Moveworks

4 big shifts IT teams are making with AI for 2024

Intro

2023 proved to be a breakthrough year for AI. ChatGPT's debut sparked a frenzy of predictions about AI's impending impact. Investors foresee trillion-dollar stock market gains. Entrepreneurs envision automating a growing percentage of jobs. Thought leaders liken slowing AI's progress to "murder."

But past the hype, how can IT leaders pragmatically ride the AI wave in 2024?

This ebook spotlights four priority shifts by IT teams at the forefront.

By the end, you'll know how to:

- Benchmark your IT team's performance against industry leaders
- Stay ahead with the latest AI tools and technology
- Establish best practices for generative AI in knowledge management
- Architect your stack to fully leverage AI capabilities

This guide is essential for:

- Heads of Digital Workplace (Director+)
- Heads of End-User Services (Director+)
- Heads of Service Desk (Director+)
- CIOs and their direct reports
- Any IT leader eager to exploit AI's potential

The future of AI is here. Don't get left behind. Discover the moves you need to make to unleash AI's power across your organization in 2024.

Table of Contents

01

How industry-leading IT teams use AI to improve employee experience

02

New ways of working and connecting with AI in 2024

03

Generative AI best practices for knowledge base management

04

Architecting your enterprise with an AI foundation

01

How industry-leading IT teams use AI to improve employee experience

Providing fast, seamless support should be the top priority for every IT help desk. But how can you objectively measure the speed and efficiency of your team? Benchmarking metrics like mean time to resolution (MTTR) against top performers.

Analyzing over 200 enterprises, we found a striking gap between average and exceptional IT teams. The best-in-class organizations resolved tickets in less than half the time of industry medians. This dramatic difference is primarily driven by their adoption of cutting-edge AI and automation.

The data reveals a clear imperative: to achieve elite service levels, IT leaders must accelerate their use of AI across the employee support process. Copilots, chatbots, virtual agents — call them what you want — the simple fact is that machine learning tools are enabling 24/7 self-service.

IT teams that leverage these innovations to work smarter can slash resolution times, boost CSAT scores, and exceed employee expectations. The future of excellent service begins with AI. By benchmarking metrics and processes against the best, IT help desks can continuously improve operations and provide five-star support.

> Analyzing over 200 enterprises, we found a striking gap between average and exceptional IT teams. The best-in-class organizations resolved tickets in less than half the time of industry medians.

How long are employees waiting for support?

In this report, we provide help desk benchmarks for high-touch issues, access requests, troubleshooting questions, approvals, and overall MTTR. Across the board, we see a growing divide between the companies that have leaned into innovation and those that haven't.

High-touch issues submitted to the help desk

Employees with high-touch issues depend on subject matter experts for solutions. Fixing a broken printer or reconfiguring Salesforce requires more than a single help desk agent.

When it comes to high-touch issues, you might assume that automation might hurt, not help. Instead, we've found automation is particularly crucial when solving high-touch issues.

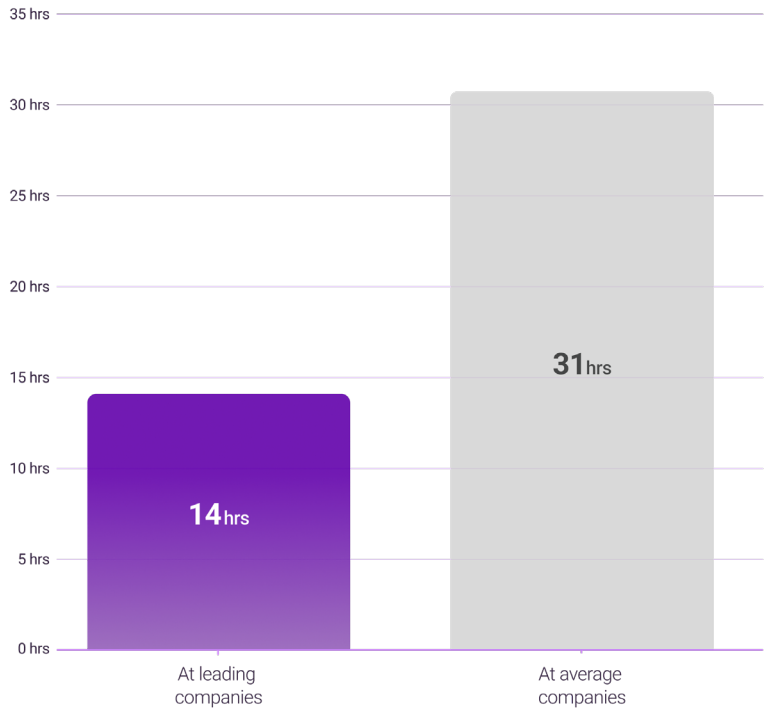

*This data represents a sample with selected types of issues.

When employees submit a high-touch issue, how long are they waiting for support?

The difference in wait time between average and leading organizations is nearly 42 hours. Why?

Best-in-class organizations triage issues to the right person, improve stakeholder communications and accelerate comments. But this can take time to manage. With AI, however, companies can take much of the friction out of solving these complex functions.

To help employees, you must first understand what they're asking for. Using NLU (natural language understanding) and NLP (natural language processing) models, industry-leading companies can significantly speed up complicated processes, including triage.

Access issues

There are several use cases — password resets or account unlocks, for example — where it might take an agent just minutes to perform the task. Still, these tickets sit in the queue for hours. Industry-leading companies use this as an opportunity to introduce a back-end automation tool to trigger workflows.

The difference is shocking: a 12-hour discrepancy that employees experience whenever they're locked out of an account.

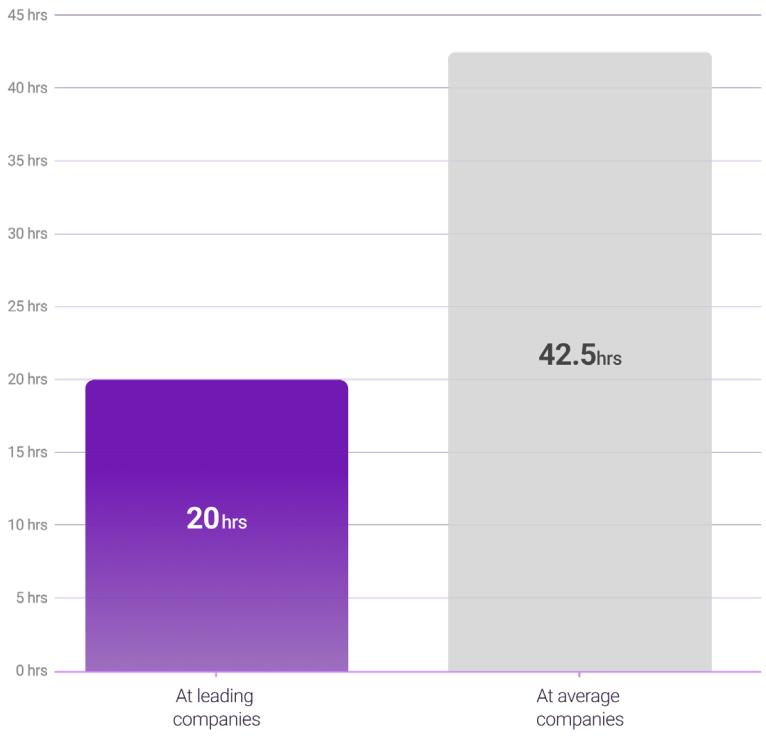

How long do employees wait for access?

In the age of hybrid work, not all employees can walk over to their favorite help desk agent and describe their problems in seconds. Leading companies use AI to transition that experience seamlessly to a digital workplace.

Troubleshooting questions

All help desks have a collection of knowledge base articles documenting solutions that only specialists at the company would know. Innovative leaders use AI to transform these long, hard-to-reach articles into bite-sized solutions.

In conventional organizations, it takes a lot of time and effort from the help desk agent to offer a customized solution to an employee, impacting wait times.

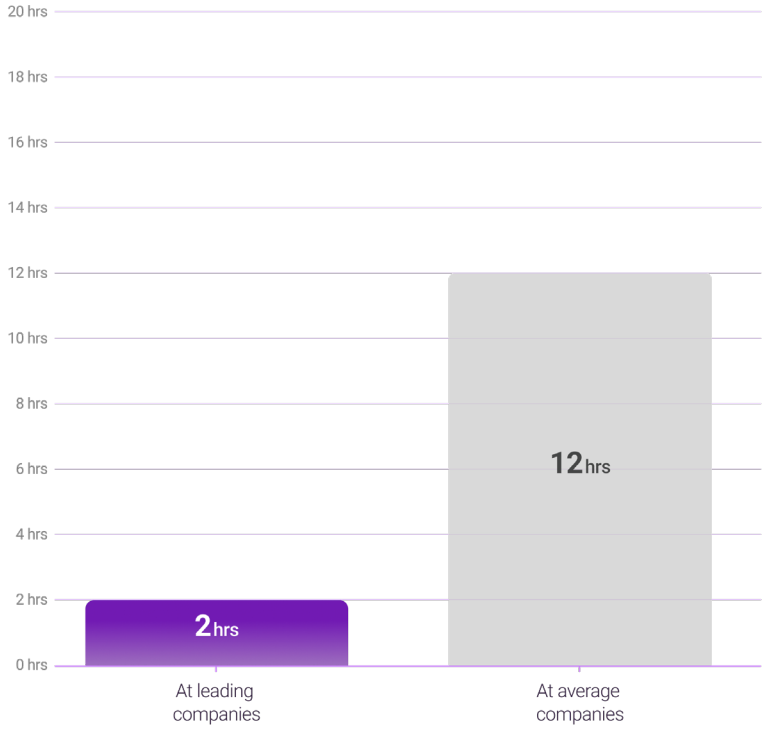

How long do employees wait to get their questions answered?

There's a 13-hour difference in wait time between companies that leverage automation and those that don't.

Employees expect personalized solutions, and IT agents spend hours gathering information before coming up with the perfect answer. At leading companies, this process is automated — AI is used to gather all context around employee requests so help desk agents can have a 360-degree view while answering questions.

The impact of approvals

Some use cases need managerial input. But getting a manager's attention can be challenging at times.

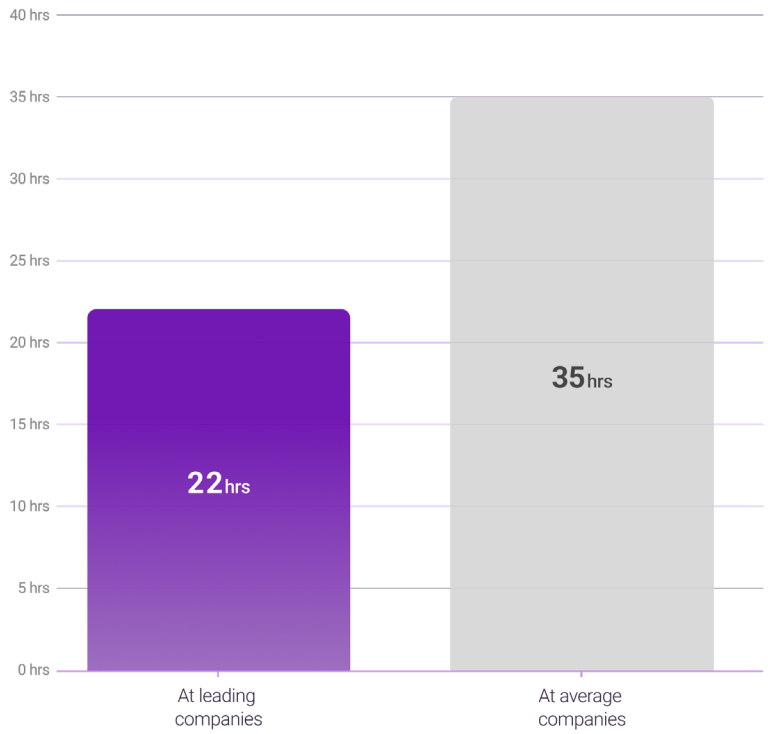

How long are employees waiting for approvals?

Employees wait nearly 12 hours for approvals due to the primary mode of communication: email and siloed support channels. But automated support happens in real time. In this category, leading organizations can attack the root of the problem: poor communication.

Industry-leading companies have incorporated conversational AI to bring these requests to an enterprise chat platform, allowing managers to give approvals as soon as they see their employees' messages.

To move from an average company to one that's leading the pack, you will have to prioritize the employee experience. Conversational AI is a powerful tool with limitless applications. It can help your help desk agents eliminate their pain points and significantly reduce employees' wait times.

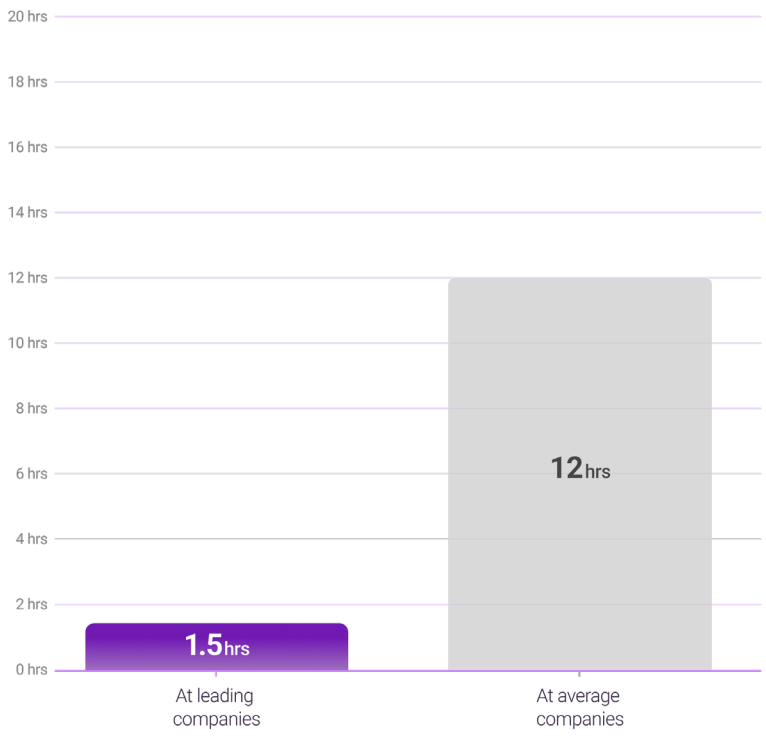

02

5 new ways of working & connecting with AI in 2024

AI is radically transforming the workplace. In 2024, leading organizations are leveraging five key AI capabilities to drive efficiency and productivity across their organizations.

From AI-generated knowledge bases that effortlessly scale to copilots that augment human capabilities, to seamless conversational access to back-end systems, AI is fundamentally changing how we get work done. By adopting a thoughtful approach to workplace AI, your company can maximize efficiency, productivity, and satisfaction across global hybrid teams in 2024. The future of work is here.

01

Enterprise-focused large language models are powering 24/7 self-service at a global scale.

02

AI copilots are emerging as the new face of workplace technology, delivering a conversational user experience to employees.

03

No-code platforms are making it easier than ever to build natural language interfaces across a wide range of use cases.

04

Natural language analytics tools are providing actionable insights from unstructured data.

05

Security and responsible AI must be proactively prioritized across all AI deployments.

Enterprise LLMs unlock 24/7, global, self-service

The emergence of large language models (LLMs) like GPT-4 into the popular consciousness heralded a new era in AI. However, these general-purpose models also have significant limitations when applied to complex business settings. **This gap gave rise to a new category of AI: enterprise LLMs specialized for workplace tasks.**

Leading organizations increasingly focus on finding or building custom LLMs tuned for specific use cases rather than generic solutions. These enterprise LLMs excel at specialized domains like customer service, IT support, and knowledge work. They provide the accuracy, speed, and reasoning required to make a cross-business impact.

However, adopting LLMs requires prudent strategies. A diverse ensemble of models mitigates overreliance on any single LLM. Customization with in-house data improves performance on unique workflows. And responsible development and testing minimize risks. With the right approach, enterprise LLMs can transform operations.

Driving business impact with enterprise LLMs

The key advantage of enterprise LLMs is driving tangible improvements in critical workflows:

Cost reduction: By automating repetitive tasks like customer service and IT support, enterprise LLMs significantly cut costs. For example, they can handle common inquiries 24/7 without human agents, allowing managers to reallocate staff to higher-value work.

Faster resolution: LLMs provide instantaneous responses no matter where someone works, resolving many issues in seconds rather than days through manual processes. Real-time support accelerates workflows.

Improved accuracy: Specialized training enables higher accuracy on domain-specific tasks than general models. LLMs can master industry terminology and processes.

Global scale: LLMs remove language barriers by providing native support in multiple languages, allowing for consistent worldwide experiences without localization needs.

Customizing LLMs for enterprise needs

Off-the-shelf LLMs struggle with niche processes and terminology. Customization closes this gap.

For example, Moveworks built MoveLM™ using proprietary enterprise data and algorithms specialized for workflows related to employee support to great success. In Moveworks' testing, MoveLM outperformed other models on domain-specific benchmarks.

Effective customization requires:

- Tailored model architecture and algorithms
- Workplace-specific training data
- Rigorous comparative benchmarking
- Iterative tuning based on qualitative and quantitative insights

This fine-tuning creates LLMs purpose-built for business requirements rather than general usage. Combined with an ensemble approach, customized enterprise LLMs drive maximum impact.

AI copilots act as a conversational UX for the enterprise

With the emergence of AI copilots, organizations now have conversational interfaces that understand natural language to help users complete tasks and access information across complex enterprise environments.

AI copilots are becoming core enterprise technologies, serving as the new face of workplace technology. Rather than navigating complex apps and tools, employees can simply chat with their AI copilot to get work done faster.

The value AI copilots bring to businesses and customers

AI copilots have changed the way businesses and customers interact with various enterprise systems, offering several advantages that lead to improved operational efficiency and satisfaction levels.

Unified enterprise systems: AI copilots connect employees to every business system, simplifying interactions through natural language processing. Leveraging machine learning models, these intelligent assistants facilitate swift access to enterprise data, enhancing employee productivity as they perform tasks more efficiently.

Omnichannel support: AI copilots adapt to the preferred channels of users, providing a seamless conversational interface across platforms like Slack, Microsoft Teams, email, or web portals. This ubiquity ensures consistent, uninterrupted support for all employees, regardless of their preferred workspace.

Multilingual support: With the capability to communicate in many different languages, AI copilots provide unparalleled global support to users in their native tongue, regardless of their location. This multilingual functionality enables businesses to cater to diverse employee demographics more effectively.

Connecting backend systems across departments: By integrating with a wide range of business systems, AI copilots can assist with various use cases across different departments. From resetting passwords to managing PTO requests and expense submissions, AI copilots bring an all-encompassing solution to any business need.

Enhanced creativity and productivity: By automating mundane tasks and providing quick access to information, AI copilots allow employees to focus on higher-value tasks, fostering creativity, and boosting overall productivity.

Upleveled employee skills: With AI copilots as reliable resources, employees can continually sharpen their skills and knowledge base, thus driving career development and contributing to a more skilled workforce.

Cost savings: AI copilots reduce the need for additional support staff, automate routine tasks, and minimize the time wasted on searching for information, resulting in significant cost savings for businesses that invest in such advanced technologies. By harnessing the immense potential of AI copilots, businesses can streamline operations, enhance workforce efficiency, and create a more satisfying work experience, while customers benefit from rapid, tailored support and communication to address their needs effectively.

AI copilot use cases across the enterprise

Clearly, adopting natural language conversations unlocks significant value across the enterprise. Conversational AI solutions like copilots boost productivity and satisfaction by:

- Handling typical service desk, HR, and other queries, accelerating resolution times
- Enabling real-time data queries via natural language to surface insights from systems and documents
- Providing instant, personalized support across channels and languages
- Replacing manual data entry with automated natural language data capture
- Delivering summarized insights from surveys, reviews, and other feedback
- Sending relevant, targeted communications proactively through preferred channels
- Simplifying complex enterprise workflows

The common thread is using natural language to make key processes faster, smarter, and easier. Conversational interfaces boost productivity, empower users, and unlock new sources of value across the enterprise.

With AI copilots, employees can focus on high-value activities rather than manual busy work. A conversational UX provides a frictionless interface for workplace productivity.

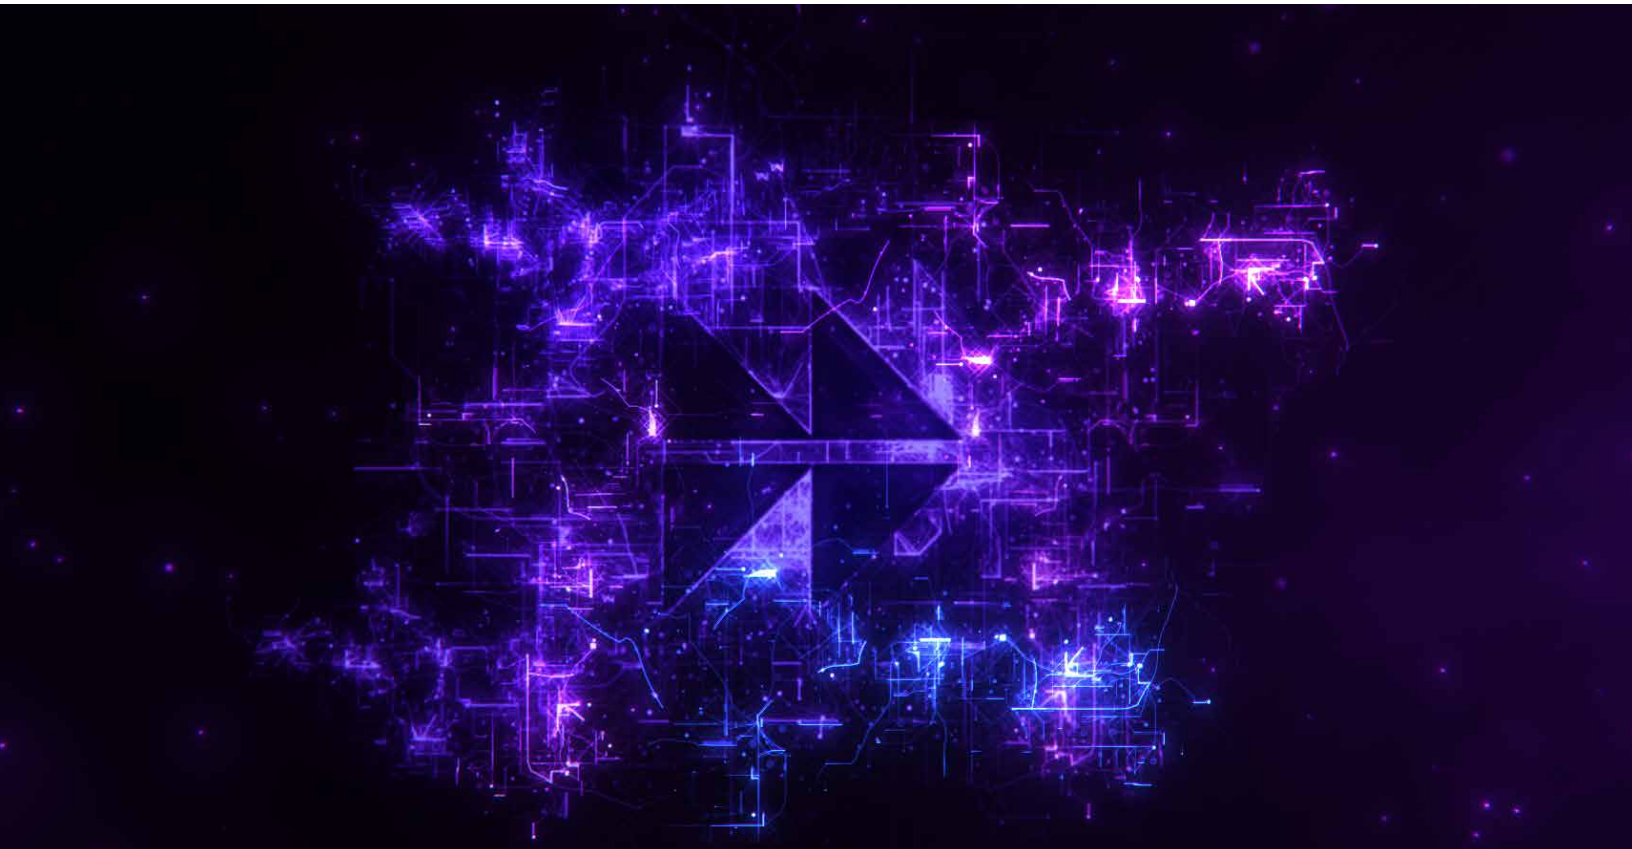

No-code tools make building conversational AI use cases simple

The rise of no-code platforms is changing how enterprises create and deploy conversational AI. **Rather than relying on technical experts, no-code tools allow anyone to build customized conversational interfaces and connect employees to data and systems through natural language.**

With intuitive drag-and-drop interfaces, non-technical users can integrate backends, define conversational flows, and launch AI tailored to their needs. Streamlined workflows replace complex coding, accelerating development from months to minutes.

Solutions like Moveworks' Creator Studio exemplify the power of no-code for conversational AI. With Creator Studio, users simply describe the capabilities they want in plain language. Advanced generative AI handles the complexity behind the scenes, generating the required training data and dialogue flows automatically.

By democratizing development, no-code conversational AI unlocks a world of use cases, including:

- Seamless access to data and insights across enterprise systems
- Automated handling of IT, HR, and other support requests
- Proactive notifications and reminders about key events
- Personalized customer service and sales engagement
- Contextual recommendations to aid decision-making

With no-code as the engine, conversational AI can expand from narrow chatbots to organization-wide AI transformation efforts. Ultimately, no-code platforms enable any workflow to become conversational, providing a natural language interface to every enterprise system.

AI-powered analytics pull insights from natural language

The rise of digital workplace collaboration tools and SaaS apps has led to an explosion of unstructured enterprise data. Crucial insights often get buried in the noise of natural language conversations and documents.

AI-powered analytics solutions are bringing order to this complexity. By leveraging natural language processing and machine learning, these tools can rapidly analyze unstructured text data to uncover actionable insights.

For example, Moveworks' Employee Experience Insights (EXI) ingests ticket conversations and documentation to identify priorities for improving IT and HR services. EXI extracts key topics and trends from natural language input to spotlight areas like repeated HR policy questions or common IT hardware issues.

Powered by advanced large language models, EXI goes beyond surface-level reporting to provide meaningful analysis from unstructured data. The benefits include:

- Identifying unseen inefficiencies in workflows
- Measuring the business impact of new initiatives
- Discovering pain points in employee experiences
- Prioritizing projects based on data-driven insights
- Simplifying reporting with AI-generated summaries

By deriving intelligence from natural language, AI analytics enable organizations to optimize operations, boost productivity, and improve employee satisfaction. Unlocking unstructured data is key to data-driven decision-making.

Security and responsible AI are clear priorities

As AI becomes more prevalent across the enterprise, ensuring its secure and ethical use is paramount. So, to fully embrace the power of AI, it's essential to understand and address the risks it may introduce into an enterprise environment. Only by mitigating potential threats can we ensure successful — and secure — implementation.

Organizations must take steps to reduce risks and biases. To secure AI systems, key tactics include:

- Strict access controls over data and models
- Ongoing monitoring for anomalies or misuse
- Testing for potential vulnerabilities like data poisoning
- Deploying defenses like multi-factor authentication
- Partnering with cybersecurity experts for audits and advice
- Maintaining human oversight and control

Special care should be taken with workplace AI that accesses sensitive systems and data. Strong data governance, access controls, and compliance processes are essential.

Ensuring responsible and ethical AI

Organizations must also proactively address risks like bias, fairness, and transparency in AI systems through:

- Careful and representative data collection
- Testing models for unwanted biases
- Providing explainability into AI decision-making
- Enabling human oversight of outcomes
- Developing inclusive teams and processes
- Partnering with experts in ethics and governance

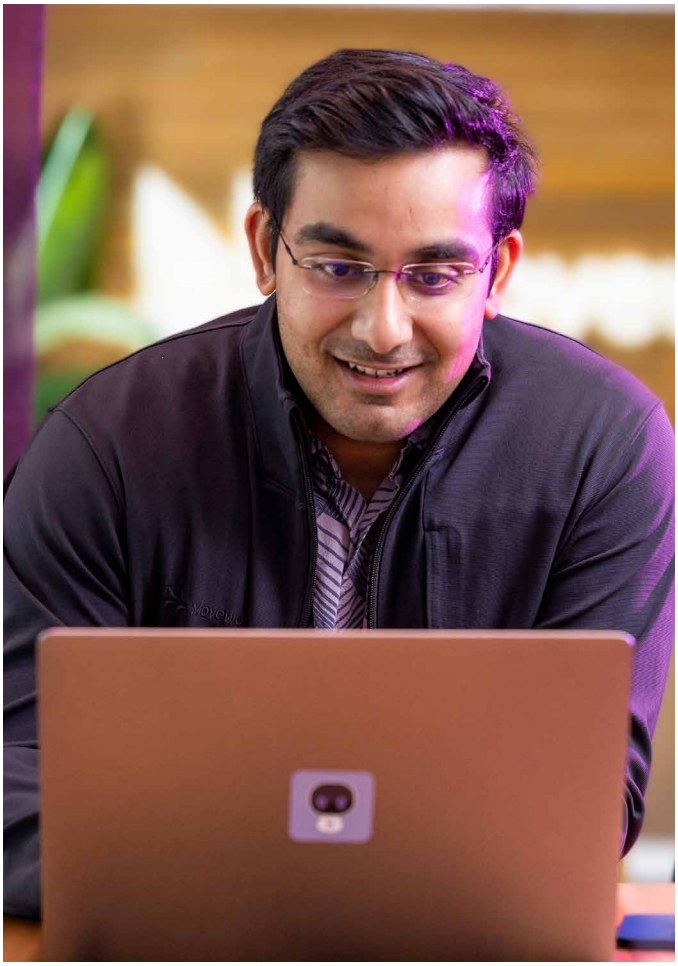

The large language model lifecycle

LLMs have forever changed how we interact with artificial intelligence. However, understanding their lifecycle and potential risks is crucial for safe, effective deployment.

The lifecycle of an LLM can be broadly categorized into four major stages, each playing a pivotal role in shaping the model's capabilities and performance:

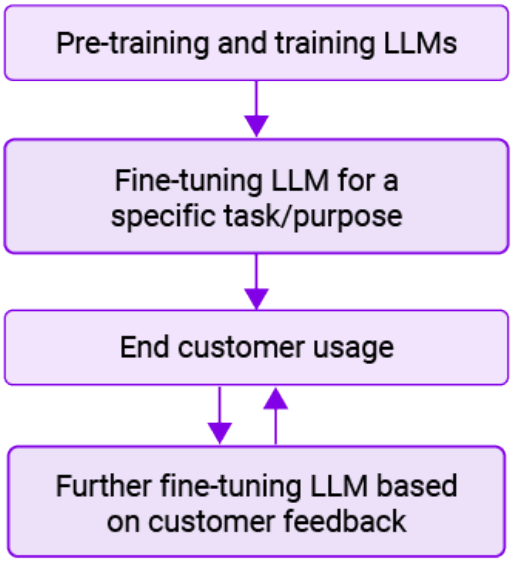

The large language model (LLM) lifecycle

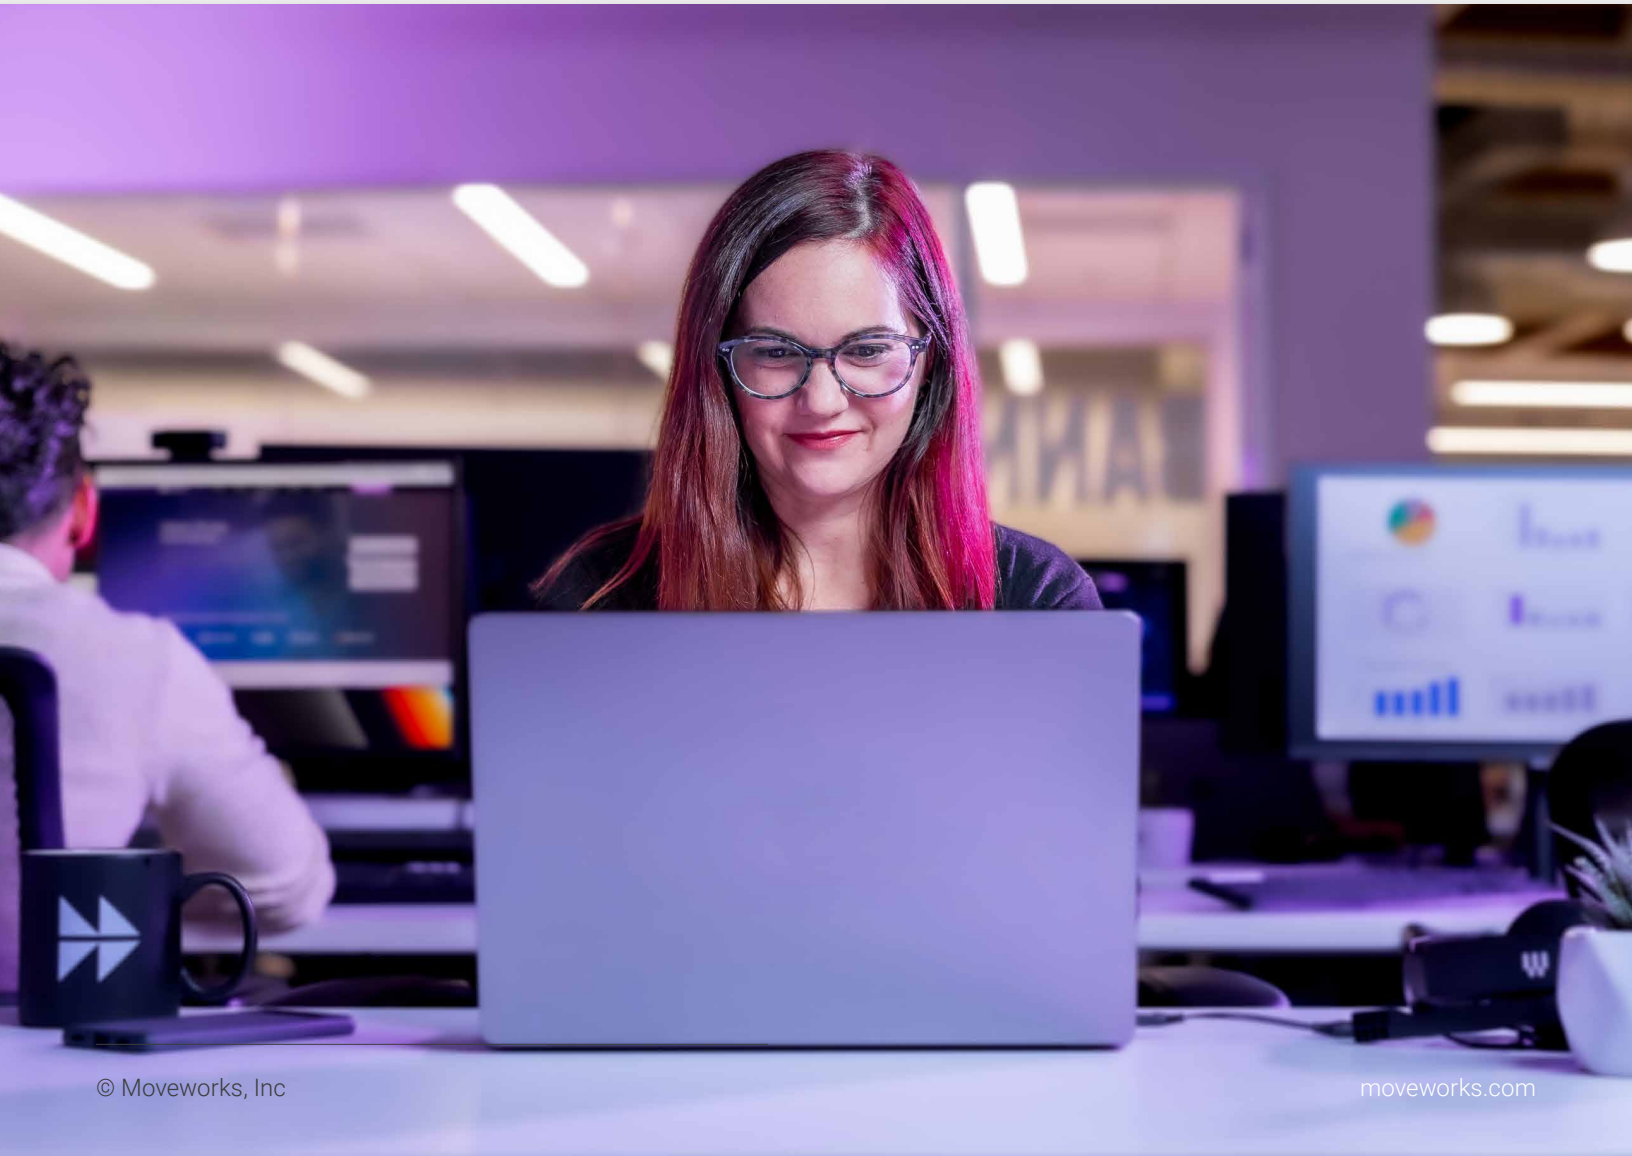

Pre-training and training LLMs: LLMs are pre-trained on vast amounts of annotated information to develop a generic knowledge base. This pre-training phase equips the model with a broad understanding of language patterns and concepts. During training, the LLM learns to predict the likelihood of a word or phrase appearing in a given context, further refining its language comprehension abilities.

Fine-tuning for specific applications: The LLM undergoes a fine-tuning process to tailor its capabilities for specific tasks or purposes. For instance, if the goal is to create a translation capability for a particular language, the LLM will be fine-tuned to understand idioms and

nuances unique to that language. Fine-tuning allows for specialization and optimization, making the LLM more adept at specific applications.

End customer usage: The LLM is deployed for end customer usage, typically through an API or a chat interface. It becomes an integral part of applications that interact with users. Although end customers do not have direct access to the underlying model, they interact with it seamlessly, benefiting from its language understanding capabilities.

Continuous improvement through customer feedback: The journey of an LLM doesn't end with deployment. This stage marks the beginning of an ongoing cycle of improvement. End customers' feedback plays a crucial role in fine-tuning the model further. The LLM is updated and refined to enhance its performance and accuracy as users report issues or provide feedback (including other models' and annotators' input). This iterative process helps the model stay relevant and adaptive to changing requirements. This stage could lead to privacy concerns for end users, which we'll discuss below.

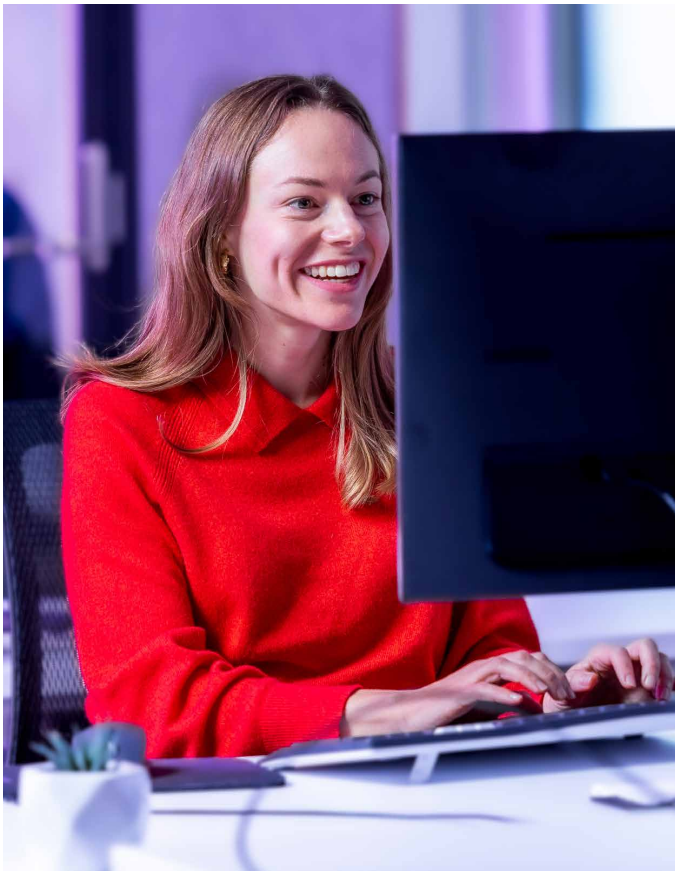

Using threat modeling to identify and address risks in the LLM lifecycle

Threat modeling is a family of techniques that help us answer crucial questions such as “What are we working on?”, “What can go wrong”, and “What are we going to do about those things?” When we’ve done so, we can decide if the dangers of using the LLM are acceptable to us and perhaps apply risk management techniques of acceptance or transfer to what remains.

Adapters help manage privacy considerations

We’ll emphasize a particular approach for handling sensitive data during fine-tuning: adapters. An adapter is a compact plugin for a global model.

It’s possible to use company-specific data (say, “we use the term “partners” rather than “staff”) to create an adapter that helps the LLM answer questions using that terminology. That adapter can be specific to one company, and so even if it contains trade secrets, those secrets are not exposed to the global model but are kept in an adapter, used for only that company.

For example, suppose the model needs to forget specific information or relearn a particular task, such as when using customer data. In that case, the adapter can be retrained with new data that omits the information that needs to be forgotten. This method proves to be more cost-effective and straightforward than retraining the global model.

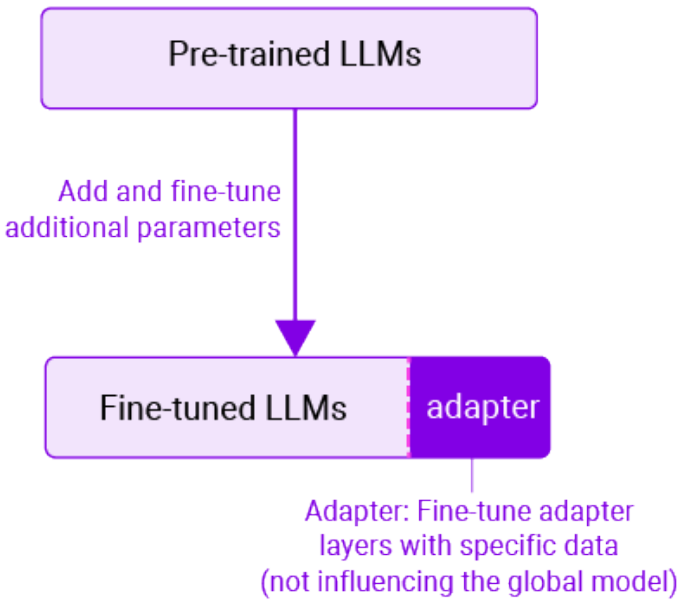

03

Tapping generative AI for knowledge base management

Knowledge is power in the modern enterprise, but effectively harnessing and disseminating knowledge remains a challenge. Traditional knowledge management approaches rely on static documents and manuals that quickly become outdated, leading to fragmented information trapped in silos and requiring expertise to find answers. Meanwhile, employees struggle to get the information they need to excel in their roles.

Generative AI promises a new paradigm for enterprise knowledge management. By leveraging large language models trained on company data, generative knowledge solutions can democratize access to information across the organization. Employees can get answers to natural language questions instantly without combing through manuals. This chapter explores the limitations of traditional knowledge management and how generative AI can transform the creation, curation, and delivery of knowledge.

Limitations of traditional knowledge management

Reliance on static documents and manuals: Traditional knowledge management depends heavily on content repositories like SharePoint, wikis, and manuals. This content is manually created, static in nature, and requires ongoing maintenance.

Difficult to update and maintain over time: With static content, organizations struggle to keep information current. Outdated material proliferates across an enterprise over time.

Information trapped in silos requires expertise to find: Vital information is scattered across documents, systems, and people. Employees waste time trying to track down answers from the right expert.

Not optimized for natural language queries: Static content isn't designed for conversational access. Employees must use rigid search to find answers buried in manuals.

The promise of generative AI in knowledge management

While traditional knowledge management struggles to stay relevant in the digital age, innovations in generative AI present new opportunities to transform the creation, curation, and delivery of enterprise knowledge.

By combining large language models with conversational interfaces, generative knowledge solutions promise to overcome the limitations of static knowledge repositories. Instead of digging through outdated documentation, employees can access the latest information through intuitive conversational experiences. Generative AI has the potential to democratize knowledge across an organization, reduce the burden on experts, and keep information current through continuous learning.

Key strengths of this emerging approach include:

Conversational interfaces enable intuitive access: Unlike static content repositories, generative AI offers conversational access to knowledge. Employees can ask questions in natural language and get accurate answers instantly.

Language models can generate responses to natural language questions: Large language models trained on company data can understand questions and generate knowledgeable responses on the fly. No need for rigid search or digging through manuals.

Knowledge is democratized across the organization: With an enterprise copilot connected to various knowledge bases, any employee can get answers easily without tracking down an expert. Knowledge becomes more accessible across the company.

Dynamic knowledge base powered by continuous learning: Generative models continually ingest new data from tickets, documents, conversations, and more. The knowledge base stays current without manual updating of static content.

Building a generative knowledge management system

Creating an effective generative knowledge management solution requires carefully building and optimizing the AI underpinnings. Key steps include:

Gathering domain documents and data: The knowledge model needs to ingest relevant domain content like support tickets, documentation, conversations, and case data to ground the model in the company's unique data.

Training language models on domain knowledge: Pre-trained language models are fine-tuned on the enterprise domain data to teach the AI about the organization's language patterns, terminology, and content styles.

Optimizing for conversational Q&A: The model is trained and tested on natural language questions and answers to make it conversational with the goal of creating intuitive, accurate responses.

Integrating with business systems and workflows: Tight integration with existing systems like ITSMs and CRMs makes the knowledge AI a seamless part of employee workflows to power widespread adoption.

For example, an IT support knowledge assistant would ingest historical IT tickets, documentation, chat logs, and case data. The model is fine-tuned on this data corpus to understand IT terminology and issues. Conversational training focuses on natural language queries and responses. And integration with the ITSM provides a smooth user experience.

Moveworks' Knowledge Studio

Moveworks Knowledge Studio exemplifies a generative approach to knowledge management. Knowledge Studio ingests historical IT support tickets and documentation from a customer's environment. It fine-tunes a powerful pre-trained language model on this unique enterprise data.

Knowledge Studio provides a conversational interface for users to ask natural language questions. The model generates accurate, up-to-date responses on the fly by grounding in the ingested enterprise content. Moveworks leverages tight integration with ITSM platforms like ServiceNow to embed Knowledge Studio seamlessly into employee workflows.

Key capabilities include:

Ingesting ticket notes, documentation, conversations, and case data

- Ingesting ticket notes, documentation, conversations, and case data
- Fine-tuning language models on the customer's unique corpus
- Conversational Q&A optimized for natural language
- Generating knowledge article drafts in seconds
- Tight integration with ITSM platforms
- Continuous learning as new data is ingested

Knowledge Studio demonstrates how purpose-built generative AI can transform enterprise knowledge management. It overcomes the limitations of traditional static knowledge bases by providing conversational access to up-to-date, grounded content.

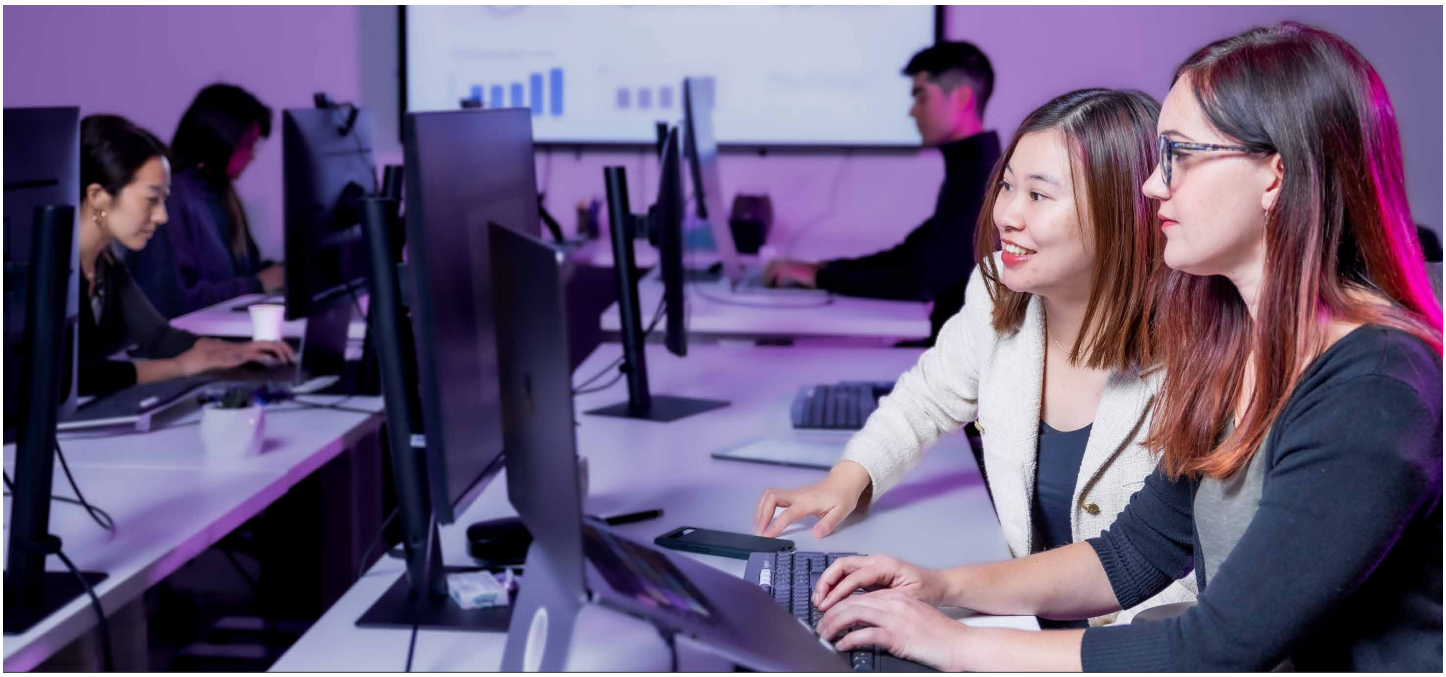

Benefits of a generative knowledge management system

Deploying a generative knowledge management system powered by AI delivers tangible benefits across the enterprise:

Increased productivity and efficiency: Employees get answers instantly without combing through manuals or tracking down experts, saving time and speeding up resolution.

Reduced costs of managing documentation: Static documentation requires ongoing resources for creation and upkeep. Generative AI reduces this burden by dynamically generating content.

Knowledge sharing and discovery: Conversational access and AI-generated insights make knowledge more discoverable across silos so that expertise is democratized.

Agility in updating knowledge bases: The self-learning knowledge model continuously ingests new data to keep content current. No manual document updates required.

For example, generative AI could field a growing percentage of support tickets automatically, freeing agents to focus on complex issues. Turnover and onboarding ramp-up would be reduced since information is more readily accessible. The speed and scale of AI would enable localized knowledge across global markets.

Generative knowledge management unlocks the expertise of an organization. By making institutional knowledge universally accessible, enterprises can increase productivity, reduce costs, share knowledge, and adapt quickly.

Case Study:
Palo Alto Networks powers up IT support with generative AI

Palo Alto Networks faced common knowledge management challenges — outdated documentation, fragmented information, and manual content creation slowed IT support. Employees wasted time searching obsolete knowledge bases.

To reinvent their IT knowledge sharing, Palo Alto Networks implemented Moveworks' Knowledge Studio leveraging generative AI. Knowledge Studio ingests Palo Alto's historical IT tickets, documentation, and conversations to train a customized natural language model.

Now, Palo Alto Networks employees get instant answers to IT support queries without combing through manuals. The AI addresses questions directly by generating responses grounded in Palo Alto's own data.

Within months, Knowledge Studio delivered a significant impact:

- Faster resolution of employee IT issues
- Increased IT agent productivity
- Boost in self-service through discoverable AI content
- Lower costs by automating knowledge creation

For Palo Alto Networks, Moveworks' generative AI transformed its static knowledge base into a dynamic, conversational system adapted to its unique IT environment. Knowledge Studio unlocked improved efficiency, productivity, and employee experience. Palo Alto Networks' success with Knowledge Studio underscores its value in overcoming outdated practices to enable cutting-edge, AI-driven knowledge management.

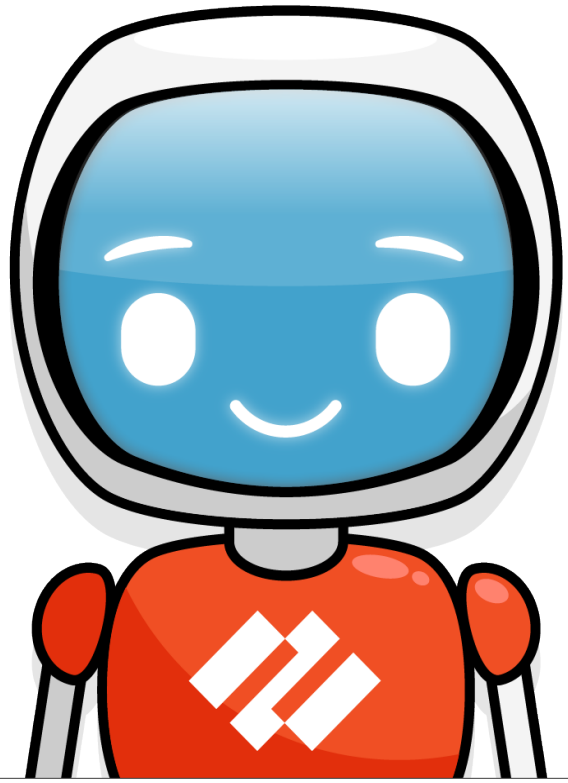

04

Architecting your enterprise with an AI foundation

We're entering an era where conversations with technology feel as natural as chatting with a colleague. Conversational AI is making this possible — and it's transforming how enterprises connect with customers and employees. Copilots, virtual assistants, chatbots — all these tools and more can handle repetitive tasks, enable self-service, and make every interaction more seamless.

Forward-thinking IT leaders recognize conversational AI's immense potential to reshape systems and processes. And they also know that integrating conversational AI requires more than just implementing a chatbot. **To truly transform customer and employee experiences, IT leaders need to architect a robust framework.**

This process starts with identifying high-impact conversational AI use cases that align to business goals. What repetitive tasks can be automated? Where could a copilot boost self-service and efficiency?

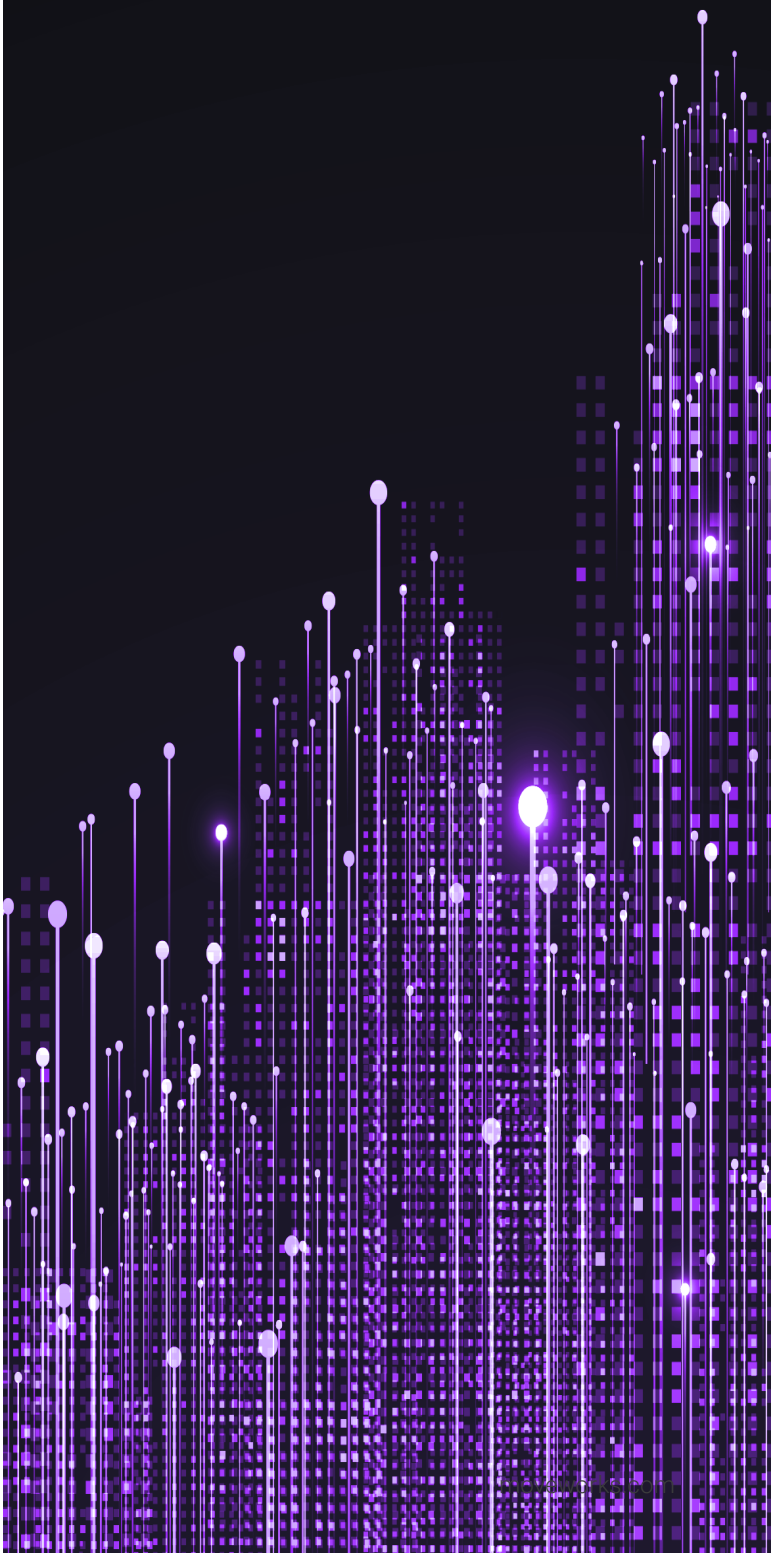

AI is increasingly important to enterprise service desk automation

Intelligent automation is transforming enterprise service desks through the application of AI and machine learning. Natural language processing (NLP) enables service desks to understand requests provided in everyday human language and automate resolutions for common issues like password resets.

For employees, this means 24/7 self-service access for simple requests without waiting for human agents. For service desks, automation significantly improves productivity and efficiency. AI can handle high-volume, repetitive tasks like password resets, freeing up human agents for more complex support issues. This is to say that AI reduces costs by lowering demand for tier-1 support staff.

For enterprise architects, conversational AI represents a powerful way to drive digital transformation of service delivery:

- **Competitive advantage:** Early adopters of conversational AI differentiate through more intuitive, humanized interactions, building loyalty and efficiency.
- **Scalability:** Intelligent assistants cost-effectively handle high volumes of repetitive tasks and queries.
- **Breaking down silos:** Conversational AI integrates with back-end systems, unlocking use cases by connecting previously siloed data.
- **Valuable data insights:** Conversation datasets provide insights to improve products, services, and decisions when analyzed with AI.
- **Futureproofing:** Conversational interfaces are rapidly becoming the norm. Architects must embrace this shift to meet consumer expectations.
- **Driving innovation:** Implementing conversational AI sparks new engagement models, unlocking additional innovations as teams realize their potential.

Conversational AI is a strategic priority for service desks seeking efficient, frictionless support experiences. The time is now for enterprise architects to build a foundation for the conversational future.

Realizing the benefits of AI for enterprise service desks

Implementing AI and conversational interfaces delivers tangible benefits for enterprise service desks:

Streamlined task automation

AI automates high-volume routine IT tasks like password resets, software installs, hardware troubleshooting, and more. In sum: AI eliminates repetitive manual work so human agents can focus on high-value activities.

24/7 self-service access

Conversational interfaces powered by AI offer employees on-demand self-service. Simple requests can be resolved anytime without waiting for IT support availability.

Improved productivity and efficiency

With AI handling common requests, service desk ticket volume is reduced. Agents gain more time to resolve complex issues. Enterprise-wide efficiency improves as employees self-serve for routine needs.

Lower costs

AI automation reduces the need for tier-1 support staff to handle basic requests. Intelligent scaling also optimizes staffing costs to match demand, significantly decreasing service desk overhead.

Enhanced employee experience

Conversational AI delivers consumer-grade experiences employees have come to expect. Satisfaction improves with instant self-service and frictionless interactions.

AI-enabled service desks create value across the enterprise. Automating repetitive tasks while enabling 24/7 self-service improves IT productivity, reduces costs, and enhances employee experiences. **AI is a strategic capability for the modern digital workplace.**

The impact of AI across service desk use cases

AI and conversational interfaces are transforming service delivery across a breadth of enterprise use cases. By automating repetitive tasks and queries, AI drives significant productivity and experience improvements

Common cross-industry automation use cases

AI delivers instant value for service desks by automating high-volume routine IT tasks:

- **Password resets:** AI authenticates users and resets passwords instantly without agent involvement.
- **New employee onboarding:** Intelligent assistants guide new hires through hardware/software setup, access requests, and knowledge absorption.
- **Hardware/software provisioning**: Automated workflows triggered by AI requests deploy equipment and applications to end users.
- **Knowledge management**: Copilots serve up answers to common questions and IT issues by tapping knowledge bases.
- **Access requests**: Requests for network/system access and permissions are handled completely through AI.
- **Incident management**: AI initiates automated diagnosis, triage, and remediation steps based on user-reported incidents.

Function-specific use cases

Conversational AI also streamlines specialized workflows and systems for different business functions:

- **Healthcare:** Clinical system access control and management.
- **Finance:** Automated expense reporting, budget lookups, approvals.
- **Marketing:** Self-service campaign analytics and insights.
- **Sales:** AI-powered lead qualification, data lookups, and CRM system access.
- **HR:** Benefits questions answered, time-off requests processed, and payroll changes triggered.

The list above barely scratches the surface of what teams, organizations, and entire industries can do. The use cases are nearly endless. Enterprise architects should thoroughly evaluate workflows across the business to identify the best opportunities to transform service delivery with AI.

Assessing processes for automation readiness

When identifying which processes are ripe for automation, there are a few key characteristics to look for:

- **High volume:** Processes that handle a high volume of transactions or requests are prime candidates for automation, as bots can efficiently handle large quantities.
- **Repetitive:** Processes that involve repetitive tasks like data entry or answering frequent questions are well-suited for automation to eliminate redundancy.
- **Prone to human error:** Processes with manual data entry or calculations may lead to human error that bots can reduce.

There are also useful techniques for assessing automation opportunities:

- **User surveys and focus groups:** Gather firsthand feedback from employees on which processes are inefficient, frustrating, or error-prone.
- **Journey mapping:** Analyze usage metrics like help desk tickets, sales inquiries, or operational data to spot high-volume automation candidates.
- **Shadowing:** Observe employees directly as they perform processes to experience pain points and inefficiencies ripe for automation.

Leveraging these criteria and techniques, you can systematically evaluate processes across the enterprise to build a prioritized automation roadmap focused on the highest ROI. This assessment lays the groundwork for automating legacy processes with new conversational AI capabilities.

Developing an enterprise AI strategy and roadmap

Once high-value automation opportunities have been identified, the next step is developing an enterprise AI strategy and roadmap. Key elements include:

Aligning to business goals: AI initiatives should tie directly to overarching enterprise objectives around customer experience, employee productivity, cost reduction, and more.

Prioritizing opportunities: Rank automation opportunities based on expected ROI to focus first on the use cases that will demonstrate value quickly.

Phased rollout plan: Given the scale of enterprise-wide automation, a phased rollout allows for piloting, testing, and iteration before full production deployment.

Success metrics: Define quantitative metrics and KPIs for each AI application to continually measure impact toward business goals.

Scalability planning: Architect the AI platform and use cases to efficiently scale up transaction volume, end users, languages, and global reach as adoption increases.

Maintainability: Establish processes for ongoing maintenance, monitoring, and model refinement to ensure applications continue meeting business needs.

Continuous improvement: Build in feedback loops to identify areas for additional automation and incremental enhancements over time.

With a strategic roadmap anchored to business objectives and value, enterprise leaders can steer an AI transformation that improves experiences, efficiency, and competitiveness over the long term.

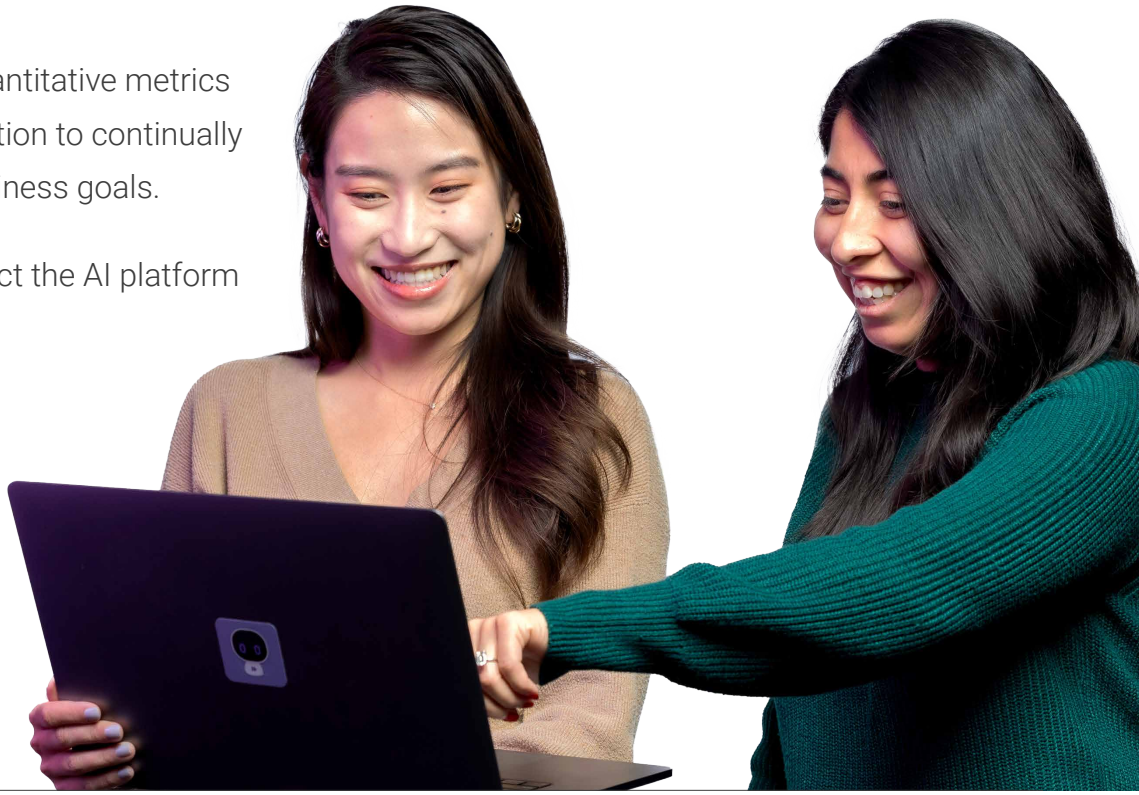

Driving user adoption through change management

When introducing new AI capabilities, thoughtful change management is key to driving user adoption. Best practices include:

Communicate the why: Explain the specific benefits AI will bring to employees' day-to-day work so they understand the value.

Training: Provide hands-on training for employees on new AI-enabled processes, tools, and user interfaces and offer ongoing support materials.

Ambassador programs: Identify power users who are excited about AI to be ambassadors who provide peer coaching and support during the transition.

Pilot groups: Test changes with pilot groups first to work out issues before full rollout. Incorporate pilot user feedback into the implementation plan.

User feedback channels: Actively collect user feedback through surveys, focus groups, interviews, and other channels before, during, and after rollout. Use these insights to refine your approach.

Celebrate wins: Highlight key employee stories and successes enabled by AI. Quantify benefits through metrics to showcase value delivered.

With proactive communication, training, and user experience focus, organizations can smooth the transition to new AI capabilities, rapidly achieve proficiency, and realize the full benefits.

Creator Studio makes building conversational AI use cases easy

As conversational AI transforms how enterprises interact, a key challenge has been the difficulty of building customized solutions tailored to each organization's unique needs. Moveworks' Creator Studio aims to change that.

Creator Studio is a no-code, generative AI platform that allows anyone to rapidly create conversational AI use cases for their enterprise systems and workflows. With an intuitive interface and advanced natural language capabilities, Creator Studio democratizes conversational AI development.

By leveraging Creator Studio's enterprise knowledge graph and integration capabilities, these AI-powered conversations can connect data across fragmented systems. Employees get a unified interface to systems through natural language.

Creator Studio builds on Moveworks' foundation of large language models. With no need for coding, the platform opens up development to citizen developers across the business. This democratization of conversational AI promises to accelerate digital transformation.

There are three main workspaces within Creator Studio:

- **Paths:** Generate custom conversations and workflows for activities like transferring sales accounts or provisioning servers.
- **Queries:** Securely access data from any app or system, like surfacing purchase orders or looking up HR information
- **Events:** Send proactive, actionable notifications, such as contract signing reminders or follow-up support surveys.

We'll leave you with this…

Architecting for an AI-enabled enterprise is an iterative process requiring ongoing learning and refinement. While this chapter provided a high-level overview, each organization's AI journey will be unique.

By starting with targeted, high-value use cases, architecting with scalability in mind, and staying nimble to accommodate new innovations, IT leaders can drive an AI transformation.

Throughout this process, stakeholder buy-in, user-centric design, and change management will be instrumental to success. Keeping the end-user experience at the forefront will ensure adoption and business value.

While there are challenges inherent to any new technology, the potential benefits make overcoming these hurdles worthwhile. With the right strategy and execution, conversational AI can redefine your employee and customer experiences, unlock new sources of value, and establish a foundation for continued innovation.

Conversational AI can truly change how businesses operate — but only with enterprise architects paving the way.

Moveworks