[Company Name] Employee AI Training in 30 Minutes

Facilitator agenda, talking points and a 10-question quiz with answer key, built on the [Company Name] AI Acceptable Use Policy. For a live or video session with employees, contractors, temporary staff and interns. If anything here and the policy ever differ, the policy applies.

Free template from CIO Support, published by LAN Service Group, Inc. Version 1.0, October 2026. Built on the AI Acceptable Use Policy template version 1.0.

How to use (delete this box before the session): fill in every [bracketed field] with the same details as your AI Acceptable Use Policy, and delete this box, the template line under the title and the footer. Swap the practice prompts for examples from your own work, keeping them free of real customer or employee data. If your policy uses different classification levels, change Confidential, Restricted and Internal here and in quiz question 3 to match. Print the quiz page on its own and keep the answer key with the trainer. Change this kit whenever you change the policy. It explains the policy and does not replace it. You may copy, edit and use this template inside your organization. General guidance, not legal advice.

Trainer: [Trainer] Session date: [Date]

Agenda (30 minutes)

Time	Segment	People leave knowing	Policy section
0:00 to 0:03	Why we have AI rules	AI is welcome at work; the rules protect our data and our customers, and you own what AI helps you produce.	1, 2, 5, 9
0:03 to 0:08	What data never goes in	The no-exceptions list, the written-approval list, and what to do when unsure.	4
0:08 to 0:12	Approved tools and accounts	Which tools, which account, how to ask for a new one, and when to disclose AI use.	3, 6, 7
0:12 to 0:17	How to spot a risky prompt	Seven warning signs and five practice prompts.	3, 4, 5, 8
0:17 to 0:21	What to do after a mistake	Stop, get an exposed credential changed or revoked at once, report within [number] hours, and say what happened.	8
0:21 to 0:30	Quiz, answers and questions	Ten questions, then go through the answers together.	3 to 8

Allow extra time if you take many questions or collect signatures in the session.

Why train at all: NIST's AI Risk Management Framework (AI RMF 1.0), which is voluntary, includes under its Govern function that personnel and partners receive AI risk management training to enable them to perform their duties and responsibilities consistent with related policies, procedures, and agreements (GOVERN 2.2).

1. Why we have AI rules (3 minutes)

- [Company Name] encourages the productive use of AI tools. The rules exist so you can use AI at work without exposing company, customer or employee information, and without passing on AI output that has not been checked.
- The rules apply to everyone who uses AI tools for [Company Name] work, on company devices or on personal devices. "AI tool" means more than chat assistants and AI search tools. It includes AI features built into software we already license, such as email, documents, CRM or meeting tools, plus coding assistants, browser extensions and plug-ins that read, summarize or rewrite content, meeting note-takers, transcription and recording bots, and image, audio and video generators.
- You are responsible for anything you produce with AI help, exactly as if you had written it yourself. AI tools can be confidently wrong, invent sources and produce insecure code. Check facts, figures, names, citations and links against a reliable source before you rely on them. A person must review AI output before it goes to a customer, is used in a contract, is sent to a regulator, is deployed as production code or becomes part of a regulated record.
- AI output must never be the sole basis for decisions about a person's hiring, pay, promotion, discipline or termination, or a customer's eligibility for [products or services]. A person must make, and be accountable for, the final decision. Any AI tool used to support these decisions needs prior [HR / Legal] approval.

- [Company Name] may monitor the use of company accounts, devices and networks, including AI tools, as permitted by law and [Company Name]'s [monitoring or electronic communications policy]. Breaking the rules is handled under [disciplinary policy] and may lead to loss of access to AI tools or other disciplinary action.

2. What data never goes in (5 minutes)

- No exceptions, in any AI tool, approved or not: passwords, API keys, access tokens, private keys or other credentials; payment card numbers, bank account numbers, Social Security numbers or other government ID numbers.
- Only with written approval: health information about identifiable people (including PHI); CUI and export-controlled technical data (ITAR or EAR); personal information about customers, employees or job applicants beyond what the approved use requires; unpublished inventions, trade secrets and [other sensitive intellectual property]; and information a customer, partner or other third party gave us under a confidentiality obligation, or that a contract says may not be shared with outside service providers. Approval means [owner] has approved that specific tool and environment for that data in writing, after [Legal / Compliance] confirms the required contracts and safeguards are in place. Without that written approval, the answer is no.
- Everything else: before you type, paste or upload anything, check its classification under [Company Name]'s [data classification policy] and the limit set for that tool on the approved list. Confidential material, such as financials, contracts, customer lists and source code, goes only into tools approved for Confidential data.
- If you are not sure how information is classified, treat it as Restricted and ask [contact].

3. Approved tools and accounts (4 minutes)

- Use only the approved AI tools, and only through company-managed accounts that you sign in to with your company credentials [through single sign-on]. Do not use personal AI accounts, free or paid, for company work. The current approved list is kept at [location or link] and is maintained by [owner].
- Do not connect an AI app, plug-in or extension to company email, files, calendars or other systems, and do not grant it access when asked to sign in or approve permissions, unless it is on the approved list for that use.
- Want a new tool, or a new use for an approved tool? Contact [owner or request link], and do not use the tool for company work until it is approved. Requests are normally answered within [number] business days.
- Be open about AI: disclose AI assistance when [e.g. a customer, regulator or contract requires it, or when content is published under a person's name]. Do not use AI to impersonate a real person, or to create content intended to mislead.
- Before using a meeting note-taker, transcription or recording tool, follow [Company Name]'s rules on notifying participants and obtaining consent, and applicable law.

4. How to spot a risky prompt (5 minutes)

Before you press Enter, read your prompt and anything you attached. Stop and rethink if you see any of these:

1. It contains a password, API key, access token, private key or other credential, or a payment card, bank account, Social Security or other government ID number. Remove it: there are no exceptions.
2. It names or identifies a real person together with their health, pay, performance or other personal details.
3. You are pasting or uploading a whole file, inbox, contract, customer list or code base when you only need part of it, or into a tool that is not approved for Confidential data.
4. It asks the AI to decide something about a person, such as who to hire, reject, discipline or pay.
5. You are in a tool that is not on the approved list, or signed in with a personal email address.
6. The words did not come from you: someone sent you text to "paste into your AI", or a document, email or web page you asked it to read contains instructions, such as to ignore its rules or to send information somewhere. Do not follow them. If the tool then acts oddly or asks for a password or other credentials, stop and report it.
7. The output will go straight to a customer, into a contract, to a regulator, into production code or into a regulated record without a person checking it first.

Why the sixth sign matters: OWASP's Top 10 for LLM Applications lists prompt injection as LLM01:2025 and says indirect prompt injections occur when an LLM accepts input from external sources, such as websites or files. Treat instructions hidden in content as untrusted.

Practice: read each prompt aloud and ask which warning sign it hits. These are invented examples; replace them with situations from your own work, without real names or data.

Practice prompt	Sign	Safer approach
"Here's our overdue-accounts spreadsheet with card numbers. Write a reminder email for each customer."	1, 3	Remove the card numbers, which never go into AI. Draft one reminder template in an approved tool and merge the customer details outside the AI tool.
"Here's an employee's doctor's note. Draft a reply to their manager about their leave."	2	Health information about an identifiable person stays out unless [owner] has approved that specific tool and environment for it in writing, after [Legal / Compliance] confirms the required contracts and safeguards. Draft a general reply with no name or health details, or ask [contact].
"Rank these applicants' résumés and tell me which ones to reject."	4	A person makes the final decision on hiring. Any AI tool used to support it needs prior [HR / Legal] approval.
"I'll polish this contract clause in my own AI app on my phone. It's quicker."	5	Personal AI accounts are not for company work, free or paid. Use the approved tool with your company account.
A supplier's PDF you asked the AI to summarize contains a line telling the assistant to send the conversation to an outside address.	6	Do not act on it. Stop and report it to [contact, email or phone].

5. What to do after a mistake (4 minutes)

If you put something into an AI tool that should not be there, or AI output you used turns out to be wrong:

1. Stop. Do not add more information or carry on with the conversation.
2. If a password, API key, access token or other credential went in, tell [contact, email or phone] straight away so it can be changed or revoked. Do not wait for the reporting deadline.
3. Report it to [contact, email or phone] within [number] hours. Report if you entered information into an AI tool that the policy does not allow, including by accident; if AI output that was used or sent turns out to be wrong in a way that could cause harm; or if an AI tool behaves unexpectedly, such as showing someone else's data or asking for credentials.
4. Say what happened: which tool and which account, what information went in, when, and whether any output was sent to anyone.
5. If others can see the content, for example through a shared link or shared workspace, stop sharing it now. Otherwise, ask [contact, email or phone] before you delete the conversation, file or history, so they can record what was shared first.
6. If wrong output has already gone to a customer or anyone else, do not send more until a person has checked it.
7. Prompt self-reporting helps us limit the damage. [Optional: An employee who promptly reports their own honest mistake will not be disciplined for that mistake alone.]

Keep the optional sentence only if your policy includes it, because the template says to include it only if HR and Legal approve it.

6. Quiz, answers and questions (9 minutes)

Hand out the quiz page. Allow about five minutes, then go through the answer key together. Close the session by pointing people to the policy, the approved list at [location or link] and the one-page Employee AI Rules, and by collecting signed acknowledgments if the policy is new or has changed materially.

Quiz: using AI at [Company Name]

Circle one answer for each question. Ten questions, about five minutes.

Name	[Date]	Score (out of 10)

1. A colleague uses a paid personal AI account with model training turned off. Can they use it for company work?

A Yes, if they also turn off chat history and use a strong password
B Yes, because training is off
C No, personal AI accounts are not for company work

2. Which of these may never go into any AI tool, even an approved one?

A Text from our public website
B An API key
C Non-sensitive internal meeting notes

3. You are not sure whether a document is Confidential or Restricted. What do you do?

A Treat it as Restricted and ask [contact]
B Treat it as Internal
C Treat it as Confidential, then use any tool approved for Confidential data

4. What is needed before health information about an identifiable person goes into an AI tool?

A Nothing, if the tool is on the approved list for Internal data
B Your manager's approval by email
C Written approval from [owner] for that tool, after [Legal / Compliance] checks

5. A browser extension that is not on the approved list asks for permission to read your company email. What do you do?

A Decline, and request it through [owner or request link] if you need it
B Allow it if it looks professional
C Allow it for one day only

6. An AI tool drafted a customer email that includes a price and a deadline. Before you send it, you:

A Send it, because the AI wrote it
B Check the facts and figures against a reliable source, and review it
C Ask the AI to double-check its own figures and send it if it says they are right

7. A manager asks you to have AI reject the weakest job applicants automatically. Is that allowed?

A Yes, if the tool is on the approved list
B Yes, if HR is copied
C No. A person decides, and the tool needs [HR / Legal] approval

8. You realize you pasted part of a customer contract into an unapproved tool. What first?

A Stop, and report it to [contact, email or phone]
B Delete the chat right away and tell no one
C Wait to see if anything goes wrong

9. You ask an AI tool to summarize a web page, and it starts asking for your company password. What do you do?

A Enter it, since the tool is approved and you signed in through company single sign-on
B Stop and report it: a tool asking for credentials is behaving unexpectedly
C Enter it, because the page may need a login

10. You want an AI note-taker to join a call with a customer. What comes first?

A Nothing, if the note-taker is on the approved list
B Tell the customer at the end of the call that the meeting was recorded
C Follow the rules on notice and consent, and applicable law

Answer key (for the trainer)

No.	Answer	Why	Policy section
1	C	Personal AI accounts, free or paid, are not for company work. Use a company-managed account.	3
2	B	Credentials such as API keys are on the no-exceptions list.	4
3	A	If you are not sure how information is classified, treat it as Restricted and ask [contact].	4
4	C	Health information about identifiable people is on the written-approval list. Without that written approval, the answer is no.	4
5	A	Do not grant access unless it is on the approved list for that use. Do not use a tool for company work until it is approved.	3
6	B	A person must review AI output before it goes to a customer. Check facts, figures, names, citations and links.	5
7	C	AI output must never be the sole basis for hiring decisions. Any AI tool used to support them needs prior [HR / Legal] approval.	5
8	A	Report within [number] hours, including accidents. Prompt self-reporting helps us limit the damage.	8
9	B	An AI tool that behaves unexpectedly, such as asking for credentials, must be reported. Passwords never go into AI.	8
10	C	Before using a meeting note-taker, follow [Company Name]'s rules on notifying participants and obtaining consent, and applicable law. The note-taker must also be an approved tool.	6

Sources

NIST AI 100-1: Artificial Intelligence Risk Management Framework (AI RMF 1.0), GOVERN 2.2: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

OWASP Top 10 for LLM Applications: LLM01:2025 Prompt Injection: https://genai.owasp.org/llmrisk/llm01-prompt-injection/