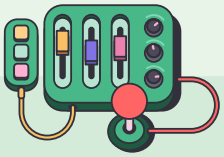

Step 4: Assess your AI needs

Organizations must approach AI implementation with a clear understanding of their strategic objectives and a well-defined framework for measuring success. This systematic approach ensures that AI investments align with business goals to deliver measurable improvements.

Before implementing AI solutions, organizations should:

- Identify core business challenges and opportunities
- Align AI initiatives with strategic business objectives
- Evaluate potential impact on customer experience
- Consider the competitive landscape and market positioning
- Assess resource availability and constraints

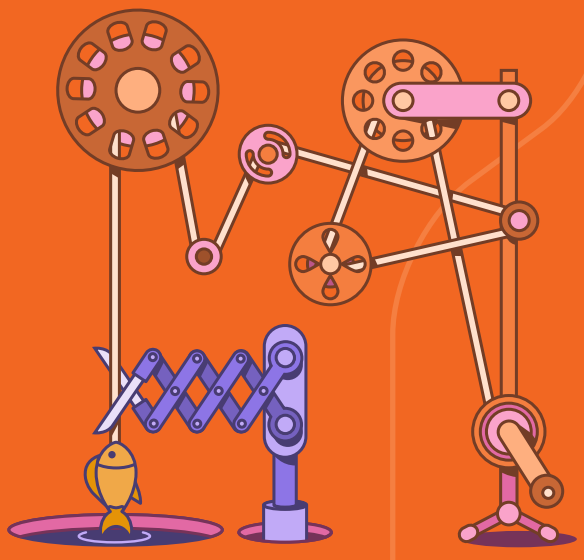

Set your milestones

6-MONTH MILESTONES

- Initial implementation completion
- Early adoption rates
- Basic functionality achievements
- Preliminary cost savings
- Initial user feedback metrics

12-MONTH MILESTONES

- Operational efficiency improvements
- Return on investment calculations
- User satisfaction metrics
- Process automation statistics
- Error reduction measurements

24-MONTH MILESTONES

- Long-term value creation
- Market position improvements
- Innovation metrics
- Scalability achievements
- Competitive advantage indicators

Establish key performance indicators

QUANTITIVE MEASURES

- Revenue impact
- Cost reduction percentages
- Time savings calculations
- Productivity improvements
- Error rate reductions

QUALITATIVE INDICATORS

- Employee satisfaction levels
- Customer experience improvements
- Process efficiency gains
- Innovation capability enhancements
- Progress on cultural objectives

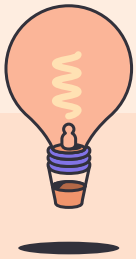

CIO perspective: Defining IT and business goals is key to ensuring AI's impact

"The most important thing to accomplish in the early stages of adoption is to reach agreement with business leaders about how to measure AI success. IT leaders may focus on things like time savings or productivity improvements, whereas business leaders may be far more concerned about reducing time to purchase or increasing customer promoter scores."

Mark Settle
Seven-time CIO

Step 5: Understanding AI tools

Sorting through an increasingly complex AI vendor landscape is no easy task but they can be broken down into three basic categories:

- **Chatbots** are virtual conversation partners that range from rule-based to AI-powered, providing basic interactions and support.
- **Agents** operate autonomously in the background, making independent decisions with minimal human intervention, capable of reasoning and adapting to changing circumstances.
- **Copilots** represent the latest advancement, functioning as digital teammates who work alongside users, providing context-relevant suggestions and automating repetitive tasks within specific product ecosystems.

Things to consider

Assessing AI tools requires a methodical approach, and, ideally, a standardized scorecard to determine the best fit for a given use case. Consider these attributes when evaluating new AI tools:

Security and privacy

Trust is crucial. Seek AI tools with robust security measures, preferably within your infrastructure. Investigate data handling, training processes, and alignment with your company's AI policy and compliance regulations.

Return on investment

Avoid purchasing flashy tools without clear purpose. Define specific business challenges, compare unique benefits over personal AI solutions like ChatGPT, and carefully assess pricing models. At the very least, AI should provide meaningful assistance in automating routine tasks.

Accuracy

Recognize that no AI tool is 100% perfect. Understand expected accuracy rates, potential false positives, and bias reduction efforts. Maintain human oversight and consider trialing the product or speaking with existing customers.

Speed

Evaluate AI response times during live demos. Slow AI can hinder team productivity, so ensure the tool provides timely, efficient outputs.

Scalability

Choose tools that can grow with your organization. Request metrics on data input volume, output capacity, error rates, and performance under expanding team needs.

Usability

Assess user-friendliness and prompt engineering requirements. Seek tools that are intuitive to all users. Using AI should enhance process effectiveness rather than add unnecessary complexity.

Model selection

Look for tools offering multiple models for different task complexities. The ability to switch between models can significantly impact ROI and performance.

CISO perspective: Vendors should empower, not dictate, with AI solutions.

"I see a desire for more transparency in the AI space. From a product perspective, it's about being explicit and letting customers use what works best and what's approved by them and helps their environment. It's not about [vendors] dictating to customers what needs to be there."

Mandy Andress
CISO,
Elastic

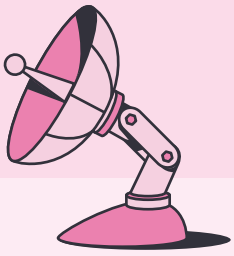

Step 6: Evaluate and purchase

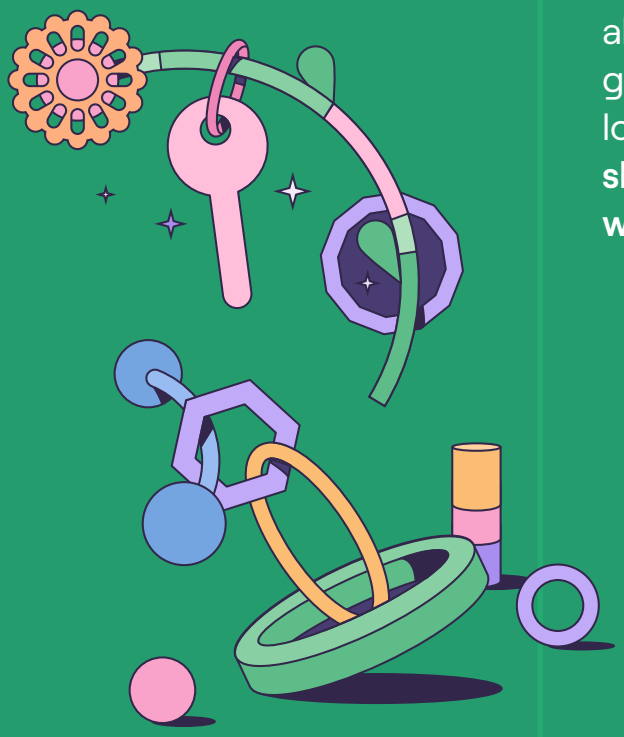

Choosing the right AI tool is a pivotal step in your AI adoption journey. With countless options available, a thorough evaluation ensures that the tool you select aligns with your organization's goals, operational needs, and long-term vision. **Next, we'll share some questions to ask when evaluating AI tools.**

Effectiveness and impact

- Why is this better than opening ChatGPT in a new tab?
- Does this work for my real use case?
- What are other customers using it for?
- Can I speak to one of your customers?

Performance and accuracy

- How long does an execution take? Can you show me in real time?
- How accurate is the AI? Can you share it in percentages?
- How many false positives can I expect?
- How are bugs and issues handled?

Security and privacy

- Is this AI via API?
- How will my data be handled?
- Will my data travel around the internet?
- How is the model trained on my data?
- How is the metadata or data being logged?
- How are humans being kept in the loop?

Customization and flexibility

- What language models are supported?
- Can I bring my own LLM?
- Are the models set globally or can I specify the model for a specific task?
- Can I easily turn off the AI-powered features if I need to?

Pricing

- What's the pricing model?
- How does pricing scale as your usage increases?
- How can I view and control how much is being spent?
- How is ROI measured for this tool?

Scalability and limitations

- What are the limitations for using the features?
- Has the solution been evaluated for bias?
- How well does it scale? Any examples from fast-growing customers?

Usability

- Can everyone on my team use the tool?
- What level of training will my people need to use the tool?
- Does the feature give you options to iterate before committing to the set prompt?
- How intuitive is it for users to do that iteration?
- What resources (best practices etc.) are available for users?

Trial and evaluation

- Is there a free trial, demo period, or proof of concept project available?
- Can we test the tool with a small dataset or a limited use case before full deployment?

Policy and compliance

- How well does this tool align with our internal AI policy?
- Is the tool equipped to address regulatory compliance (e.g., GDPR, CCPA)?
- What would be my next steps if there are major changes to AI policy or compliance requirements?

Future proofing

- Is the feature designed to improve over time?
- How does AI figure into your product roadmap?
- What else is on your product roadmap?

Step 7: Launch an AI pilot program

A well-designed AI pilot program serves as a critical first step in understanding the potential benefits, risks, and operational implications of AI tools while maintaining robust security protocols. Your pilot program should prioritize organizational learning, tech assessment, and risk management.

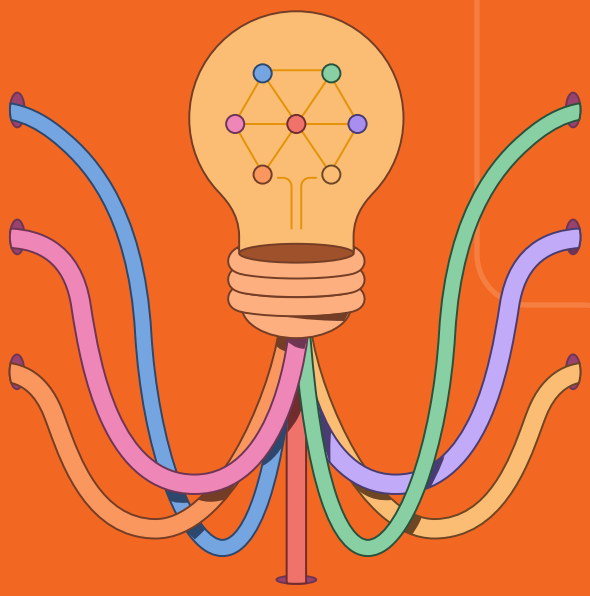

1. Choose a manageable scope

Select a specific, well-defined use case that aligns with organizational goals.

- Avoid overly complex initial implementations that might overwhelm your team or introduce unnecessary complexity
- Start with a narrow, high-impact area where AI can demonstrate clear value and minimal risk

2. Provide access to quality data

Relevant, high-quality data is crucial for AI success. Conduct a thorough data audit to verify:

- Data accuracy
- Comprehensive representation
- Compliance with privacy regulations
- Minimal bias
- Appropriate anonymization and protection mechanisms

3. Monitor performance regularly

Implement robust monitoring mechanisms to track AI tool performance, accuracy, and potential security anomalies.

- Establish key performance indicators (KPIs) and security metrics
- Use automated and manual monitoring techniques
- Create real-time alert systems for unexpected behaviors

4. Plan for failures and iterations

Recognize that initial implementations will have limitations and potential failures.

- Develop a robust risk mitigation framework
- Create clear protocols for identifying and addressing AI system failures
- Establish a culture of continuous improvement and learning
- Maintain flexibility in your approach

5. Ensure security

Before turning on AI technology, be sure to carry out security due diligence. Some key checklist items include:

- Assess data handling and privacy protection mechanisms
- Review compliance with industry standards (NIST, ISO 27001)
- Implement multi-layered security protocols
- Use the principle of least privilege for AI system access
- Establish robust authentication and authorization mechanisms
- Encrypt sensitive data
- Create incident response plans

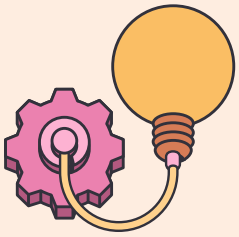

6. Create a controlled setup for organizational learning

A controlled setup allows organizations to learn and refine their AI approach before wider implementation. This approach provides a low-risk environment to:

- Test AI technologies in a contained setting
- Assess performance and compatibility with existing systems
- Identify potential security vulnerabilities
- Build internal expertise and confidence
- Develop adaptable implementation strategies

7. Gather comprehensive feedback

Create multiple channels for stakeholders to provide insights, concerns, and observations. Develop anonymous and transparent feedback mechanisms, and include perspectives from:

- IT security teams
- End users
- Legal and compliance departments
- Business unit leaders

8. Document everything

Maintain clear records of your pilot program.

- Track implementation details, performance metrics, security assessments, and lessons learned
- Create standardized documentation templates
- Ensure documentation is accessible to key stakeholders
- Use documentation to refine future AI implementation strategies

A well-executed AI pilot program is more than a technological experiment — it's a strategic approach to understanding and safely integrating AI into your organization.

Next steps: report and repeat

Use the priorities, goals, and KPIs established in Step 4 to assess the report and provide the results to stakeholders. Based on the outcomes, start the process again at Step 7 to pilot a new use case or AI tool.

CIO perspective: Don't rely on a single use case to prove AI's value.

"It's critical to not place 'all your eggs in one basket' in experimenting with initial use cases. You want to spread your bets across a cross-section of opportunities and not simply hope for the best with one or two initial prototyping experiments. Business leaders may be enthralled by AI at the moment but experience has shown that their enthusiasm for new technology quickly evaporates if it fails to produce immediate tangible results."

Mark Settle
Seven-time CIO

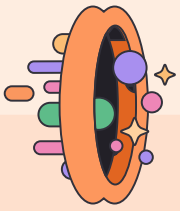

Realising AI's promises securely

AI has the potential to revolutionize how enterprises operate.

But successful AI implementation is not always as seamless as some vendors might claim. And AI is undeniably complicating the process of protecting an organization's data, mitigating risk, and ensuring compliance with evolving regulations. By adopting a disciplined approach to AI governance and security practices, organizations can not only realize AI's potential but also enhance their overall security posture in the process.

Organizations that are thoughtful in establishing AI leadership, policies, and security protocols will be able to thrive without creating new risks, helping the enterprise adopt AI technologies with speed and confidence.

At Tines, we're committed to delivering meaningful AI applications that build on the robust capabilities of workflow orchestration and automation. This includes native AI features that are secure and private by design yet powerful enough to drive meaningful business outcomes, and products like Workbench — a Tines-powered AI chat interface where you can take action and access proprietary data in real-time, privately and securely. This approach ensures AI becomes an IT leader's trusted ally rather than a source of new challenges.

In our conversations with Elastic's CISO Mandy Andress, she reflected, "I want to be able to look back five, eight, 10 years from now and look at this as the dark ages of security. How did we even operate? We do things so differently now. How did we even succeed with how we did things back then?" This is the future state we're excited about — a future where AI has proven its value, seamlessly integrated across all business units as an indispensable tool for enhancing team effectiveness.

Achieving this outcome begins with the steps your organization takes today, tomorrow, and over the coming year — steps that position AI as a catalyst for innovation while maintaining the highest standards of security and trust.

→ Learn more about AI in Tines and Tines Workbench: tines.com/ai

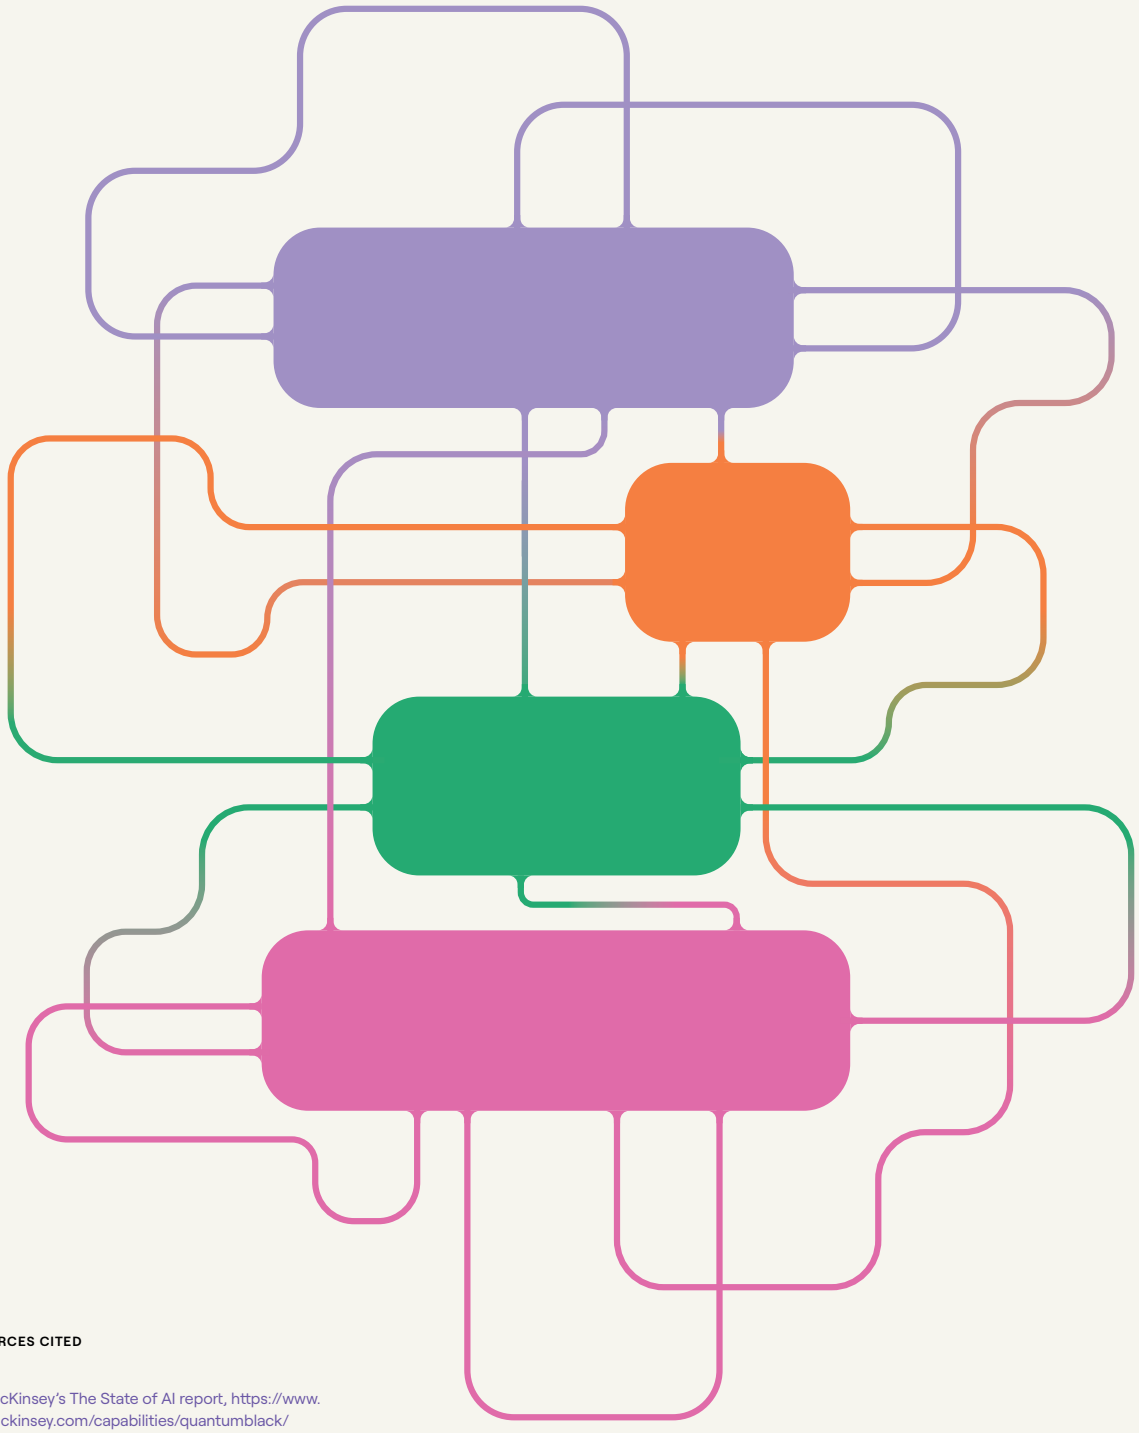

RESOURCES CITED

1. McKinsey's The State of AI report, https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
2. Salesforce survey of 150 CIOs, https://www.salesforce.com/news/stories/cio-ai-trends/
3. Tines, CISO perspectives: Separating the reality of AI from the hype, https://www.tines.com/reports/ciso-perspectives-ai/

tines