



Moveworks

The IT leader's guide to investing in an AI copilot

Intro

You're slammed with emails, meetings, and a million little tasks eating away at your day. Meanwhile, your team is drowning in repetitive queries and tedious processes. Sound familiar? Then you need to keep reading...

The future of work is already here, and AI copilots are leading the charge. Today's conversational AI doesn't just answer simple questions – it understands your enterprise inside and out. With deep integrations into all your systems and data, a copilot can search across systems, automate workflows, take action, and level up your team's productivity in ways you've only dreamed about.

But building a true enterprise-grade copilot is no simple feat. Get it wrong, and you'll just have an expensive novelty on your hands. That's why we've assembled this guide – to ensure you invest in the right solution to supercharge your business.

In the following pages, we'll dive into the evolution of AI copilots, the pitfalls to avoid, and real-world examples of companies crushing it with this game-changing technology. You'll get an inside look at the four tiers of copilot sophistication and what it really takes to implement one successfully.

Most importantly, you'll learn how to determine whether building or buying is the right path for your needs. After all, time is money – and your employees' time is incredibly valuable.

So grab a coffee, and let's explore how an AI copilot can finally start working for you, not the other way around. The future of work is waiting...let's get started!

Table of Contents

01

The evolution from conversational AI to AI copilots

04

The cost of building vs. buying an AI copilot

02

The challenges of building an AI copilot

05

Successful AI copilots in action

03

How to avoid AI investment pitfalls

The evolution from conversational AI to AI copilots


AI is transforming how businesses operate and deliver value. At the forefront is conversational AI with technologies like chatbots given new life by ChatGPT bursting onto the scene.

While earlier forms of conversational AI focused on simplistic tasks like FAQs and customer service, a new class of solutions is emerging to truly revolutionize the enterprise: AI copilots.

This chapter explores the evolution from basic conversational AI to sophisticated enterprise AI copilots, answering:

- What is conversational AI and how is it used in business today?
- What are AI copilots and how do they advance conversational AI?
- What are the different types of AI copilots?
- What are key considerations when selecting an AI copilot?

Equipped with this foundation, you'll understand what sets copilots apart and how they can transform efficiency, productivity, and experience.



What is conversational AI?

Conversational AI enables machines to interpret and respond to human language, creating a more natural interaction between humans and machines.

Leveraging natural language processing (NLP), machine learning, and artificial intelligence to mimic human-like conversations, conversational AI can be used in various business applications, including employee service, customer service, HR functions, and internal communication.

Businesses of all types are using conversational AI

In a world where efficiency and speed are paramount, business leaders are constantly seeking new ways to streamline their processes and better serve their customers and employees. That's where conversational AI comes in.

By using copilots and other conversational tools, companies can **provide instant help** to their customers and employees, **reduce wait times** and **free up support teams** to focus on more complex tasks.

The benefits of conversational AI go beyond support. For example, in HR, conversational AI can automate tasks like scheduling interviews, answering employee queries, and managing leave requests, taking the burden off of HR professionals and allowing them to focus on more strategic work.

Another major advantage of conversational AI is the potential to **improve the employee experience**. By

automating tedious and repetitive tasks, AI can help employees focus on more high-value activities that require human expertise, ultimately increasing job satisfaction and productivity.

All in all, it's no surprise that businesses of all types are eagerly adopting conversational AI in a variety of different forms. With the potential to increase efficiency, improve support, and enhance employee experience, conversational AI is quickly becoming a must-have tool for modern businesses.



Three examples of conversational AI in the enterprise

01

Unity levels up its tech support with conversational AI on Slack

[Unity](#), a leading platform for creating and operating interactive, real-time 3D content, successfully implemented conversational AI to enhance its employee experience.

The company's CIO emphasized the importance of employee experience and how it plays a crucial role in enhancing overall organizational performance. With employees submitting their IT issues conversationally on an #ask-IT Slack channel, Unity's support team had to keep track of dozens of ad-hoc issues.

Thanks to conversational AI, Unity's service desk was able to address a significant portion of these issues automatically. All employees had to do was type in their question, just how they would engage with a help desk agent. With conversational AI, instead of waiting days for support, employees now get help in less than a minute resulting in a happy team.

In fact, more than 90% of Unity's employees are satisfied with their IT experience – and that's nothing to scoff at!

"Employee experience is my top priority. Having employees submit IT issues on a Slack channel posed some unique challenges, but we do it because employees like it. And with AI addressing a lot of those issues automatically, our service desk likes it, too."

Brian Hoyt
Former CIO, Unity

02

Luminis Health uses conversational AI to support frontline healthcare workers

Conversational AI is transforming various industries, including healthcare. It offers numerous benefits, from improved patient care to enhanced operational efficiency. One such example is [Luminis Health](#), where an IT Director implemented conversational AI to up-level his team and provide better services to patients.

One of the most significant advantages of conversational AI in healthcare is its ability to automate routine tasks.

For instance, conversational AI is more than capable of handling password resets, appointment scheduling, and other repetitive tasks, freeing healthcare workers' time to focus on more critical responsibilities.

With AI resolving issues remotely and instantly, often without agent intervention, IT saw a 25% reduction in IT call volume two weeks after launching their conversational AI solution.

It is worth noting that implementing conversational AI is not about replacing human resources; instead, it is an opportunity to up-level team members by allowing them to focus on high-value tasks. Investing in AI is an investment in the team's upskilling, enabling them to work more efficiently and productively.

Luminis Health's success in implementing conversational AI highlights the technology's potential to transform the healthcare industry.

“Investing in AI is not about replacing FTEs. It’s an opportunity to up-level your team. You don’t use your people, your most valuable resource, to reset passwords. AI can take care of that.”

Andre Green
IT Director, Luminis Health

03

Albemarle doubles employee productivity with conversational AI

For global enterprises like the [Albemarle Corporation](#), providing consistent, high-quality IT support to all employees, regardless of location or language, can be daunting. But with the help of conversational AI, Albemarle has achieved this goal.

The company's CIO explains that the conversational AI solution provides 24/7 help to employees in their [native language](#), allowing everyone to get the support they need quickly and easily without the need for localized service desks around the world.

In addition to providing IT support to employees, conversational AI can [pull insights from backend IT systems](#), helping Albemarle turn thousands of requests into a simple, actionable to-do list.

While current ITSM insights tools focus primarily on tickets and SLAs, conversational AI can help companies identify trends and issues before they become major problems and proactively address them. This approach offers Albemarle's support leaders granular insights, allowing them to immediately see and [address inefficiencies](#) across the company.



“With AI, we can provide the same high-quality support to every employee, no matter which languages they speak.

It gives our people 24/7 help in their native language, conversationally. Now, they can get support right away without us needing localized service desks all over the world.”

— **Patrick Thompson**, Former CIO, Albemarle Corporation



Making conversational AI better for the enterprise

Clearly, conversational AI is already transforming the business landscape in unprecedented ways, and its adoption is only accelerating with the rise of large language models.

Especially with ChatGPT and other similar tools making headlines, conversational AI has gained popularity across industries, with dozens – if not hundreds – of startups popping up. But simply making API calls or integrating with a singular large language model won't give you the results you want in an enterprise setting.

For a conversational AI to create a lasting impact within an enterprise, it needs to connect each system and application across an enterprise, allow them to communicate seamlessly, and then take action based on the user's needs. For instance, while you could ask a chatbot like ChatGPT to add you to a sales distribution list, it doesn't have the knowledge or ability to understand and act on your request.

Without deep integrations with company-specific data and the systems and apps within your organization, conversational AI use cases will be lackluster at best and downright useless at worst.

This is where AI copilots come in and take conversational AI to the next level.

What is an AI copilot?

An AI copilot is a conversational interface that uses [large language models \(LLMs\)](#) to support users in various tasks and decision-making processes across multiple domains within an enterprise environment. By leveraging LLMs, AI copilots possess the capability to understand, analyze, and process vast amounts of data.

AI copilots play a crucial role in enhancing productivity and efficiency by adding to conversational AI:

- 1. Providing context-aware assistance:** AI copilots can proactively respond to user needs based on contextual information, ensuring relevant and timely support during critical decision-making processes.
- 2. Automating mundane tasks:** By taking charge of repetitive and time-consuming tasks, AI copilots allow users to focus their efforts on strategic and creative work, significantly boosting overall productivity.
- 3. Analyzing data:** AI copilots can quickly process large amounts of information, identify patterns and trends, and present actionable insights to drive effective decision-making.
- 4. Enabling seamless communication:** AI copilots facilitate effective interactions with various stakeholders, including employees, customers, and vendors, streamlining

communication processes and reducing delays or misunderstandings.

- 5. Unifying disparate systems:** AI copilots can be the cohesive force connecting multiple platforms, tools, and software applications under one umbrella, ensuring data integrity, accessibility, and compatibility across the enterprise.

In a nutshell, an AI copilot acts to simplify complex tasks and provide valuable guidance and support, ultimately elevating the user experience and driving businesses toward their goals effectively and efficiently. As AI copilots continue to evolve with enhanced capabilities and deeper integration into enterprise ecosystems, they hold the potential to redefine the way businesses operate and compete in the coming years.



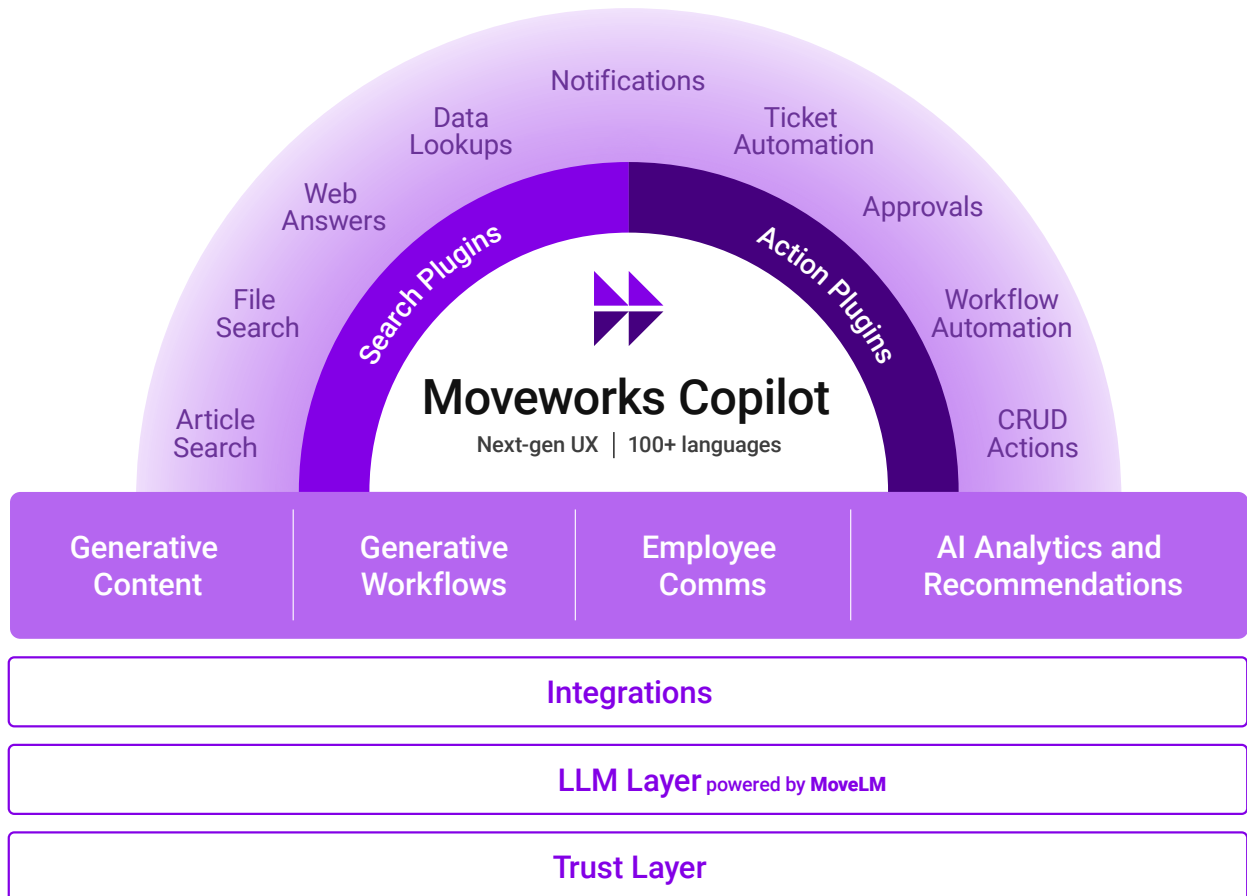
AI copilots are built to connect the enterprise

An enterprise copilot is a fluid conversational interface that connects your employees with every business system. It's built on hundreds of machine learning models, fine-tuned to your enterprise data. Available across every channel and fluent in more than one hundred languages, your enterprise copilot makes it easier than ever for your employees to get things done.

As businesses become increasingly complex and rely on a myriad of software solutions, employees often face the challenge of navigating and managing diverse systems.

Traditional, isolated solutions often fail to resolve cross-system communication problems, which may lead to reduced productivity and inefficiencies. An enterprise AI copilot is the answer to overcoming these challenges.

Integrating all enterprise systems under a single conversational interface allows employees to access information and complete tasks more efficiently. The AI copilot simplifies collaboration, making it easier for employees to excel in their functions and significantly boosting overall productivity.



Understanding the four tiers of enterprise copilots

As AI copilots become essential for automating workflows, enhancing productivity, and remaining competitive across nearly every enterprise function, understanding the options for building vs buying copilots is crucial.

When exploring whether to build or buy an AI copilot, it's crucial to understand the four tiers of copilots and the unique considerations for each:

Tier 1: Basic LLM Integration


Tier-one copilots provide basic conversational abilities through simple API integration of large language models (LLMs).

A tier-one copilot simplifies various everyday tasks by leveraging AI-powered assistance. Some common use cases include:

- Generating content suggestions for social media posts
- Auto-completing email drafts or responses
- Providing answers to frequently asked questions
- Summarizing general content
- Identifying and fixing grammar and style errors
- Performing sentiment analysis on sales call transcripts

Tier-one copilots are relatively easy to kick off, requiring minimal resources and offering a low barrier to entry, making them an attractive starting point for organizations exploring AI tools.

Launching the copilot boils down to accessing an LLM, such as GPT-4. LLM-as-a-service providers, like Hugging Face and OpenAI, make this process even more accessible. You simply subscribe to a reliable API provider for your chosen LLM and integrate their API into your software or platform. This streamlined approach needs little beyond developer resources dedicated to implementing the API integration, ensuring a smooth and cost-effective way to introduce AI-driven assistance into your organization.



Tier 2: Customized LLM implementation

Tier-two copilots are customized LLMs fine-tuned and grounded on an organization's data to improve domain-specific accuracy, requiring more investment in data and machine learning expertise.

A tier-two copilot is better prepared to manage some domain-specific tasks. Some common use cases include:

- Translating IT support tickets
- Drafting an FAQ for the Finance department
- Summarizing legal documents
- Generating industry-specific content for marketing materials
- Assisting with medical diagnosis based on patient symptoms and medical history
- Identifying trends in customer feedback data and providing actionable insights

Developing a tier-two copilot involves a substantial upfront investment in resources and expertise. Key elements include pre-trained models, supporting infrastructure such as GPUs, and a skilled team of developers and machine learning engineers capable of selecting a suitable pre-trained model, like LLaMA, RoBERTa, MPNet, or Flan-T5, and fine-tuning it with the organization's domain data for a more customized LLM implementation.

Annotation is also critical for this more focused approach. And while the annotation process can commence with external service providers like Scale.ai, it's essential to transition in-house as you advance to higher tiers with increased domain specificity and scrutiny.

Although tier-two copilots demand more resources and investment than their tier-one counterparts, the payoff is a high-performance, domain-specific AI solution that tackles your organization's unique challenges.

Tier 3: Advanced LLM pipelines

Tier-three copilots involve chaining multiple LLMs together, creating sophisticated pipelines optimized for multi-step use cases that leverage the strengths and capabilities of each LLM involved and adding significant technical complexity.

By incorporating multiple LLMs and advanced techniques, tier-three copilots can tackle a broader range of use cases, enhance productivity and efficiency, and address challenges in more sophisticated domains. Some common use cases include:

- Analyzing medical literature
- Automating account team follow-up

- Providing multilingual IT and HR support
- Moderating content
- Assisting in the creation and evaluation of financial models and projections
- Analyzing and optimizing supply chain operations for businesses

To get started with a tier-three copilot, there are several key elements and investments you need to consider. First, you must create a multi-LLM stack consisting of various pre-trained models designed to work together for more complex tasks. It is essential to have connectors in place to enable seamless system integrations and facilitate the interactions between different LLMs.

As is also the case with tier-two copilots, investing in fine-tuning is crucial for optimizing the models in the multi-LLM stack. Along with this, system integrations help the copilot assimilate into your organization's existing workflows and automation.

Allocating resources for annotators to generate high-quality domain-specific data will ensure improved performance for your copilot. Likewise, assembling dedicated AI and machine learning teams is vital for developing, implementing, and optimizing the tier-three copilot.

Tier 4: Enterprise-wide LLM adoption

Tier-four copilots are full-fledged enterprise platforms that work to address the challenges inherent in providing extensible employee support and facilitating autonomous decision-making.

As a sophisticated LLM system specifically designed for enterprise-wide deployment, these tier-four copilots encompass advanced features like a reasoning engine, analytics, security, and privacy, as well as out-of-the-box connectors catering to the demanding requirements of large organizations.

A tier-four copilot can handle issues across multiple functions, channels, languages, and departments.

Here are just a handful of examples:

- Providing enterprise-wide support
- Assisting with decision-making by providing insights, predictions, and recommendations
- Monitoring compliance and security
- Managing intellectual property
- Upleveling customer service and engagement
- Curating and generating content across an organization



Accuracy and factuality are of paramount importance when incorporating a tier-four copilot in enterprise settings such as legal, corporate development, or finance departments. Organizations must invest in multiple specialized teams to effectively implement a tier-four copilot, including design, UX, annotation, machine learning, systems integration, and security compliance and privacy infrastructure.

Advanced reasoning techniques further enhance the performance of a tier-four copilot, enabling it to tackle complex problems and improve employee productivity. Organizations can leverage robust language models to optimize their decision-making processes and streamline operations across departments by grounding the copilot to specific use cases.

For Tier 4, buying an end-to-end enterprise copilot solution is overwhelmingly recommended.

Building a production-ready system with enterprise integrations, controls, and support is likely out of reach for most companies. Leaving it to specialized vendors is best.

What makes the Moveworks next-gen Copilot a tier-four copilot?

Moveworks sets itself apart as a tier-four copilot by both harnessing hundreds of machine learning models, specifically fine-tuned to enterprise data, and deeply integrating with the organization's disparate tech stack to fully connect the enterprise ecosystem.

There are six crucial building blocks that enable our copilot to deliver this seamless, frictionless experience:

1. Reasoning

Reasoning is the brain powering our copilot's advanced capabilities. It mimics human cognition to tackle complex challenges that stump traditional, transactional bots. This way, our copilot can strategize solutions and plot the best course of action, just like a support agent might.

When an employee poses a question or request, reasoning kicks into gear. This engine reviews the inquiry and begins an internal dialogue to determine the optimal way forward, much like a person working through a problem. It taps knowledge resources to understand context and history. It engages in conversation with the user to clarify ambiguities in the original request. And it interfaces with various systems to orchestrate actions across enterprise infrastructure.

2. Planning

Planning is the set of steps that the brain devises to complete the task at hand. It's the project plan laid out to take the conversation through to resolution.

Once reasoning has worked out the optimal approach to the employee's request, planning maps out each specific step required to get there. This could involve clarifying questions to ask, systems to search, actions to take, and responses to provide.

Planning carefully sequences each of these steps, ensuring they build off each other effectively and defining any logic or conditions needed to handle different scenarios along the way. Without planning, you're just hoping for the best. This building block grounds our copilot by providing a clear roadmap to the solution, removing guesswork.

By leveraging reasoning to strategize and planning to map out details, our copilot can tackle multifaceted issues smoothly and systematically. The combination of advanced cognition and meticulous planning enables truly intelligent conversations and automation.

3. Action-oriented plugins

The next step is execution — and that's where plugins come in. Plugins are specialized components that actually carry out the planned actions.

You can think of plugins as the hands that do the work. Just as hands have unique fingerprints, each plugin is custom-built to handle specific use cases and integrate with certain systems. For example, plugins may reset passwords in Okta, provision users in Google Workspace, onboard staff in BambooHR, manage access in AccessHub, deploy servers, update Salesforce, pull reports from analytics platforms, and sync product data between ERP systems. The list could go on.

The magic happens when plugins work together. Like a skilled chef using multiple utensils, our copilot combines plugins to unlock tremendous possibilities. They let us automate intricate processes across the systems you already use with Moveworks.

Plugins are the powerhouse that turns thoughtful plans into tangible actions. With these specialized components at our disposal, we can resolve nearly any request via seamless integrations. And the best part — no reconfiguration needed. Plugins leverage your existing setup.

4. Clarifications

Employees rarely provide every detail upfront. Clarifications are critical when the initial request lacks enough context to determine a solution.

Our copilot proactively seeks clarification to enrich vague requests. It can have natural language conversations with employees to ask thoughtful follow-up questions, or it can search integrated

systems to uncover additional helpful information.

Clarifications enable our copilot to dynamically gather the complete picture so it can craft an accurate response. Rather than make assumptions or ignore gaps, it leverages two-way dialogue and system data to fill in the blanks.

Clarifying vague, brief requests allows our copilot to better offer customized solutions. The process of adding more context gives the copilot more information to develop strategies and more specifics to take focused action. Asking for clarification turns unclear questions into clear answers.

5. Summaries

Summaries are the final responses our copilot delivers after executing the action plan.

Without generative AI, you end up with static, rigid responses — just copying and pasting text. Summaries empower our copilot to synthesize dynamic answers in an easy-to-digest way.

The copilot analyzes all the information gathered, including clarifications, system data, and actions taken. It then uses its generative capabilities to summarize the most important details in simple language.

Then, the copilot can tailor the response to that specific employee's original question and context.

They get a summarized explanation that is personalized, focused, and human-friendly.

Summaries unlock the true potential of AI to have natural conversations. Employees receive crafted responses vs. chunky copied text.

6. Citations

Citations are a critical complement to summaries. They list where the summarized information is coming from — be it knowledge base articles, system data, or other sources.

Citations serve the vital purpose of establishing trust and credibility. Employees can instantly see the responses aren't fabricated — they're grounded in real corporate information and facts.

This verification gives employees confidence in the accuracy of the summaries. If our copilot users ever want to dig deeper or learn more, they can trace back to the original sources.

Together, summaries and citations provide the best of both worlds: synthesized responses that are easy to consume yet fully transparent. Citations unlock the ability for generative AI to provide constructive, trustworthy responses. Employees get AI's advantages without losing human oversight and validation. This balancing act is essential to user adoption and continued reliance on virtual support.

7. Security

Moveworks takes a proactive, multi-layered approach to AI security and privacy. We implement technical safeguards like content moderation, fact verification, prompt protection, and query risk assessment. Our systems adhere to enterprise security best practices like sensitive data masking, hardened infrastructure, least privilege access, and identity validation.

Beyond that, our AI models are developed with principles of responsible AI in mind: preserving customer data privacy, promoting fairness, ensuring safety, and maintaining robust accountability and governance practices. We leverage AI security and privacy by design across our development lifecycles.

Moveworks holds multiple compliance certifications and meets rigorous enterprise security and privacy standards. We treat customer trust as our highest priority, employing comprehensive guardrails to mitigate risks and uphold ethical AI operations.

With guardrails in place, customers can feel confident allowing our copilot to fully automate processes and resolve issues on their behalf. Employees get best-in-class virtual support while organizations maintain control over security posture and data governance.



Key considerations for selecting AI copilots

When selecting an AI copilot, it's crucial to evaluate the platform based on specific factors that contribute to seamless implementation and overall performance. As discussed in the previous section, there are different copilot tiers, each with different core considerations and capabilities.

When evaluating AI copilots, focus on platforms built specifically for enterprise needs. Key selection criteria include:

- **Enterprise context:** The copilot should understand company data, systems, and workflows to deliver relevant, accurate assistance. Seek copilots trained on proprietary enterprise data.

- **Security and compliance:** Rigorous security protections, access controls, and compliance with regulations like HIPAA and SOC 2 are essential. Prioritize enterprise-grade copilots designed for sensitive data.
- **Scalability and integrations:** As use cases and data volumes grow, the copilot must integrate with hundreds of business systems. Opt for an AI copilot that can scale as your business grows, providing a robust and adaptable solution.
- **Reasoning and learning:** Look for copilots that continuously learn from new data to handle evolving tasks. Advanced reasoning capabilities are key for complex issues.

- **Extensibility:** Modular copilots that allow expanding to new use cases through mix-and-match plugins offer flexibility as needs change.
- **Steerability:** Advanced controls give admins visibility into how the copilot operates and the ability to guide its actions, ensuring alignment with policies and governance.
- **Transparency:** Detailed audit trails, confidence scores, and citation data provide visibility into the copilot's logic for trust and verification.

- **Flexible LLM architecture:** Access to customize and chain the latest and greatest LLMs for each use case optimizes performance and cost-efficiency for the long term.

Prioritizing these enterprise-focused capabilities will ensure your organization selects a scalable, secure, and adaptable copilot that drives real business value.



The future of AI copilots

As AI copilots continue to improve and evolve, collaboration between humans and machines will become more seamless, leading to greater productivity and enhanced problem-solving capabilities.

By bridging the gap between various enterprise systems and offering contextually relevant assistance to users, AI copilots serve as powerful allies in the modern, fast-paced business environment.

As you consider implementing an AI copilot for your organization, it's essential to understand its distinctive features, the value it can bring to your operations, and how it differs from other AI-driven solutions, such as chatbots and virtual agents.

Ultimately, the key to harnessing the full potential of AI copilots lies in selecting a platform that aligns with your organization's unique requirements and is built on robust security, scalability, and learning capabilities.

By choosing the right AI copilot, an organization can fuel its growth, transform its operations and ensure a smoother, more productive journey for both employees and customers.



02

The challenges of building an AI copilot

At Moveworks, we know firsthand the monumental effort required to build an enterprise-class copilot from the ground up. For years, our team has invested significant time, talent, and resources into developing a full AI platform that meets the highest standards for security, scalability, and performance across complex enterprise use cases.

In this chapter, we'll share insights into the challenges we faced and overcame in building the Moveworks [next-gen Copilot](#), so you can understand the level of investment required to create an enterprise-grade solution.

By sharing our experience, we aim to provide transparency into the blood, sweat, and tears required to build a secure, robust, and extensible AI copilot tailored for the enterprise. This inside perspective reinforces why leveraging an established vendor can help most organizations realize value quickly while benefiting from years of knowledge and investment.



Questions to ask before your start

Building your own AI copilot can seem tempting, but it's important to weigh the pros and cons before jumping in.

Here are two key questions to consider:

- **What are you wanting to achieve with an AI copilot?** What pain are you solving for? Be clear on your goals and how AI will help you get where you want to be.
- **With those goals in mind — would it make more sense for you to make API calls or get an enterprise-wide solution?** Leveraging an existing API service may be more cost and time effective than developing something yourself,

depending on your use case. If you just need basic chat capabilities, an existing public model might suffice. But if you need something that can deal with more specialized domains or take action across multiple systems? You're going to need a more comprehensive solution.

As we explore the challenges of building your own copilot in this chapter, carefully consider these two questions to determine if building your own solution is the right approach for your needs and resources.

The undertaking is substantial, so clarity of goals is essential.

Let us look at a number of points that should be considered when pondering whether to build or buy a copilot.



Challenge 1: Internal resources required to build an AI copilot

When evaluating building versus buying an AI copilot, assessing the internal resources required to construct an enterprise-grade solution in-house makes clear why purchasing is often the most prudent option.

Here's why:

- **Distraction from core competencies:** If your core competency is not building AI, then buying an AI solution from a vendor is the most sensible thing to do. Otherwise, you need to allocate extra resources on it and do a significant number of other activities too. It might also involve making your employees move away from doing their most meaningful work.
- **Team capabilities and headcount:** Developing an enterprise-class AI copilot necessitates a multidisciplinary team with expertise in data science, engineering, and machine learning. The team should be capable of not only understanding complex enterprise use cases but also devising innovative solutions to address them using the latest and greatest in AI tech.
- **Infrastructure requirements:** A robust AI copilot demands a substantial investment in specialized infrastructure, including powerful hardware capable of handling intricate computations and cloud services that provide the necessary scalability and flexibility.

- **Domain expertise:** Having individuals with domain knowledge is critical. Understanding the intricacies of the enterprise use cases, business processes, and user interactions across teams and departments is fundamental to developing an AI copilot that seamlessly integrates into the daily workflows of your organization.
- **User experience (UX) design:** A user-friendly and intuitive interface is vital for user adoption. Having UX design expertise within the team is essential to create an AI copilot that not only functions effectively but is also easy for end-users to interact with and understand.

These substantial internal resourcing requirements reinforce why purchasing a mature enterprise AI copilot is the strategic choice for most organizations without years invested in developing AI-specific teams and infrastructure. Building their own solution often spreads resources thin across non-core areas.

Moveworks' strategic internal advantages

Building a world-class, enterprise-grade AI copilot in-house is an immense undertaking that requires substantial resources and expertise outside of most companies' core competencies.

Purchasing an AI solution from a dedicated vendor like Moveworks provides several strategic advantages:

- **Stay focused on your core business:** Developing AI capabilities internally runs the risk of distracting from the work that truly drives your organization's success. By buying an AI copilot, you avoid taxing resources on non-core activities.
- **Gain from established AI expertise:** Moveworks has assembled a specialized team with decades of deep and cutting-edge expertise in data science, machine learning, natural language AI, annotation, and enterprise domain knowledge. Replicating this cross-functional AI skill set in-house is extremely difficult.
- **Leverage turnkey infrastructure:** Our solution provides powerful machine learning operations infrastructure purposely built for training, hosting, and scaling AI models.
- **Benefit from AI domain expertise:** With years of experience deploying our copilot across global enterprises of many shapes and sizes, Moveworks brings unparalleled understanding of automating business processes, workflows, and service experiences.
- **Drive user adoption a conversational UX:** Our conversational user experience makes it easy for employees to get what they need, ensuring a delightful and productive experience that encourages employee adoption.

By purchasing an enterprise-proven solution, your team can focus valuable resources on leveraging AI capabilities aligned with your core business,

rather than building the underlying AI technologies from scratch.

Challenge 2: Technical considerations for building an AI copilot

Developing an enterprise-grade AI copilot requires overcoming complex technical obstacles across many dimensions. Deep expertise in areas like machine learning, data engineering, and infrastructure architecture is requisite.

Here's why:

- **Steerability and control:** Advanced controls and oversight are critical to ensure the copilot's actions align with policies and governance. Mechanisms must be implemented to provide visibility into the copilot's decision-making and allow authorized users to guide its behavior without compromising reliability or safety.
- **Transparency and verifiability:** Detailed audit trails, citations, and explanations are essential to provide transparency into the copilot's responses. For use in regulated enterprises, the copilot must ground its answers in available data sources rather than hallucinations.

- **Handling model drift:** The dynamic nature of enterprises means the copilot's performance can degrade over time as patterns shift. Adaptive algorithms and continuous monitoring are required to detect and mitigate model drift.
- **Customization and extensibility:** Enterprises have unique systems, workflows, and requirements. Designing the copilot to be customizable and extensible to accommodate specific use cases is a technical challenge.
- **Integrations and interoperability:** Seamless integrations with hundreds of HR, IT, and other systems are needed to unlock end-to-end automation.



- **Scalability and flexibility considerations:** The copilot should seamlessly scale to handle increased workloads and adapt to diverse enterprise environments. This necessitates technical solutions that balance performance, security, and accessibility across various infrastructures.
- **Sourcing high-quality annotated data:** Acquiring and curating real-world enterprise data needed to fine-tune and ground a large language model is a time-consuming process that demands meticulous attention to detail. The team must dedicate thousands of hours to collecting and labeling data to ensure the copilot's effectiveness across diverse scenarios.
- **Omnichannel:** To match user expectations, the copilot needs to handle conversations across various channels like chat, email, voice assistants.
- **Multilingual support:** For global enterprises, the copilot must support conversations in multiple languages.

Taken together, these technical obstacles reinforce why purchasing a mature, enterprise-ready AI copilot is often the most prudent path forward for an organization looking to leverage AI, rather than taking on the immense challenge of building a production-grade solution themselves. Very few companies have the expertise and resources to construct an enterprise copilot from scratch without significant tradeoffs in quality or capabilities.

Moveworks' strategic technical advantages

Moveworks provides a continuously improving, enterprise-grade AI copilot solution that overcomes the immense technical hurdles that challenge in-house development efforts.

Our state-of-the-art technology:

- **Scales with flexible architecture:** The Moveworks Platform is powered by two core engines that enable it to scale effectively: a generative reasoning engine driven by large language models (LLMs), and a robust action engine. Working together, these engines allow the platform to parse incoming requests, communicate naturally with employees, interface with other systems as needed, and dynamically devise tailored solutions. This flexible, modular architecture provides the scalability and adaptability to handle increasing workloads and evolving enterprise needs over time.
- **Ensures controlled, steerable results:** Our prompt engineering techniques allow authorized users to precisely steer the copilot's actions through careful input design, ensuring alignment with policies and governance. You maintain full oversight and control.
- **Verifies results:** Moveworks provides detailed audit trails, citations grounding responses in data sources, and clear explanations of the copilot's reasoning. A fact-checking model validates outputs against queries.
- **Continuously adapts to change:** Our AI incorporates feedback loops and continuous monitoring to detect and adapt to changing data patterns, preventing model drift from impacting performance over time.
- **Extends to meet your needs:** Our copilot works within your unique workflows, using plugins to accommodate evolving use cases and incorporating information from the systems that you use.
- **Automates end-to-end:** Moveworks has pre-built integrations across hundreds of HR, IT, and business systems, enabling true end-to-end automation out-of-the-box.
- **Leverages quality training data:** Moveworks has invested thousands of hours into curating enterprise data across service domains like IT, HR, facilities, and more. Our data annotation team follows rigorous guidelines to precisely label queries, responses, and metadata. This provides a rich, constantly-growing dataset for model fine-tuning.
- **Engages across channels:** Moveworks' copilot provides a seamless, unified experience by understanding conversations in their native context across channels like chat, email, and more.

- **Supports global languages:** Our copilot engages in fluent conversations across over 100 languages, seamlessly switching between languages within the same dialogue. Models are customized on localized enterprise data to ensure accurate understanding of languages worldwide.

With deep machine learning expertise and enterprise-hardened technology, Moveworks provides a production-ready, self-improving AI copilot solution that solves the toughest technical challenges.

Challenge 3: Security and privacy implications of building an AI copilot

Developing an enterprise-grade AI solution that handles sensitive company and customer data necessitates rigorous security and privacy precautions. This represents a substantial undertaking for organizations attempting to build their own bot.

Here's why:

- **Vulnerabilities and data privacy concerns:** Building an in-house copilot introduces potential vulnerabilities if security is not baked into the design from day one. Teams must proactively protect against risks like data leaks, account hijacking, disinformation spreading, and toxic content generation.
- **Precautions needed around data access:** Copilots built internally also require meticulous access controls around company data. The

team must implement granular permissions, encryption, anonymization, and monitoring to prevent exposure of confidential information like financials, IP, or PII.

- **Extensive security testing and auditing:** Rigorous penetration testing, red team exercises, and third-party audits are essential to validate the copilot's defenses. Few companies have the dedicated resources to ensure adequate code review, infrastructure hardening, and vulnerability management.
- **Ongoing maintenance and patching:** After launch, the team must dedicate resources to continuous security monitoring, updates, and fixes. Falling behind on patches or ignoring emerging threats leaves the copilot vulnerable to compromise over time.

Purchasing a copilot from an established vendor allows relying on their years of security investment and expertise. Maintenance, auditing, penetration testing, and meeting compliance requirements all become the vendor's responsibility rather than an internal burden.

For most companies beyond technology giants, buying rather than building their own copilot is often the more prudent path to achieve enterprise-grade security and privacy. Attempting to match specialist vendors' mature safeguards is an immense challenge requiring substantial dedicated resources.

Moveworks' strategic security advantages

Moveworks provides an enterprise-grade AI copilot solution with robust security baked in from the ground up, allowing you to confidently automate while preserving data privacy and compliance.

Our secure-by-design approach handles key aspects like:

- **Rigorous security practices:** Moveworks engages in rigorous security testing including SAST analysis, vulnerability scans, and penetration testing by third parties. We follow security best practices like the OWASP Top 10 to prevent risks like data leaks or account hijacking.
- **Proven compliance:** Our solution meets the highest global security and privacy standards, achieving certifications like ISO 27001, ISO 27017, ISO 27018, ISO 27701, and SOC 2 Type 2 attestation and more. You can deploy with confidence.
- **Data access controls:** Granular access controls, data masking, encryption, and monitoring govern how the AI copilot interacts with your sensitive data. Moveworks enforces least privilege principles and tight logging oversight.
- **Secure architecture:** Our multi-tenant microservices architecture enables secure and performant data isolation. Containerization allows elastic scaling while maintaining high security posture across hybrid environments.

- **Continuous security:** Moveworks manages all security updates, patches, and ongoing monitoring for you, ensuring the copilot's defenses never lapse against evolving cyber threats over time.
- **Responsible AI practices:** Our AI models are developed from the ground up following Responsible AI tenets like preserving data privacy, ensuring fairness and safety, and maintaining robust governance processes.

By purchasing an enterprise-proven AI solution, you get a secure copilot while offloading the immense burden of building, testing, and maintaining rigorous security guardrails in-house. Moveworks provides the AI automation benefits with the highest standards of security and compliance.



Challenge 4: Financial implications of building an AI copilot

The budget required to build an enterprise-grade AI copilot from scratch is substantial, while buying from a vendor offers faster time-to-value and lower costs.

Here's why:


- **Major upfront investment:** Constructing a copilot in-house necessitates large upfront investments in specialized infrastructure, engineering talent, compute power, and extensive data annotation. Significant financial resources are required to develop, train, and refine high-performing natural language models.
- **Implementation:** Implementing a self-built AI copilot is an immense undertaking, requiring extensive work to collect and annotate data, select and train AI models, integrate with existing systems, and drive user adoption through comprehensive enablement.
- **Ongoing expenses:** In addition to initial costs, maintaining a custom-built copilot demands significant ongoing expenses for continuous improvements to the model, infrastructure, and skill sets. Dedicated staffing, cloud services, licensing, and security audits represent perpetual costs.
- **Extended timeline:** Realizing value from an internally built copilot takes years due to long development cycles and the challenges

of refining accuracy. The extensive timeline results in a prolonged period before users experience benefits.

- **Adoption:** Driving widespread adoption of a self-built AI copilot is challenging, as employees may resist an unfamiliar tool built in-house without extensive user experience expertise. Lack of adoption leads to employees reverting to older systems, diminishing the potential ROI.
- **Opportunity costs:** Tying up budget, talent, and focus on developing core AI capabilities often distracts from an organization's core business goals. The opportunity cost of not allocating resources towards higher-impact initiatives is substantial.

In contrast, purchasing an AI copilot from an established vendor provides quicker time-to-value, lower costs, and enterprise-grade capabilities out of the box. The vendor has already invested in R&D, infrastructure, and skills development so customers avoid overheads.

For most organizations, buying is often the more prudent financial decision compared to building in-house unless AI is their primary focus. Partnerships with vendors who provide enterprise readiness on day one offer faster ROI at lower risk.



Moveworks' strategic financial advantages

Developing an accurate, enterprise-grade AI copilot in-house requires considerable upfront investment in data, infrastructure, and specialized AI expertise that may distract from your core business.

Moveworks eliminates these financial barriers by:

- **Avoiding massive upfront costs:** Building your own copilot necessitates major outlays for compute infrastructure, engineering teams,

data annotation, and model training before any value is realized. With Moveworks, you get a production-ready solution without these upfront sunk costs.

- **Accelerating time-to-value:** Our AI solution is designed for fast deployment using pre-built skills and integrations tailored to meet enterprise needs. You achieve automated support and productivity gains in months, not years.
- **Optimizing operating expenses:** We handle all the ongoing expenses like cloud hosting, software licensing, model retraining, and security audits. You only pay for Moveworks rather than staffing and maintaining it yourself.
- **Maximize ROI through adoption:** Moveworks applies AI/UX best practices honed from deployments across leading enterprises. As a result, we're able to prioritize delightful user experiences that promote high adoption and maximize your ROI.
- **Avoid distraction and opportunity costs:** With Moveworks, you avoid the opportunity costs of tying up resources on non-core artificial intelligence development cycles. Your teams can stay focused on driving business impact.

By purchasing an enterprise AI solution, you achieve all the benefits of intelligent automation while minimizing costs and accelerating time-to-value versus undertaking the enormous endeavor of building an AI system from the ground up.



The strategic case for buying over building an AI copilot

As illuminated throughout this chapter, developing an enterprise-grade AI copilot in-house is an immense undertaking filled with complexity. From talent requirements to technical obstacles and security considerations, constructing a production-ready solution requires years of focused investment, resources, and expertise.

For most organizations beyond major tech giants, attempting to build their own copilot from the ground up typically proves infeasible or spreads resources too thin. The myriad challenges explored make clear why buying is often the most prudent strategic decision.

By purchasing from an established vendor with years focused on enterprise AI, organizations can avoid the pitfalls and costs of internal construction. Vendors have already tackled the hard problems like security, scalability, customization, and accuracy. Their accumulated real-world data, infrastructure,

and skill sets enable providing an enterprise-ready solution from day one.

Buying allows enterprises to concentrate resources on their core business goals and domain expertise rather than underlying AI capabilities. While some integration work remains, buying transfers the heavy lifting to partners who live and breathe enterprise AI development.

Unequivocally, for organizations seeking to leverage the power of AI copilots, buying is generally the fastest and lowest risk path to impact. Vendor solutions avoid the distractions and costs of internal build efforts. Partnerships allow leveraging existing models and years of refinement.

Unless AI is your primary competency, purchasing an enterprise-grade copilot is often the most prudent decision strategically. The challenges covered make clear why buying can propel your organization ahead, letting specialists handle the complexities of enterprise AI behind the scenes.

03

How to avoid AI investment pitfalls

Say it with me: AI is no longer just nice to have. It's as important a building block for your organization as any other.

AI has become so ubiquitous that Forrester estimates that by 2025 [100% of companies will leverage it](#). AI's economic contributions have become so vast, PwC forecasts that AI will contribute [\\$15.7 trillion](#) to the global economy by 2030. The momentum behind AI is accelerating at a breakneck pace, and more are taking notice — haven't you heard of [ChatGPT](#)?

AI has emerged from merely a technology concern amongst data scientists and machine learning wizards. All of a sudden, discussion of large language models (LLMs) has become commonplace in boardrooms, where leaders deliberate over the right conversational AI strategy for their organizations.

As a business leader, you can't afford not to be evaluating AI use cases.

The opportunity cost is too high, and the potential benefits are too significant. However, in this economy, you also can't afford to make big investments in AI unless they pay off. The problem is, as research shows, more than [90% of AI projects fail](#).

Clearly, you need AI, but there's no room for experimentation or educated guesses. You're stuck between a rock and a hard place. So what do you do? experience is rooted in a strong belief.



Now is the time to invest in an AI copilot

When evaluating AI for your organization, you first should identify the business objectives you seek to execute against. This is something that should be done collaboratively between your Technology (IT, Data Science, etc.) and Business teams. Ensuring this alignment is critical.

Second, **define the specific use cases** to implement and the level of effort, technology, people, process,

and data capabilities involved.

And third, **establish the success criteria and KPIs** that you'll use to measure outcomes. Having a framework in place to both assess the value opportunity and measure the value realized will help inform where to invest initially and over the long run.

For this chapter, we'll focus on the third step, walking you through a framework for evaluating the impact of AI on your organization.

Aligning AI investments to business outcomes

Countless research studies highlight the importance of tying technology investments to measurable business outcomes. We've all heard the saying, what doesn't get measured, doesn't get done. We take that same approach at Moveworks, grounding our AI value framework in three business outcome categories:

1. Improved operational efficiencies: The ability of an organization to effectively and efficiently utilize its resources (time, people, capital) to achieve its goals.

2. Enhanced experiences: Delivering delightful, optimal experiences across your employees, partners, and customers.

3. Accelerated business transformation: Fundamentally changing how your organization operates to improve its performance and competitiveness.


Once you define your business outcome categories, consider establishing impact areas and associated KPIs to baseline, target, and measure against over time. These will help inform your initial investment, gauge the investment's success over time, and re-inform additional investment as needed.

Improve Operational Efficiencies



Effectively and efficiently utilize enterprise resources (time, people, capital) to achieve goals

Enhance Experiences



Deliver delightful, optimal experiences across your employees, partners, and customers

Accelerate Business Transformation



Fundamentally change how your organization operates to improve its performance and competitiveness

Moveworks' AI value framework is grounded in three key business outcomes: improving operational efficiency, enhancing experiences, and accelerating business transformation.

When working with our customers, we collaboratively define KPIs that map to eight common business impact areas that Moveworks' conversational AI platform delivers against:

- **Improved service desk efficiencies:** Making the service desk more effective while reducing costs.

Implementing new technologies such as AI, automating processes, and optimizing existing technology infrastructure (hardware and software) give support resources more time to focus on higher-value projects, reduce the time and effort required to resolve IT issues, improve the quality of service provided to end users, and increase employee satisfaction, leading to cost savings for the organization and a more productive environment for the workforce.

- **Reduced manual processes:** Replacing manual, time-consuming tasks with automated or digital solutions.

By automating certain processes, you can eliminate the need for manual intervention altogether. Organizations capable of reducing the time and budget required to complete basic tasks free up staff to focus on more valuable, strategic work. In addition, automated processes are often more reliable and consistent than manual ones, leading to improved quality and employee and customer satisfaction.

- **Increased tech stack ROI:** Getting the most value out of your technology investments, from your knowledge bases to your ITSM, to your HCMs, CRMs, ERPs, and more.

By improving the performance and efficiency of the technology stack, organizations can reduce change management costs and increase productivity in the long term, resulting in increased competitiveness, revenue, and profitability.

- **Improved employee productivity:** Increasing the amount of work or output produced by employees within a given period.

Quickly providing employees with the resources and support they need to excel in their roles, automating or streamlining tasks, and improving communication and collaboration among teams all increase profitability per employee.

- **Improved employee onboarding:** How new employees are introduced to a company and its policies, procedures, and culture is deeply connected to their – and their company's – success.

The right tools provide new employees with the knowledge, skills, and support they need to succeed from day one. By helping new employees feel comfortable, confident, and engaged in their work, overall job satisfaction and performance improve, ultimately benefiting the organization.

- **Reduced employee churn:** A decrease in employees who leave an organization.

Employee churn, also known as turnover or attrition, can be costly for businesses, leading to a loss of institutional knowledge, reduced productivity, increased hiring and training expenses, and ultimately reduced profitability. Organizations can improve retention, reduce turnover costs, and create a more stable and engaged workforce by reducing employee churn.

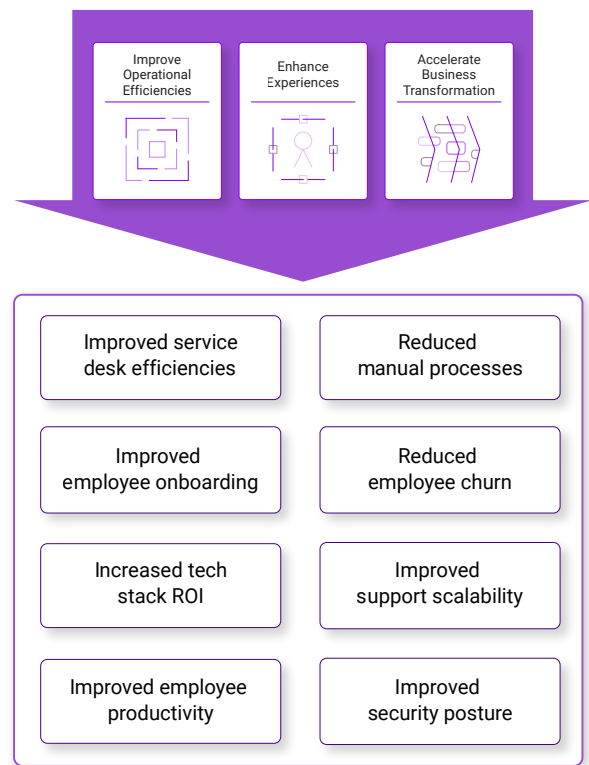
- **Improved support scalability:** The ability of a support team to effectively handle an increasing number of customers or users without a corresponding increase in resources.

Automating routine tasks, implementing self-service support options, and using data and analytics to identify and prioritize high-impact initiatives and projects all contribute to a support organization’s ability to grow. By improving support scalability, organizations can provide better support to their customers and users, even as the number of people they serve increases.

- **Improved security posture:** The overall effectiveness of an organization’s security measures in protecting its assets and data.

A strong security posture is essential for businesses, as it helps prevent systems and software vulnerabilities that can have serious consequences. To improve security posture,

organizations can implement various measures, such as robust password policies, regularly updating software and security protocols, and providing regular training to employees on security best practices. Improving security posture has numerous benefits, including reducing costs, reducing reputational risk, and protecting revenue.



Moveworks breaks down each business outcome into eight measurable impact areas. Know the value of AI in your organization.

Identifying the business objectives, defining the use cases and capabilities required to deliver against those business objectives, and establishing a value framework with specific KPIs to inform initial and ongoing investment are three important steps to take when evaluating AI for your organization.



Know the value of AI in your organization

Investing in AI requires a strategic approach tied to clear business outcomes. Rather than making haphazard technology bets, organizations should:

- Identify the specific business objectives AI can help execute against. Get alignment between technology and business teams on where AI can drive impact.
- Define use cases that deliver on those objectives, evaluating the capabilities and effort required. Prioritize high-value AI applications.
- Establish an evaluation framework with measurable KPIs mapped to business impact

areas. Use data-driven metrics to inform investments and track outcomes over time.

With this outcome-driven approach, organizations can make informed decisions about AI investments based on real business needs and measurable results. They can demonstrate ROI, course-correct as needed, and expand AI usage to new areas based on value delivered.

The days of AI experimentation are over — strategic alignment of AI to business goals is imperative in today's economic climate. By following an outcomes-based methodology, organizations can confidently invest in AI copilots knowing that the time, effort and cost will pay dividends.

The costs of building vs buying an AI copilot

So, you need a solution that doesn't yet exist within your business — one that will transform workflows, enhance analytics, and enable new offerings.

You have a critical decision to make: build a custom solution from the ground up, or buy a ready-made option from a vendor.

When considering whether to build or buy an AI copilot, one of the most critical factors is cost.

Developing a robust enterprise-grade copilot from scratch requires extensive investment — not just in dollars, but also in time and technical resources.

In this chapter, we will break down the key costs involved in building your own AI copilot solution versus partnering with an established vendor. We will cover:

- Industry insights on building versus buying
- The pros and cons of building a copilot
- Understanding what it takes to build a copilot-like experience in your enterprise
- Navigating the hidden costs of a DIY approach
- The importance of adopting a buy-first approach for AI copilots
- Tips to ensure the success of your AI copilot


By exploring these key factors, you'll gain insight into the substantial investment required for an in-house build versus leveraging external AI expertise. We'll provide real-world data on costs so you can make an informed decision on the best path for your organization.

What the data tell us about building versus buying

In the ever-evolving landscape of technological endeavors, the journey into machine learning (ML) projects has been marked by both triumphs and setbacks. According to Gartner, a sobering statistic emerges: 85% of AI projects meet an untimely demise. The sunk cost fallacy, an idea that sticking to an in-house project despite escalating costs and diminishing returns is the only path forward, can lure organizations into a trap of unrealized potential.

Consider the human element in this digital landscape alone. Glassdoor reveals that the average salary for a software engineer in 2023 stands at \$106,227, with top-tier engineers commanding well over \$300,000. A profound revelation surfaces: the annual expense of purchasing software can often be eclipsed by the annual salary of the engineers needed for an in-house venture.

Forbes adds its voice to the conversation, emphasizing the efficiency gained by working with reputable providers deploying proven solutions and underscoring the futility of reinventing the wheel and advocating for leveraging the wealth of experience that vendors bring to the table. **Why spend countless hours architecting a solution that already exists and is proven in the market?** The benefits of best practices, derived from numerous successful implementations, await those who opt for pre-made solutions.



Diving deeper into the realm of IT project management, the [CHAOS Report's findings](#) are nothing short of eye-opening. A staggering 31.1% of projects fail. Furthermore, a significant 52.7% of projects end up costing a staggering 189% of their original estimates. [The Standish Group](#), in its extensive survey, classifies projects into three resolution types, revealing that the overall success rate is a modest 16.2%, with challenged projects making up 52.7%, and the rest being impaired or canceled.

Yet, amidst the statistics painting a picture of challenges and pitfalls, a beacon of hope emerges: the promise of artificial intelligence (AI). AI, with its transformative capabilities, augments and automates cumbersome manual processes.

For example, AI holds the potential to relieve support ticketing burdens. As evidenced by a

staggering [70% of industry leaders](#) in agreement, AI is set to significantly change how companies create, deliver, and capture long-term value. The [estimated savings of \\$15,000](#) per IT worker per year, once empowered with effective AI for their work, underscores the profound impact AI can have on operational efficiency and cost-effectiveness. This paradigm shift toward embracing AI not only streamlines processes but also reflects a strategic alignment with the evolving landscape of technological innovation.

At Moveworks, we stand on the precipice of this technological frontier, recognizing the potential that AI holds. It has been our guiding premise since our inception – to harness the power of AI to not just meet challenges but to exceed expectations and redefine what's possible in the realm of copilots and beyond.



Pros of building a copilot	Cons of building a copilot
In-house control over development; Tailored to specific needs and requirements.	Requires skilled in-house development team; Potential delays in project timelines due to development complexities and challenges.
Customizable to the organization's unique language requirements.	In-house implementation may struggle to match the extensive language support provided by established external solutions.
Greater control over resource development and customization.	Potential higher upfront costs for hiring skilled developers and allocating resources. Longer development timelines compared to using ready-made solutions.
More control over licensing agreements and costs.	Higher upfront costs for development and licensing. Ongoing maintenance and updates may require continuous financial investment.
Customizable to specific infrastructure needs; Full control over security measures.	Requires a significant initial investment in hardware and infrastructure. Ongoing maintenance and scalability challenges may arise.
Full control over operations, maintenance, and enhancements.	Resource-intensive. May require ongoing training and skill development for internal teams; Risks falling behind in features and capabilities compared to dedicated solutions.

Weighing the pros and cons of building a copilot

When considering a copilot, the decision to build in-house or opt for a pre-made solution carries significant implications.

Purchasing a pre-made copilot presents the allure of immediate deployment and potential cost savings, aspects often overlooked when delving into the complexities of creating custom software. The in-house development process can encounter delays as the scope expands and the copilot continuously adapts to unforeseen modifications.

Choosing a pre-made copilot means acquiring a product crafted by experts. Successful copilot vendors leverage technologies and feedback from a diverse array of businesses, ensuring the delivery of

a superior product. However, managing resources in the copilot development process can be as intricate as the depths of the development itself.

It's crucial to acknowledge that there is no one-size-fits-all copilot, and pursuing the perfect solution may lead to indecision or disappointment. Attempting to build an all-encompassing copilot that handles every task is unlikely to excel in any individual function without expert support.

Vendors of pre-made copilots, with a vast client base, may provide integrations with your existing software at a relative cost advantage. Depending on the copilot type, you may tap into cutting-edge machine learning and AI technologies exclusive to large user-based software.

What it takes to build a copilot-like experience into your enterprise

Embarking on the journey to integrate a copilot-like experience into your enterprise requires a nuanced understanding of the foundational elements.

While tools like ChatGPT serve general purposes, transitioning to a fully embedded enterprise copilot demands a meticulous approach.

Unlocking the capabilities of a copilot to handle enterprise-focused tasks necessitates addressing key components:

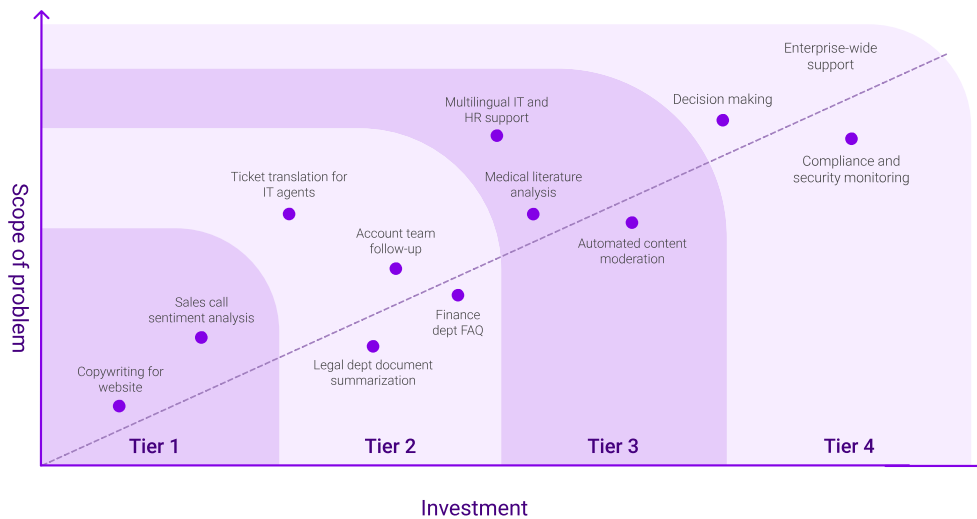
- **Domain-specific requirements:** Understanding the unique needs, challenges, and objectives of each department involved in the copilot is crucial. That includes a wide range of enterprise service needs — identity, permissions, security, compliance, and more.
- **Relevant data inputs:** The key to building a copilot is to gather and process relevant data specific to each domain. This requires a team of [annotators](#) to provide tags, descriptions, ratings, translations, and well-researched examples so the copilot can provide the best possible support.
- **System integrations:** Connect the copilot to the right tools, technologies, and systems for each department, ensuring seamless collaboration across domain-specific applications.
- **Custom AI models:** Fine-tune the LLMs according to domain-specific data sources,



- desired outcomes, and unique challenges specific to each domain for more accurate decision-making.
- **Strong analytics and data science team:** Staff with expertise in data analysis and machine learning development are required to support the copilot's learning and optimization and ensure that the experience improves over time instead of degrading.
- **Consistent, intuitive user experience:** To maintain a straightforward and seamless interaction with the copilot, you might need a [conversational AI system](#) that lets users easily take full advantage of its capabilities.
- **Robust security infrastructure:** Especially for an enterprise copilot, engaging with security experts to safeguard sensitive enterprise data, protect user privacy, and adhere to necessary compliance measures is mandatory.

Given the above, it should be clear that incorporating a copilot experience into your organization takes considerable effort. While some use cases might be lighter lifts, others will require heavy, long-term maintenance. The scope of the problem you're trying to solve will require varying degrees of investment. For example, using AI to provide enterprise-wide support will, by definition, require a much better understanding of and investment in the bullets above than using AI to write web copy.

This is to say that an AI copilot is not something you can build and maintain on your own, particularly if you want to expand its functionality across the enterprise. Investing in the right partner with the right resources is the key to successfully implementing a helpful and efficient AI copilot.



When looking at AI copilot use cases, we start to see the relationship between the scope of the problem and the investment required to solve it, enabling decision-makers to make informed choices when embracing LLMs for their organizations.

Navigating the hidden costs of doing it yourself

Investing in technology — including AI copilots — a fundamental cost continuum emerges, compelling decision-makers to balance the expenses of development time against licensing costs.

This is to say that regardless of the chosen model, a financial investment is inevitable — either upfront in development or through ongoing licensing fees. The more extensive the development requirements, the lower the upfront costs tend to be. Conversely, opting for a pre-built copilot solution often incurs higher licensing fees but substantially reduces development expenses.

Here are some considerations:

- **Copilots are complex:** Building copilots is a complex task that demands considerable time and expertise. It becomes evident that utilizing pre-made copilot solutions is advisable whenever possible.
- **ROI concerns:** The success of a copilot hinges on its utility to users. If the copilot lacks usefulness or faces low adoption rates, the organization struggles to achieve a satisfactory return on investment (ROI).
- **Long-tail applications:** Understanding how to apply and having a long-term investment in optimizing these solutions is critical. Organizations making continuous optimizations and automations find that each change requires

more time cumulatively, especially for edge cases, to ensure ongoing usefulness to end users.

- **Interaction-based charges:** Scaling charges based on copilot usage can present challenges for finance teams, leading to unpredictable and potentially higher-than-expected costs.
- **Training time impact:** Prolonged training times diminish user goodwill, highlighting the importance of a fast onboarding process for user retention.
- **Professional services dependency:** Many organizations opting for a DIY copilot approach require extensive professional services due to the lack of in-house expertise, turning a seemingly “free” copilot solution into one that costs significantly more than anticipated.

Many organizations have spent months developing an AI solution, only to retract it due to inadequate recognition capabilities — underscoring the risks and complexities of the DIY approach. Navigating the financial landscape of copilots demands strategic decision-making to ensure not just a cost-efficient solution but one that aligns with organizational goals and user satisfaction in this evolving technological frontier.

Why you should have a buy-first approach for AI copilots

When it comes to acquiring an AI copilot for your business, taking a buy-first approach by purchasing an enterprise-ready solution from a vendor should be the default strategy for most organizations. Building a copilot in-house from scratch requires massive investment and is extremely high-risk.

There are substantial cost advantages both upfront and over time when buying rather than building a copilot solution. Purchasing a copilot eliminates unpredictable research and development expenses, while outsourcing the complex integration, maintenance, and enhancement work to experts.

Total cost of ownership (TCO) is an estimate of all the direct and indirect costs involved in acquiring, deploying and operating a product or system over its lifetime. Let's analyze the key elements of TCO to understand why buying beats building for AI copilots:

- **Development and implementation:** The upfront engineering costs in year one to develop the initial copilot and use cases
- **Translation services and multilingual support:** The ongoing cost of translating answers surfaced by the copilot into other languages
- **Resource development:** The continuous cost of creating new knowledge articles, forms, and catalog items

- **Software licenses:** Any annual software license fees for a copilot platform
- **Hardware and infrastructure requirements:** The yearly hardware and cloud infrastructure costs
- **Ongoing operations, maintenance, and enhancements:** The recurring cost of maintaining, fixing issues, and enhancing the copilot

When evaluating TCO, buying a proven enterprise copilot is overwhelmingly more cost-efficient than building one yourself across planning, acquisition, integration, maintenance, upgrades, and retirement. AI copilots are a complex software problem — leave it to the experts.



Cost	Building a copilot	Buying Moveworks
Copilot development and implementation	Internal copilot development and implementation costs required. Costs may include hiring or allocating internal resources for development.	No copilot development required. Implementation cost included.
Translation services and multilingual support	Translation services and multilingual support Internal investment required for translation services or additional support. May involve hiring multilingual support staff or implementing language support tools.	No additional cost.
Resource development	Higher costs associated with traditional bot toolkits for resource development. Internal development resources may be more expensive.	Up to 90% less cost compared to bot toolkits.
Software licenses	Internal development and licensing costs; May require purchasing licenses for software components separately.	Moveworks licensing fee.
Hardware and infrastructure requirements	Internal investment required for hardware and infrastructure; Costs may vary based on the chosen on-premises solution.	Fully hosted solution; No additional cost.
Ongoing operations, maintenance, and enhancements	Internal resources needed for ongoing operations, maintenance, and enhancements; May require multiple FTEs for these activities.	Approximately 1 FTE, 25% of their time.

Tips for ensuring the success of your AI copilot


Choosing and implementing an AI copilot successfully requires a strategic approach that incorporates essential tips to navigate potential challenges and maximize efficiency. Here are valuable insights to ensure the success of your solution:

- 1. Establish a dedicated team:** Building a team focused on the copilot's development and optimization is crucial. This dedicated team should comprise experts from various organizational departments, fostering cross-organizational collaboration and ensuring diverse perspectives and needs are considered.
- 2. Prioritize user experience:** User experience is paramount in the success of any solution, but it is especially important for disruptive solutions, such as AI. Find a copilot with a user-centric approach, emphasizing simplicity and effectiveness. Regularly gather user feedback to make continuous improvements and refine the copilot's capabilities based on real user interactions.
- 3. Avoid reinventing the wheel:** Leverage existing solutions and technologies whenever possible. Integrating out-of-the-box solutions can significantly streamline the development process, reduce costs, and accelerate the copilot's time to market. Focus on customization where necessary, but embrace

proven technologies to enhance efficiency.

- 4. Plan a practical rollout:** Planning a practical and phased rollout is essential. Prioritize functionality and ensure the copilot's core features are robust before making grand promises or expanding its scope. A step-by-step approach allows for adjustments based on real-world performance and user feedback.
- 5. Measure what matters:** The success of your copilot should be measured by its impact on users. Develop key performance indicators (KPIs) that align with user satisfaction, task completion rates, and overall efficiency improvements. Regularly assess and refine these metrics to gauge the copilot's ongoing success.

Keeping an eye on these tips ensures a thoughtful and effective approach to the deployment and enhancement of your AI copilot. By fostering collaboration, prioritizing user experience, embracing existing solutions, planning pragmatically, and focusing on user-centric metrics, your organization can navigate the complexities of AI copilot implementation with confidence and achieve lasting success.



Choose the right path forward for your business

In the ever-evolving realm of AI, the choice between crafting an in-house AI copilot and procuring an established external solution holds far-reaching implications for organizations. This chapter meticulously unraveled the intricate web of costs, spotlighting the profound differences between internal development and the utilization of proven external alternatives.

Our goal is to steer organizations away from the pitfalls of the DIY approach. By dissecting financial

intricacies, unveiling hidden costs, and emphasizing real-world data, the facts advocate for a strategic “buy-first” mindset and underscores not only the fiscal efficiency but also the strategic wisdom of tapping into the expertise of established vendors, ensuring the seamless integration of copilots into organizational workflows.

As industries navigate the transformative potential of AI, this chapter serves as a guiding beacon, urging decision-makers to choose a path that not only minimizes costs but strategically aligns with the expertise needed for AI copilots to thrive in the dynamic landscape of enterprise solutions.

05

Successful AI copilots in action

In the dynamic landscape of enterprise operations, AI copilots have emerged as a transformative force, streamlining workflows and enhancing user experiences.

We would like to shine a spotlight on the real-world success stories of organizations leveraging Moveworks' AI copilot technology. From Broadcom's rapid IT issue resolution to Wellstar's instant ROI with Moveworks, and Palo Alto Networks' remarkable time savings during the shift to remote work, these cases exemplify the tangible benefits copilots bring to diverse industries.

Keep reading for all the juicy details...

Broadcom's employees get help in <60 seconds

Broadcom, a global infrastructure technology company with offices worldwide, has addressed its IT integration woes brought on by rapid growth and acquisitions using AI. The company had to deal with information scattered across 10 different knowledge bases which made it challenging for employees to find accurate information. As employees had to rely on a growing number of SaaS applications, efficiently answering IT queries was becoming critical.

To rectify this, in September 2018, Broadcom launched 1.Bot, the company's internal name for the Moveworks Copilot. Through their enterprise chat platform, employees could resolve a variety of issues like resetting passwords, unlocking accounts, creating tickets, checking ticket status, and finding solutions from knowledge base articles.

Moveworks' combination of advanced conversational AI and natural language understanding (NLU) enabled 1.Bot to resolve IT support issues conversationally, directly in chat, understanding the employee's problems and providing precise answers they need.

As a result of integrating the knowledge base using Moveworks' AI platform, 57% of Broadcom's IT issues are now solved in less than a minute. The solution not only enables real-time visibility over knowledge base gaps but also allows employees to self-serve their support. Thus, achieving Broadcom's vision to integrate the help desk

experience for its employees, save time, and reduce the volume of work by automating tasks with AI.

“You don't want your users going all over the place to search for information. The advantage of Moveworks is that there's one familiar go-to place for employees to get their answers.”

Stanley Toh
Head of End-User Services & Experience, Broadcom

Wellstar's copilot delivered ROI in days

Wellstar, the largest healthcare system in Georgia, transformed its employee experience with an AI solution that works out of the box. Looking to extend its focus on specialized support to its employees, Wellstar faced the challenge of choosing the right automation solution.

Instead of building a bespoke solution using a chatbot “toolkit”, which would have demanded significant upfront work, they chose Moveworks.

In May 2020, Wellstar deployed its Moveworks Copilot, known internally as WALi, achieving ROI on day one. The copilot understood company lingo, automatically synced with backend systems, and adapted to user feedback, which allowed the IT team to focus on other priorities.

Wellstar's adoption of WALi ensured that employees got the help they needed through natural conversations, responding to each request with the most relevant solutions in a matter of seconds.

With this people-first approach, Wellstar refrained from manually building a bot and succeeded in launching an instant ROI-delivering copilot that genuinely helps their employees do their best work. This AI solution aids from frontline medical professionals dealing with tech problems to office-based administration personnel with policy queries, easing the burden on employees across the board.

“I’ve built bots, and I’ve bought bots. Unless you have a huge development team, I recommend you buy. That way, you can get all the expertise, the knowledge, the wisdom, and the talent of the people involved.”

Larry Ross
ServiceNow Manager, Wellstar

Palo Alto Networks saved its hybrid workforce 351,000 hours with AI

Palo Alto Networks rolled out its Moveworks Copilot, Sheldon, at a pivotal moment, for both the company and the entire planet. At the start of the pandemic, employees needed to get comfortable working from home, which created a flood of IT requests, HR requests, policy questions, and more.

Nicole Tate-Pappas, Senior Director of Digital Experience Management, remembered telling the team that “Moveworks needs to go live now!” She found that implementing the copilot was “incredibly simple” and “happened without a lot of womanpower from us.” Sheldon began solving issues on day one, attributed to Moveworks’ understanding of enterprise language.

“Our deployment was extremely fast,” Steve Januario, Former VP of Digital Employee Experience, said. “You can get an impact from Moveworks in under a month, if your team is aligned on the goal.”

However, the Palo Alto Networks team was most impressed with Moveworks’ ongoing commitment to customer success, long past the day the copilot went live. Moveworks takes complete ownership over its results, meeting regularly with the team to review metrics, discuss new potential use cases, and identify knowledge gaps.

In fact, the IT help desk went from 10 knowledge articles to more than 250 “just like that,” according to Tate-Pappas, because of the visibility and guidance that Moveworks provides.

“In terms of the experience for employees, there’s nothing even close to Moveworks. It’s the best, period.”

Steve Januario
Former VP of Digital Employee Experience, Palo Alto Networks

No more password resets: Why Procore brought AI to IT

Procore launched its Moveworks copilot, known internally as Bolt. An AI solution that lives on Slack, Bolt provides one extremely simple place for employees to get help, whether they’re locked out of their account, need access to software, or want to edit an email group.

Top IT teams now recognize the importance of automation in delivering faster tech support, given the shortcomings of a manual approach. The question then becomes whether to build such an automated solution in-house using a chatbot “toolkit.”

However, Steve LeBoeuf, Senior Director of IT, found that “these tools require a small army to get online” and then typically struggle to gain adoption from employees. What he and his team needed was a solution that worked immediately, using AI they didn’t have to manage themselves.

“Moveworks starts right out of the box, bringing the automation we need to fill in the gaps,” explained LeBoeuf. “I don’t have to program or be responsible for anything. We leave that work to the experts.”

Moveworks arrived at Procore already speaking the language of IT. When an employee has a question, by definition they don’t know the answer, and as a result, knowing what they mean can be a challenge. But by understanding the employees’ underlying intent, Bolt can connect the symptoms they describe to the most relevant solution available.

“Procore is in the business of construction—not building the best AI solution in the world. That’s why we use Moveworks.”

Steve LeBoeuf
Senior Director of IT, Procore

We'll leave you with this...

AI copilots have the power to radically transform how work gets done across your organization.

By connecting all your systems and data under one conversational interface, copilots automate tedious tasks, provide on-demand expertise, and unlock entirely new levels of productivity and efficiency.

But as powerful as copilots are, choosing the right one successfully requires careful planning and the right approach for your specific needs. Hopefully, this guide has given you a clear framework for evaluating the build vs. buy decision.

If you go the build route, you now understand the technical challenges and investment required to create an enterprise-grade copilot from scratch. It's a massive undertaking.

For most companies, buying a proven copilot solution will likely be the faster, more cost-effective path. With a market-leading product, you get all the sophisticated AI models and deep system integrations ready to go. That means you can start realizing unmatched productivity gains and competitive advantages much sooner.

Whichever route you choose, the time to add an AI copilot to your workforce is now. Employees and executives alike are demanding the kinds of experiences copilots provide. Those who resist will quickly fall behind more innovative competitors.

Don't let that be you. Be a leader who empowers your people with the future of collaborative AI.

Let them focus on creativity, strategy, and impact — while the tedious tasks get handled instantly. That's the promise of AI copilots.

