



The CIO Playbook: Proving the ROI of AI

Introduction

AI adoption in the enterprise has entered a new phase. Generative AI tools—chatbots, assistants, copilots—have rapidly moved from experimental pilots to embedded parts of daily workflows. For CIOs, that’s created a more urgent mandate to not just ensure security and compliance, but prove whether these tools are actually delivering business value at scale.

Yet measuring the ROI of gen AI remains a challenge. Most tools promise productivity gains, but few come with built-in ways to validate their impact across quality, user adoption, or business outcomes. And with the growing sprawl of AI point solutions, CIOs risk managing a fragmented ecosystem that adds cost and complexity instead of clarity.

At the same time, the next wave of AI—agentic systems that can reason, make decisions, and orchestrate complex tasks—is emerging. CIOs must start applying what they’ve learned from current AI deployments to prepare for this next shift. That means establishing ROI benchmarks, defining usage standards, and scaling trusted tools that work across teams and platforms.

This playbook gives CIOs a practical framework and flexible spreadsheet to do just that. It helps IT leaders develop a formula to identify what’s working, cut what isn’t, and lay the groundwork for a more strategic AI ecosystem.



Table of contents

Introduction	00
The AI ROI Scorecard	03
Measures of Success	04
Needs and Usage Mapping	06
Decision-Making	07
Scaling AI With Confidence	09
About Grammarly	10



The AI ROI Scorecard

The AI ROI Scorecard is a practical, data-driven spreadsheet designed to help your IT team quickly assess whether to keep, reject, or scale AI tools across your enterprise.

This scorecard simplifies decision-making by focusing on key areas that, together, provide a holistic view of each tool’s performance and potential to scale. Here are the key areas that we will walk through in more detail in sections to come:

- **Tool information:** On the left, input the key information you’ll need to evaluate any AI tool, including the vendor name, who at your company is requesting it, and what type of AI it provides.
- **Measures of success:** At the top of the spreadsheet are the measures of success. These are the top-line indicators you’ll use to assess AI performance, including compliance, quality, employee experience, and impact.
- **Usage mapping:** Next, map the tool’s usage to clearly identify where and how it is currently being used across the organization while assessing its alignment with business needs.
- **Decision:** Finally, evaluate the results and determine how to move forward. You should also set a reevaluation date to ensure continued alignment with evolving business goals.

Evaluating AI Impact and ROI

Vendor name:

Requested by:

AI type:

Measures of Success: KPIs and Outcomes

Compliance	Quality	Employee experience	Impact

Usage Mapping: Need and Usage

Business need	Use cases and functionality	Power users

Decision

Decision	Reevaluation date:	Additional context

[Download the Scorecard](#)



Measures of success

Effectively evaluating AI tools requires a nuanced understanding of success. Traditional metrics like productivity, time savings, or raw output often fail to tell the full story. High productivity may come at the expense of quality, while time savings may indicate poor work done faster. CIOs must break down productivity into specific, measurable components to ensure these tools deliver lasting value today and create a foundation for more advanced use cases, like agentic AI, tomorrow. Here are the four key measures of success, each building on the previous one. If a tool fails at any stage, it's time to move on and evaluate another.

1

Compliance

Definition: Compliance refers to whether the AI tool meets your organization's security, privacy, and regulatory requirements. This is the first and most crucial checkpoint; no other factor matters without compliance.

How to assess: To evaluate compliance, ask the AI vendor to complete your organization's security and privacy questionnaire. This assessment should cover data handling, encryption practices, and adherence to regulations like GDPR or the California Consumer Privacy Act (CCPA). If the tool doesn't meet your compliance standards, it's an immediate no. If it does, you have the green light to move on to the second measure of success.

Tools that integrate naturally into existing employee workflows—without requiring extensive onboarding—are especially valuable. This ease of use not only drives faster adoption but also builds familiarity with AI in a way that prepares employees for more advanced AI use down the line.

Want to see how Grammarly stacks up against your security requirements?

[Download our answers to your AI vendor questions. →](#)



2

Quality

Definition: Quality refers to the AI tool's ability to produce outputs that meet your organization's standards and users' needs. Outputs are the text, images, or analysis AI generates. High-quality output is essential for the tool to be valuable.

How to assess: Quality can be assessed both quantitatively and qualitatively. A/B testing is a useful method here: Have one group of employees use the tool and another work without it. Compare their performance in terms of the quality of the output produced. You can also gather qualitative feedback through surveys or interviews, asking users if the tool meets their expectations and delivers the intended results.

3

Employee experience

Definition: Employee experience measures how effectively employees use the tool and whether it improves their productivity and workflows. Even a high-quality tool can fail if it doesn't drive employee adoption.

How to assess: Employee experience can be gauged using surveys and feedback loops, ideally integrated into the same A/B test used for quality. Look at key indicators such as time spent using the tool, perceived ease of use, and whether the tool enhances employees' productivity. You can also look at whether employees are actually adopting the tool and integrating it into their regular workflows.

Example employee experience survey questions

- How often do you use the AI tool?
- How intuitive are the interface and functionality?
- Do you feel that the tool makes your job easier or harder?
- How would you rate the quality of the AI's output?
- How much time do you estimate this tool has saved you?

4

Impact

Definition: Impact refers to whether the AI tool drives measurable business results. Ultimately, the success of any tool is determined by its contribution to core business metrics.

How to assess: To assess impact, look at how the tool moves the needle on key business priorities. Metrics to consider include customer satisfaction scores, time to resolution for customer issues, customer retention or churn rates, lifetime value, and brand sentiment. If the tool positively affects these metrics, it's a strong indicator that it delivers real value. Establishing these impact benchmarks now is essential—not just for proving ROI today, but for informing early evaluation frameworks as you pilot agentic AI systems.

By understanding and applying these four assessment points, CIOs can more effectively evaluate AI tools, identifying which ones are worth further investment and scaling. The goal isn't just to measure ROI—it's to build a scalable evaluation model that evolves alongside your AI strategy.



Needs and usage mapping

Once you've measured success, the next step is mapping how AI is actually being used across the organization. In today's landscape—where many teams have adopted AI tools independently—usage mapping is no longer just about identifying value. It's about understanding where AI is embedded, where it's fragmented, and where it can be unified to scale securely and strategically.

Usage mapping is crucial for identifying how well the tool fits within your current and future business needs, uncovering potential redundancies, and recognizing power users who gain exponential value. By understanding these areas, you can determine whether a tool is worth scaling and how to maximize its potential across teams.

Here are the three key areas to focus on in the usage mapping process:

1. Business need

Goal:

Identify the core business problem the AI tool is solving, whether it's addressing existing pain points or unlocking new value.

How to assess:

- Evaluate how the tool aligns with your organization's top priorities.
- Determine whether the tool solves an existing business issue (e.g., improving efficiency, reducing costs) or opens up new opportunities (e.g., driving innovation or creating new revenue streams).
- Consider whether the tool supports evolving needs, like automation of complex workflows or integration into future agentic AI systems.

2. Use cases

Goal:

Identify the different use cases of the tool across teams and evaluate potential redundancies or new opportunities for its application.

How to assess:

- Map out the various ways the tool is being used across departments.
- Look for any overlaps between tools (where different teams may use similar tools for the same purpose).
- Consider whether redundancies are necessary or if a consolidation is possible.
- Explore if there are untapped use cases across teams where the tool could deliver additional value.

3. Power users

Goal:

Identify who in the organization is deriving exponential value from the tool and assess whether their usage patterns can be replicated to benefit others.

How to assess:

- Seek out power users who are highly active with the tool and vocal champions of its usage and who request or advocate for additional functionalities.
- Analyze how their usage differs from others in the organization and whether those behaviors can be scaled to other teams.
- Use insights from power users to better inform AI workforce fit and unlock new potential use cases.

Decision-making

With the rest of your scorecard completed, it's time to make a call: Is this tool ready to scale, worth keeping in limited use, or best left behind? As AI adoption matures, CIOs must not only evaluate what works but also actively shape a streamlined, secure, and scalable AI ecosystem. Every decision you make today helps determine how ready your organization is for what's next, including the rise of agentic AI.

Reject: The tool fails to meet critical requirements, such as compliance standards, lacks quality output, doesn't integrate well into employee workflows, or shows no meaningful impact on business priorities.

Reasons to reject:

- Doesn't meet risk and compliance standards
- Unjustifiable redundancy
- Irrational cost
- Low adoption

Next steps:

- Discontinue the use of the tool, freeing up resources to focus on more promising options.
- Document the reasons for rejection to guide future decision-making processes and avoid similar missteps with new tools.
- In some cases, it may be worth monitoring the vendor for significant updates that could address your organization's concerns in the future.

Allow as is: You've identified value in the tool, but it hasn't yet demonstrated enough potential to be scaled across the enterprise. It may meet compliance standards and provide solid output but lacks strong adoption or impact on business-critical metrics.

Reasons to allow:

- Meets risk and compliance standards
- No or low cost
- Narrow user base, high usage
- Promising explorations and power users

Next steps:

- Keep the tool in use for its current applications while monitoring for further developments or improvements.
- Set a reevaluation date for six months to see if the tool's capabilities or fit with your organization evolve, allowing it to better meet the bar for scaling.

Scale: The tool has demonstrated clear, measurable value across the four measures of success—compliance, quality, employee experience, and business impact. It’s ready to be deployed more widely within the organization to drive further ROI.

Reasons to scale:

- Measurable business impact
- Scalable with enterprise-grade management
- Consolidates functionality
- Product roadmap alignment
- Potential for future orchestration or agentic integration
- Supports cross-functional or multisystem workflows


Next steps:

- Begin scaling the tool across additional teams, departments, or processes.
- Ensure adequate training and support for employees to maximize the tool’s effectiveness at scale.
- Continue tracking the tool’s performance to ensure it maintains its value as usage expands.

Setting a reevaluation date

For tools that you decide to allow as is or even reject, it’s important to set a reevaluation date, typically on a six-month cycle. AI tools evolve rapidly, and what doesn’t meet your needs today may do so in the near future. Vendors frequently update their offerings, and new use cases may emerge as the tool is refined.

When setting the reevaluation date, consider:



New features or updates

Stay informed about updates or new features released by the vendor that could improve the tool’s value.



Changing business needs

Assess whether evolving business priorities create new opportunities for the tool to add value.



User feedback

Collect updated feedback from employees and power users to gauge any shifts in the tool’s fit or performance.

The more intentional your evaluation cycles are now, the better positioned you’ll be to adopt—and confidently scale—the next generation of AI.

Scaling AI with confidence

With the right tools, measurements, and mindset, IT can lead the next phase of AI transformation, delivering value today and building the infrastructure for tomorrow.

The decisions you make now will shape the maturity and impact of your AI ecosystem. With the right approach, IT doesn't just support AI adoption. It defines it.

[Download the Scorecard](#)





About Grammarly

Grammarly is the trusted AI assistant for communication and productivity, helping over 40 million people and 50,000 organizations do their best work. Companies like Atlassian, Databricks, and Zoom rely on Grammarly to brainstorm, compose, and enhance communication that moves work forward. Grammarly works where you work, integrating seamlessly with over 500,000 applications and websites. Coda, the maker of powerful AI productivity tools, is now a part of Grammarly. Together, they are redefining collaboration and transforming how we get work done. Learn more at grammarly.com/enterprise.