FORUM3

What Our Clients Ask Us Most

What leaders ask us about AI, and what we tell them

GUIDES & RESOURCES

GUIDE · 10 PAGES

These are the questions leaders ask us most, with straight answers and the data behind each one. Every number links to its source. Tool details change quickly, so treat product-specific answers as accurate as of today and confirm them in your own plan.

Have a question that is not here? Send it to learning@forum3.com. We love a good one.

What's inside

Getting started
4 questions · 01–04
PAGE 2

Tools
10 questions · 05–14
PAGE 3

Skills
6 questions · 15–20
PAGE 6

Connectors and data
7 questions · 21–27
PAGE 7

Safety and admin
4 questions · 28–31
PAGE 9

Getting started

01 How do we get started with AI?

Start with people, not the tool. The first steps are cultural: leaders use AI out loud, someone owns it, and the rules are written down before access widens. Then choose your tools and get your data in order, train a few champions, point AI at real workflows, and check in regularly to adapt. It runs as a loop, not a straight line.

THE DATA

McKinsey (1,719 respondents in 97 countries) found that 80% say AI improved their own productivity, but only 37% say it contributed to company profit. Individual wins do not add up on their own, which is why starting with people and workflows matters more than picking a tool.

02 Adoption is bumpy. How do we get the rest of the team to actually use AI?

Leaders go first, and say so out loud. People copy what they see their leaders do, so make AI use visible in meetings and decisions. Then bring in champions and eager early adopters to show others how, and show everyone where AI saves time in their own job. Name an owner for AI so someone is accountable for momentum.

THE DATA

BCG found that employees who feel positive about generative AI rise from 15% to 55% with strong leadership support, yet only about one in four frontline employees say they get that support. In the Wharton and GBK Collective survey (self-reported), 58% of leaders report strongly positive ROI where a chief AI officer was newly named, versus 24% where AI was added to an existing executive's job and 8% where there is no such role and none planned. The smaller groups make this directional.

03 How should we roll AI out, and who goes first?

Leaders and your AI task force go first, then your champions and the people who are eager, then everyone, over roughly 30 to 90 days. Widen access only after your AI use policy is in place, and give everyone enterprise-grade access and basic training so nobody is tempted to use a personal account.

THE DATA

In a BCG survey of more than 10,600 employees, more than half said that if their workplace did not give them the AI tools they need, they would find alternatives and use them anyway.

04 How do we measure whether AI is working?

Using AI is not the same as getting value from it. Pick the number before you pick the tool: let each team choose one metric, agree on a concrete target at the start, and track at least one growth measure alongside efficiency. Start with one growth measure for one team rather than a company-wide scorecard.

THE DATA

McKinsey found that 80% say AI improved their own productivity, 37% say it contributed to company profit, and about 6% are high performers (at least 5% of profit from AI and a significant impact). In the Wharton and GBK Collective survey (self-reported, and a link rather than a proven cause), 35% of departments that track a growth measure report strongly positive returns, versus 13% of those that track efficiency only. BCG found that companies tracking AI value directly in the P&L realize about three times the value of those that do not formally measure it.

Tools

05 Which AI tool should our team use, and can we use more than one?

Pick one general-purpose, enterprise-grade tool as your team's main tool. It should cover roughly 75 to 80% of your work, with specialized tools as add-ons only when needed. Claude and ChatGPT both work well as a main tool. If you are a Microsoft shop, Copilot is often the lowest-friction secure option, so test it against your top pick on real tasks first. Put every tool, main or extra, under your AI use policy.

06 What uses up credits fastest, and what if we run out?

Bigger models, higher effort settings and more context use the most, meaning many files, whole folders or very long chats. Agentic work burns through limits faster than chat, and scheduled tasks add up, especially when left on a big model. Give your AI the one file it needs rather than a folder. If you run out, admins can usually raise limits or add credits for one person or everyone, and usage alerts (we suggest 50% and 75%) keep you from being surprised. Models keep getting more efficient, so older comparisons can overstate today's cost.

07 How does memory work?

Memory lets your AI remember details about you and your work across chats. It is often off by default for privacy, so check your tool's settings to turn it on. Most tools have two switches: reference past chats, and generate memory from chat history. Deletion works differently by tool, and in some, deleting a chat does not delete what was saved to memory, so review saved memories and delete them separately. On a business account, ask your admin what is available.

08 Can my AI triage my inbox and draft replies?

Yes, with Claude, ChatGPT or Copilot. Connect your email, let your AI sort the inbox and write draft replies, then review and send. Drafts wait in your drafts folder, so nothing goes out without you. Treat every draft as a first pass and read before you send. Check the triage against your real inbox for the first week or two, and save each correction to your skill so it improves.

09 Which model should I use, and why not always the biggest?

Use the everyday model for most of your work and step up only for harder jobs. In Claude that means Sonnet first, then Opus for deeper analysis and research. In ChatGPT, start with the default and raise the thinking level, or use Pro, for harder jobs. In Copilot, use Auto or pick a Claude or GPT model. Bigger models use up your usage faster and take longer to answer, so they are not automatically better. Not sure? Try the same task in both, and if the smaller one works, stay with it. Model names change quickly, so treat these as a snapshot.

10 How much usage do I have left, and is there a per-person cap?

Look in your tool's settings. Claude shows your usage under Settings, and in Copilot's Cowork you can type /cost to see the credits used on the open task. On a business account, your admin decides what you can see and can set limits, so if you cannot find it, ask them. Check before a big task so you do not run out halfway, and note that reset windows differ by plan, so confirm yours.

11 When should I use chat, an agentic workspace or projects?

Chat is for conversation and quick questions. An agentic workspace (Cowork, for example) takes actions with your files and tools. Developer tools such as Claude Code are built for people who write software. Projects suit team collaboration. The lines overlap, the right choice is partly personal preference, and it has changed over time as the products blend together.

12 What is an AI agent, and how is it different from a skill or a scheduled task?

“Agent” is a loose term, so start with the job instead of the label. Describe the problem and the workflow, and let your AI recommend what fits. As a rough ladder: a skill is saved instructions for one task, a scheduled task is a skill on a timer, an agentic workflow chains several steps from trigger to result, and a shared agent has its own email or chat identity and remembers across users. Our line is to reserve the word “agent” for that last one. Each step up adds autonomy, and you often do not need to go all the way up.

THE DATA

McKinsey found that agent adoption is concentrated in large organizations: 40% of those with $1 billion or more in revenue are scaling AI agents, up from 27% a year earlier, compared with 22% of smaller organizations, flat year over year.

13 Why does our AI output look AI-generated, and how do we fix it?

Do not accept the first draft. Push back, ask for changes and treat the first pass as raw material. Give your AI examples of your real writing and the standards you hold, once, in a skill, including the phrases and habits you do not want. Then edit the last stretch yourself. The more of your own voice and judgment you put in, the less the result reads as AI-generated.

14 Can I move my ChatGPT history and memory into another tool?

Generally yes. Ask the old tool to export what it knows about you, review the text, and import it into the new one. If your organization blocks the data export, ask the old tool to write out its saved memory as plain text instead. Import options change often and may be controlled by your admin, so check what your new tool offers today.

Skills

15 How do I create a skill, and when should I?

A skill is a saved set of instructions your AI reuses, so you stop re-explaining the task. Do the task in a chat until the result is what you want, then tell your AI to create a skill for this task and save it (Save skill in Claude, Install in ChatGPT, and in Copilot, ask Cowork). If you have done it more than twice, make it a skill. It does not have to be perfect first: a skill covers most of the repeatable work, you check the rest, and you can update it any time. Where you can create skills depends on how your admin set things up.

THE DATA

In BCG's 2026 survey of close to 12,000 workers, 42% of frontline employees who use AI regularly report saving about a workday a week, yet 66% say they get limited or no guidance on what to do with the time. A saved skill is how one person's time savings becomes something the whole team repeats.

16 Why don't my scheduled tasks run when my computer is off?

Local scheduled tasks only run while your computer is on and the app is open. Tasks missed while it was asleep tend to run all at once when you reopen it, which uses up your usage, so many tools have a keep-awake setting. Some tools now also run tasks in the cloud, but cloud tasks cannot read local folders, so move the source files to SharePoint, OneDrive or Drive first. Where these settings live has moved before.

17 How do I share a skill with my team?

Most tools offer three levels: personal skills, a shared library anyone on the team can add to, and admin-installed skills for everyone. Sharing a skill shares the instructions, not your chats or data, so a colleague still needs their own access to the same files and tools. A good pattern is to build skills centrally and have an admin install them, such as a brand skill for marketing and job description templates for HR.

18 Can I update or retrain a skill instead of starting a new one?

Yes. Give your AI the final output and tell it to update the skill, and avoid creating a duplicate. Say it explicitly: do not count on corrections you make in conversation being saved automatically. In some tools you save the updated skill and choose to replace the old one.

19 How do I get my AI to follow our brand?

Give your AI two things: instructions and files. The instructions cover how your brand sounds and looks, with exact hex codes and the colors to avoid, saved as a skill so you reuse them every time. The files are your logos and templates, kept alongside the skill (in some tools, a plugin or shared folder is the way to package them), because AI will invent a logo if you do not hand it the real one. Expect close, not exact: logos and templates are the hardest parts for AI to match. If the output starts to drift, rebuild the skill from a clean start.

20 What is the difference between a skill, a plugin and a project?

A skill is saved instructions for one way of working, such as your brand guidelines. A connector links your AI to another app. A plugin is a bundle: skills, connectors and files such as logos, which an admin can install for everyone. A project is like a file cabinet for a client or team, holding its own files, skills and connectors. They overlap, so use this test: build a skill when you want something done a specific way, and use a project when you want a shared workspace. A skill holds what your AI cannot find on Google, like your hex codes and your own process.

Connectors and data

21 How do I give my AI access to my files?

SharePoint is the one place that works with all three of the big tools, so store shared files there. Claude can use local folders, or Google Drive and SharePoint through a connector. ChatGPT takes uploads or connects to Google Drive, SharePoint and OneDrive. Copilot's Cowork sees only OneDrive and SharePoint. Files on a network drive should move to SharePoint or OneDrive first. Many teams keep working files local and shareable outputs in the cloud. Just know that scheduled tasks running in the cloud cannot read local folders.

22 Why won't the Microsoft 365 or Outlook connector work, and why can't my AI send email?

Check the per-connector permissions first, because sending may be turned off or may need your approval each time. A good workaround is to have your AI save a draft that you review and send, or to send a Teams or Slack message instead. Admin security settings such as multi-factor authentication can also get in the way, so a vendor support ticket is the next step if you are stuck.

23 How do I share an artifact or dashboard?

Sharing options are still limited and change often: view-only links, sharing inside your organization and PDF export. People you share with may need their own connections to see live data. A downloaded HTML file is a snapshot, while a live artifact pulls fresh data every time it loads, so spot-check the numbers before you share. We suggest leaving external sharing off unless you have a clear use for it.

24 Is there a connector for our tool, and what if there is not one?

Our order of preference is a vendor connector first, then a custom connector, then scheduled exports into shared storage such as SharePoint or Drive, with a browser extension as a last resort. Connectors are usually vendor-led and slow to arrive, so ask your vendor first and do not plan around one in the short term. Browser-based workarounds can break when a session times out.

25 Why connect a database instead of just uploading our data?

Uploading a file, or asking your AI directly, usually gives it a sample, like asking one person for their best guess. A connected database holds the full history and refreshes on a schedule, so every answer reflects everything. A shared definitions layer makes sure terms like same-store sales mean the same thing for everyone, so two people asking the same question get the same answer. The data owner keeps control of how your AI reads the data.

26 Our data is messy and spread across systems. Is that holding us back?

Do not wait for perfect data to start. Writing, research and summaries need no clean data at all. When answers depend on your data, make it accurate, complete and connected, and fix the data behind your top use case first. Work from your latest data, and ask your AI to spot the gaps. Remember that your AI can open anything you can, so check file permissions before you point it at a shared drive.

THE DATA

Gartner predicts that through 2026, organizations will abandon 60% of AI projects that lack AI-ready data. In a July 2024 Gartner survey of 1,203 data management leaders, 63% said their organization either lacks, or is unsure it has, the right data management practices for AI.

27 Can I make connectors read-only or require approval?

Yes. A connector lets your AI reach into another app, like email or SharePoint, and most tools let you set what it can do. Start read-only: reading is always allowed, writing needs approval and deleting is blocked. Re-check permissions every quarter, because new tools get added over time. Your AI can open anything you can, so check your own SharePoint and file permissions first.

THE DATA

Keeping a person in the loop is still the norm. In the Wharton and GBK Collective survey (self-reported), 17% of departments let an agent finish a task without a person approving each step.

Safety and admin

28 Is it safe to put confidential company data into AI?

Yes, with rules. As we understand it, business and enterprise plans from the major vendors do not train on your data by default, while free plans can differ, so confirm your own plan's terms. Write a never-enter list before anyone asks: payment card data, customer personal information, passwords and API keys. Chat is the lowest-risk way to use AI. Connectors and the browser can act as you, so more access means more reach. Before you turn on a new connector, run a worst-case tabletop and ask what the worst thing is that it could do.

THE DATA

In IBM's 2025 Cost of a Data Breach research, 13% of organizations reported a breach of their AI models or apps, and 97% of those lacked proper AI access controls. High levels of shadow AI added about $670,000 to the average breach cost. BCG also found that more than half of employees would find and use AI tools anyway if their workplace does not provide them.

29 What can admins restrict?

Admins typically control access through domain locks, role-based permissions and single sign-on, and can limit features, connectors and who gets agentic tools. Controls differ by vendor and change often, and some are more built out in one tool than another. We suggest naming a member of your AI task force as an admin, so someone owns the settings and keeps up with the changes.

30 What can agents handle today, and what still needs a person?

Let agents read freely, ask for approval before they write, and keep deletes blocked. Add autonomy only after the controls are in place. Agents are good at drafting, sorting, summarizing and gathering information. A person should still approve anything that sends, posts, spends or changes records.

THE DATA

BCG found that companies with all six critical agent controls in place create about three times the agentic AI value of companies with one, yet only 5% have the full set. Gartner forecasts that more than 40% of agentic AI projects will be canceled by the end of 2027, citing rising costs, unclear business value and weak risk controls. And in the Wharton and GBK Collective survey (self-reported), 17% of departments let an agent finish a task without a person approving each step, while 67% use agents in some form.

31 Can my AI use my browser, and is it safe?

Yes, but a browser extension acts as you, with your logins, so treat it as the riskiest option. A connector may only be able to save a draft, while a browser can press send. Keep a person approving actions, start with limited, structured tasks and limit it to certain tabs or windows. Some teams use a read-only audit account and sign out of everything else first. Admins can turn browser control off entirely.

KEEP LEARNING

Questions? Email learning@forum3.com. More guides live in the Forum3 Learning Hub.