AI and Information Security: 42 Answers to Workplace Questions

Quick Reference Guide · extended-mind.ai · sources current as of October 19, 2026
Guide edition: November 18, 2026

These are concise answers to all 42 questions. The full version on the website contains detailed explanations, examples, and sources. Below each answer is a link to the corresponding question in the full guide.

Core principle: the decision depends on the specific task, the data composition, and how it is processed. Small business size or millions of service users do not protect your data: an API key (secret connection code) with read access to a customer database exposed in a chat allows anyone to export that database, and technical incidents (such as in ChatGPT in March 2023 or the DeepSeek database in January 2025) expose chat histories and keys to outsiders without targeted surveillance.
Baseline rule. Before uploading, ask yourself: "What exactly will the external service receive, and under what terms am I entitled to share it?" Check the rights to the document, the AI provider's processing terms, and your organization's internal policies (if you work independently - the law and your obligations to the client). If the document cannot be shared, but the service itself is permitted, ask your question without uploading the file and fill in the real numbers and names yourself without involving the AI.

Grey Zones: 12 Everyday Actions

Action	What People Overlook	#
Only translated text / checked commas	A cloud service receives the entire submitted material - even for a single comma	2
Enterprise software has its own built-in AI agent	An external third-party cloud provider may sit behind the interface	4

Action	What People Overlook	#
Bought a Pro subscription, working from home	A personal subscription does not replace corporate authorization; PC security and access method matter	6
Deleted chat / turned on Temporary Chat	Does not guarantee deletion of files, technical logs, backups, or search index records	9
Described in words / gave summary numbers	Aggregate figures and prompt wording itself can expose specific individuals and company plans	21
Asked AI to remove last names and prices from a file	To sanitize the file, the service first receives the original; sanitize only where the original is permitted	25
Replaced names with "Employee 1" / screenshot / paragraph by paragraph	Indirect identifiers reveal the person; codes do not guarantee anonymization; format does not hide text	26, 27
Blacked out amounts in PDF / deleted an Excel sheet	Extractable text remains under black boxes; spreadsheets retain hidden data and Pivot Table caches	28
Only opened the AI browser sidebar	The assistant may read the active tab and include fragments in web search queries	31
Asked for a summary of an incoming email	A hidden instruction inside the email can alter the assistant's actions within its granted permissions	32
Pasted an error message with an access key (API)	The key grants system access; deleting the message is not enough - the key must be revoked or rotated	35, 39
Forwarded a summary / shared a chat link	A summary retains confidential details; a link may expose the entire chat history and uploaded files	38

1. Task, Data, and Tool

1. What should I check before my first prompt?

- **Task:** what outcome is needed and what information is genuinely required. Fictional values suffice for a formula; a real contract requires a processing method officially

approved for such contracts.

- **Data and rights:** whether there is personal data (PII - information that identifies a person directly or indirectly), confidential information, client documents, or third-party copyrighted works, and whether you are entitled to use them this way.
- **Tool and mode:** which service will receive the information, through which account, and with which active integrations. If working independently, account for the law, client obligations, and service terms.

Full breakdown and sources

2. I am only asking to fix typos or translate text. Is that also a data transfer?

Yes, if the feature sends text to the cloud: during translation, summarization, spell-checking, or optical character recognition (OCR), the service receives the entire submitted snippet, image, or document. Find out where the operation is physically executed; there is no external transfer only when processing happens locally on your device or on a company server.

Full breakdown and sources

3. What should I do if the company has no AI policy or its rules contradict each other?

The absence of an AI policy creates neither a total ban nor permission to send secrets outside: applicable data protection laws, client contracts, confidentiality rules, and workplace device policies already apply.

- Clarify with your manager and InfoSec which tasks, data, and services are permitted, or which conflicting document takes precedence.
- While waiting for an answer, assign AI only tasks that involve no corporate data (if the service itself is permitted). If it is unknown whether the service itself is permitted, complete the task without AI.

Full breakdown and sources

2. Corporate Service

4. Are a corporate chat and an internal model the same thing?

No: a web address and company logo do not determine who actually receives the data. Architectures:

1. **On-premises / local LLM (language model):** deployed on the organization's own infrastructure.

2. **Enterprise cloud:** the organization manages accounts and settings, but an external provider performs the compute.
3. **"Wrapper":** the chat interface sits on a company server, but prompts are forwarded to a third-party organization.

A "built-in agent" in enterprise or local-vendor software may route to an external third-party cloud model (the interface vendor's jurisdiction does not determine where data is processed). Even with a local model, file OCR, web search, and integrations may remain external. Confirm with IT/InfoSec which data categories, in which mode, and with which features may be processed.

Full breakdown and sources

5. If a service is approved by the company, can I upload any work files to it?

No: authorization applies to the combination of **"task + data + service + mode."** If you rephrased a prompt within the same task and data scope, no new review is needed. If the data type, integrations, or assistant permissions change, check the procedure; if it does not cover the scenario, clear it with InfoSec/IT.

Full breakdown and sources

6. Do a paid subscription, a work email address, or a home computer change anything?

They may change certain settings, but they do not grant permission to transfer corporate data.

- **Personal paid subscription:** remains a consumer public offer, not an enterprise agreement with the company.
- **Work email:** links the account to the organization, but does not by itself make it a corporate account.
- **Home computer:** does not waive obligations to your employer or client; an unmanaged device may lack corporate security controls.

Use only the account and access method designated by your organization.

Full breakdown and sources

7. The provider promises not to train its model on my data. Is that enough?

No: model training, live request processing, and data storage are three distinct operations. A service may retain chat history, files, and logs, and grant its staff access to them for incident investigation and policy enforcement. Verify separately: training terms, log retention periods, staff and subcontractor access, server geography, deletion procedures, and add-on features (Microsoft Foundry: data privacy).

Full breakdown and sources

8. Does a domestic model or a local-region server automatically allow transferring work data?

By themselves, no. They do not automatically resolve questions of a lawful processing purpose, a Data Processing Agreement (DPA), provider staff access, backup locations, or onward transfers; moreover, where a file is stored and where the reply is computed may differ. Statutory data protection and security requirements still apply; processing on behalf of a controller requires formal DPA clauses and security measures. If the terms are unknown, consult your Data Protection Officer (DPO).

Full breakdown and sources

9. If I delete a chat, does the data disappear?

Not necessarily: information may remain in file storage, technical logs, backups, and the search index. Evaluate Temporary Chat against the documentation as well: no visible history does not mean connected tools kept no logs. Example: in Azure AI Search, without specific configuration, chunks of a deleted file remain in the search index. Verify retention periods and deletion mechanics before uploading.

Full breakdown and sources

3. Personal Data

10. Can personal data be processed in a company's internal LLM?

Yes, personal data may be processed in an internal language model if a lawful purpose and legal basis are established for the specific operation, role-based access is configured, security controls are in place, and an operating procedure is approved. Local model installation by itself grants no rights; an employee's right to view a file ≠ the right to process it in any manner (such as uploading it to a shared knowledge base or fine-tuning a shared bot).
The organization defines (under applicable data protection principles): purpose (compatible with the initial collection purpose); legal basis (mere automation convenience is not one); minimal data scope; circle of persons with access (including OCR, search, logs, and backups); security controls, retention periods, and deletion. Local hosting does not exempt the controller from internal governance policies, security controls, and required regulatory records or notifications.
For the employee: check the task against the procedure; if it is not covered or is unclear, consult the DPO before uploading. For a case already covered by the procedure, no new legal review is needed. Example: for a vacation schedule, submit full name, department, and dates - without passport scans, salaries, or medical diagnoses.

Full breakdown and sources

11. Is a person's consent always required to use AI?

Not always. Adopting AI does not automatically prove either a need for new consent or the sufficiency of old consent - check the actual purpose and processing method.

- **Within the company:** the basis may be performance of a contract with the individual or a statutory obligation.
- **Engaging an external provider:** as a general rule requires a valid legal basis or data subject authorization and a Data Processing Agreement (DPA) with mandatory security obligations; simply buying a subscription does not establish a compliant DPA.
- **Requirements for consent (under applicable data protection law):** specific, informed, freely given, and unambiguous; distinguishable and executed separately from other documents (and where required by law, in explicit written form); does not waive data minimization and security.

Full breakdown and sources

12. Can resumes, salaries, and health information be analyzed with AI?

- **Resumes and salaries (standard personal data):** analysis only in a specifically approved environment for a defined purpose (the right to store a resume for hiring does not authorize uploading it to any external AI).
- **Health information (special category):** baseline prohibition on processing except in strictly defined statutory cases under data protection law. Lawfully storing a sick-leave certificate in HR does not permit adding it to an AI knowledge base.
- **Practice:** when scheduling around an employee's illness, submit only the dates of absence without the diagnosis.

Full breakdown and sources

13. Can I send my own personal data from a work computer?

You can if company rules permit such use of the device and service, and the document contains no other people's data or employer secrets (a payslip, work email thread, or performance report may contain them). In a corporate AI account, chat history and storage may be visible to administrators; even from your own resume, remove passport details and home address first.

Full breakdown and sources

14. Can we record meetings and make AI transcripts? Is voice considered biometric data?

You may record and transcribe if the organization's procedure covers that meeting, tool, and processing. Simply announcing the recording at the start of the call is not enough, and consent from a corporate client entity does not override the rights of the individuals on the call.
An audio recording counts as biometric data under data protection law when voice characteristics are used to uniquely identify or authenticate a person. A text transcript is a different task, but the recording and transcript may still contain standard personal data, health details, and confidential work data.

- Use only a tool and mode for which the organization has verified the purpose, legal basis, retention periods, and recipient scope.
- If the procedure is unknown, check with your manager or process owner before recording; until answered or if no approved method exists, do not connect the bot and take minutes using the standard approved method.

Full breakdown and sources

15. Are uploading a file, adding it to a knowledge base, and training a model different actions?

Yes: the purpose, retention periods, and circle of persons with access differ fundamentally.

- **One-off chat analysis:** retention depends on service terms; the word "one-off" does not guarantee deletion after the reply.
- **RAG knowledge base (retrieving document chunks and attaching them to a prompt):** the model itself is not fine-tuned; who sees the data depends on access permission filtering (Azure AI Search).
- **Fine-tuning:** data modifies the model itself; the purpose and technical ability to cease using the data must be evaluated separately.

Approval for a one-off chat does not authorize adding a document to RAG or a fine-tuning dataset. Before launch, determine what derivative data is stored and how to fulfill statutory deletion and destruction requirements: deleting the source file does not guarantee deletion of derivative records.

Full breakdown and sources

16. Can we delegate decisions about candidates, employees, or clients to AI?

Drafting an assessment or performing a preliminary comparison against criteria is permitted (when data processing is authorized). Making a final decision that produces legal consequences or otherwise significantly affects an individual's rights and legitimate

interests based **solely on automated processing** is prohibited under data protection regulations, except with explicit consent or where expressly provided by statute (the evaluation procedure and consequences must be explained and the right to object and obtain human review ensured, while HR decisions must also comply with employment law). A responsible human specialist makes the final decision after personal review.

Full breakdown and sources

17. Can personal data be transferred to an external or cross-border AI service?

Only in an organization-approved workflow after legal and technical review. Neither an individual's consent, nor paying for a plan, nor keeping a local copy of a spreadsheet by itself makes the setup lawful.

- **Data residency and localization:** where data localization rules apply, primary recording, systematization, accumulation, storage, updating, and retrieval of personal data during collection must occur within the required jurisdiction (except under narrow statutory exemptions). Individual consent does not waive mandatory localization rules, and keeping a secondary local copy of the database does not cure an unlawful primary transfer.
- **Cross-border transfer:** before initiating a cross-border transfer in a commercial scenario, the controller must verify the destination jurisdiction's adequacy status or implement required transfer safeguards, obtain data-protection assurances from the external provider, and file any required transfer impact assessments or notifications with the data protection authority. Whether and under what terms the transfer may proceed depends on the destination jurisdiction's regulatory status and authority decisions.

Until the setup is confirmed by your DPO, legal team, and InfoSec, use an approved internal or compliant regional environment or ask questions without real personal data.

Full breakdown and sources

4. Internal, Client, and Third-Party Materials

18. Can I upload internal documents that have no confidentiality stamp?

You can if the content, tool, and action are approved for the task; the absence of a stamp does not by itself confirm this. A formal trade secret / commercial secret regime requires statutory and organizational measures (inventory of confidential information, access restriction and logging, NDAs, and confidentiality markings), but even without a stamp a file may contain personal data, client terms, financial plans, or information protected by an

employment contract. If the status is unknown, ask the information owner; until resolved, request a blank structure without confidential data if the tool and scenario are permitted.

Full breakdown and sources

19. Can I send a client's or contractor's document to AI?

Only after checking the client contract, non-disclosure agreement (NDA), and rights to personal data and copyrighted materials inside the file.

- Receiving a file from a client to perform work does not grant the right to share it with a third-party AI provider.
- Your NDA with an AI provider does not override a ban on third-party transfers in your client contract.
- Written permission from a corporate client does not replace a lawful basis for processing the personal data of individuals mentioned in the document. Individual consent is required wherever mandated by law.

Full breakdown and sources

20. What if the document, resume, or contact details are already published online?

Public publication does not authorize every subsequent form of processing. Publicly posted resumes and contacts remain personal data: respect any conditions and restrictions set by the data subject upon dissemination and ensure a lawful basis for your task under applicable data protection rules (although new consent is not required in all cases). Verify that you are uploading a clean public version of a document without internal colleague comments, and respect copyright.

Full breakdown and sources

21. Is it safe to send only summary numbers or a text description of a task?

Not always. The "average department salary" of a one-person unit directly reveals their income; revenue trends without a company name can expose confidential financials in a specific niche; a prompt asking *"Help me draft a branch-closure announcement"* reveals a management plan if the organization is identifiable from the account or chat history. Remove real names and rare circumstances, and use fictional numbers with the same mathematical structure for calculations.

Full breakdown and sources

22. Can I upload a corporate template, logo, or image?

You can if the content, usage rights, and processing in the service are approved. A presentation template often retains internal company details, old comments, and sample

data; an office photo may capture people's faces and monitor screens in the background. A stock image license may not permit uploading it to third-party neural networks. For design work, it is safer to describe brand colors and layout in words.

Full breakdown and sources

23. Can I upload a signed document or power of attorney?

The text itself is governed by general rules (personal data, confidentiality). **Never share a private electronic signature key, hardware token password, or verification codes with AI.** Do not confuse a public certificate with a private key: anyone with the private key and access credentials can sign documents in the owner's name. A language model neither technically nor legally verifies electronic signature validity - cryptographic tools are required for that.

Full breakdown and sources

24. Can I upload a purchased course, a paid article, or a recording of someone else's talk?

Purchase grants rights only within the license or contract and does not by itself authorize copying material into an external AI, uploading it to a shared knowledge base, or distributing an adaptation.

- Reproduction and other uses require the rightsholder's permission under applicable copyright law unless a statutory exception applies; the absence of an explicit ban is not permission.
- Exceptions for private personal use cannot be automatically extended to commercial team workflows.
- Permitted quotation or fair-dealing rules require a qualifying purpose, justified extent, and attribution of author and source, and do not legalize uploading an entire book or course to the cloud.

If you lack the required rights, discuss the topic with AI without uploading the file.

Full breakdown and sources

5. Preparing a File

25. Can I ask AI to anonymize a document and then use the cleaned version?

No, not if that tool is not already authorized to receive the raw data: to strip out names and prices, an external service first receives the entire original. Sanitize a document only in a tool approved for the original or manually, and always double-check the text by hand

afterward: software can miss indirect identifiers. If sanitizing removes the context needed for analysis, work in a system authorized for the original; if none exists, work without AI.

Full breakdown and sources

26. Is it enough to replace last names with "Employee 1"?

Anonymization means making it impossible without additional separately held information to determine which individual the data belongs to. Replacing last names with codes can be part of de-identification, but does not by itself confirm its adequacy or authorize uploading to AI. A person is easily identified by a unique title, exact dates, a rare situation, or a small team; moreover, removing names does not lift commercial secret status from prices and deal terms. Under data protection and de-identification standards, when introducing identifiers (replacing identifying details with codes), mapping tables and raw data are stored separately from the de-identified dataset, the key is not shared with third parties, and the adequacy of the method and actions are assessed and logged. Apply the method established by your organization (which determines the de-identification standards applicable to the workflow); if none exists or you are unsure which attributes to change, consult your DPO before sharing.

Full breakdown and sources

27. What if I send only a snippet, a screenshot, or altered numbers?

- **Snippets:** even a single paragraph can contain a key commercial term, and pasting piece by piece in one chat assembles nearly the entire document at the provider.
- **Screenshot:** a vision-capable model can read its contents, including names, phone numbers, the address bar, and neighboring windows accidentally captured on screen.
- **Altered numbers:** tweaking a couple of digits guarantees neither personal data anonymization nor removal of commercial secret protection. To debug a formula, create a completely fictional example.

Full breakdown and sources

28. What should I check in a file besides the visible text?

- **Office:** comments, unaccepted tracked changes, author properties, hidden rows and sheets, slide speaker notes, embedded objects, and Pivot Table caches. The built-in Document Inspector does not catch everything - for instance, it misses white text on a white background (Microsoft).
- **PDF:** a black rectangle over text does not remove the underlying text. Use the dedicated Redaction feature, apply redactions, save a copy, and test it via search and copy-paste (Adobe).
- After deleting rows from spreadsheets, verify that required formulas have not broken.

Full breakdown and sources

6. Integrations and Assistant Actions

29. Can I connect AI to my email, cloud drive, or CRM / customer database?

Only after evaluating the access scope with technical specialists. A connected assistant can autonomously find and pull in emails and files within its granted permissions - including restricted folders mistakenly shared on the drive "with everyone." Request access only to the specific folder/project needed and strictly in read-only mode (searching and summarizing do not require permission to send emails or modify records).

Full breakdown and sources

30. How do we ensure a knowledge base does not expose other people's documents?

Search must strictly enforce the specific user's permissions at data retrieval time (the same applies to isolating data across different clients). Converting text into embeddings (numerical representations) does not automatically anonymize data: risks of access-control bugs, cross-group leaks, and technically reconstructing source text from certain vector representations remain (OWASP LLM08:2025). The technical team must test the system across roles with different permission levels; if you find someone else's data in a reply, report it to the owner and do not forward it.

Full breakdown and sources

31. Do built-in AI features in a browser, phone, or office app also need to be checked?

Always check them: such features can transmit screen or open-document contents to the cloud in the background without manual copy-pasting.

- **Browser:** Copilot in Edge can read the active tab depending on settings, sometimes without a separate confirmation prompt (Microsoft).
- **Web search:** when search is enabled, the assistant can form queries from the accessible context of a work document - sending the topic and substantive fragments of the file to an external search engine (Microsoft).
- **Smartphone:** check on-screen assistant and smart keyboard settings before highlighting work text or photos.

Before enabling an assistant, check the combination of "task + data + tool + mode" against organizational rules. Verify the account type, active integrations, and enabled features. If the scenario is not covered, clear it with InfoSec or IT.

Full breakdown and sources

32. Can someone else's document trick the assistant into breaking the rules?

Yes, this is **Prompt Injection** - a hidden text command inside an incoming email or document (for instance, to find financial data in other files and send it outside) that the model may mistakenly treat as your high-priority instruction. A prompt instruction like "Do not send anything to anyone" does not guarantee protection; restrictions are enforced at the architecture and permission level (OWASP LLM01): an assistant summarizing emails has no technical need for access to all client folders or email-sending capabilities.

Full breakdown and sources

33. Should we allow AI to send emails, modify records, and delete files?

Only for actions approved in the workflow; for operations with significant consequences, we recommend human confirmation before execution (OWASP LLM06:2025 Excessive Agency).

- **Least privilege:** permission to create an email draft rather than send it automatically.
- **Human review:** recipient, text, attachments, and proposed changes are verified before confirmation. The organization defines the confirmation procedure and responsible roles.
- **Target-system controls:** permissions are restricted technically, actions are logged, and rollback is provided for reversible operations (reliably recalling an already delivered email cannot be guaranteed).

Full breakdown and sources

34. Can I show source code, error messages, and technical logs to AI?

You can if the tool and submitted data are approved. Always strip out active passwords and keys; exclude other details if they cannot be shared in the chosen mode.

Raw technical logs often retain user queries with full names, internal IP addresses, and tokens; source code may contain proprietary algorithms or belong to a client; AI plugins in IDEs can automatically pull neighboring project files into context (check the exclusion list with your administrator).

For debugging, prepare a minimal reproducible example and replace confidential values with fictional ones wherever doing so preserves the cause of the error.

Full breakdown and sources

7. Passwords and Keys

35. Can I paste a password or key into a chat so AI can help me configure something?

Absolutely not under any circumstances. Secrets include passwords, API keys (secret connection codes), access and session tokens, one-time and recovery codes, and private electronic signature keys.

- Use a placeholder like YOUR_KEY in your prompt text, and enter the real value yourself into a secure field of an approved system or corporate secrets vault.
- If the assistant must call a service itself, a technical specialist configures the key inside the application so it never enters prompts, outputs, or logs, with least-privilege permissions (OWASP Secrets Management).
- If a secret lands in a chat, deleting the message is not enough - the key must be immediately revoked or rotated (see Question 39).

Full breakdown and sources

36. What if I put secrets in a separate file and ask AI to store it encrypted?

Uploading an unencrypted file for encryption gives the service your secrets in plaintext right away; sending an encrypted archive along with its password lets it decrypt the file immediately. For a local AI agent with computer access, a "separate folder" does not mean "inaccessible": .env files, environment variables, or files outside the project can be read by the assistant's tools. A text request saying "do not read" does not block access technically.

1. Store secrets in a corporate password manager or secrets vault; if a file is required, encrypt it with a vetted local tool **before sending** (do not use encryption schemes invented by the language model itself).
2. Keep the decryption key separate from the data and out of the assistant's reach.
3. Have a technical specialist configure direct secret retrieval by the application itself and technically block the model from reading key files or outputting their values in replies and logs (OWASP Cryptographic Storage).

Full breakdown and sources

8. Verifying and Sharing the Output

37. If a service is secure, can I trust its conclusions?

No: **data-processing security and factual accuracy of the output are verified separately**.

- Check calculations against source numbers and formulas, legal citations against current statutory text, and conclusions against the document text; open and verify

every link provided.

- When told "no errors found," clarify what specifically was checked.
- Run code and commands in an isolated test environment first; for decisions with real consequences, assign a domain specialist to review.

Full breakdown and sources

38. Can I forward an AI reply to colleagues or share a public link to a chat?

Check what information remains in the reply, with whom it may be shared, and by what method. A brief summary can retain names, amounts, and the core of a commercial secret - model paraphrasing does not lift confidentiality restrictions. When creating a shared link to a conversation, check what exactly it exposes (a single reply, the entire chat history, or uploaded source files) and who can open it.

Full breakdown and sources

9. Mistakes or Missing Rules

39. What should I do if data has already been sent to the wrong place?

Employee's immediate actions

Immediately stop the transfer and notify InfoSec, your Data Protection Officer (DPO), and your manager.

1. **Details for investigation:** time, service, account, data scope, operation/chat ID. **Do not re-send the confidential file itself or the exposed password in your report email!**
2. **Revoking secrets:** if keys or passwords are compromised, immediately arrange to revoke (rotate) them, update dependent services, and audit usage logs (GitHub guide).

Subsequent assessment and notifications

3. **Incident assessment and regulatory breach notification deadlines:** not every mistaken submission is a reportable public breach, but the DPO promptly assesses the situation against statutory breach-notification thresholds. If an unlawful or accidental transfer/access to personal data resulting in a risk to individuals' rights is established, the controller notifies the relevant data protection authority within the statutory reporting window (often **within 24 to 72 hours** of discovery) and reports internal investigation results.
4. **Client notification (when processing on their behalf):** strictly follow the contractual procedure; do not promise "complete data deletion" until officially confirmed by the service provider.

Full breakdown and sources

40. What should I do if the approved tool is weak and the deadline is burning?

Urgency does not suspend security rules. Split the task:

- If using an external service for such a task is permitted by your organization, use it to prepare the structure, checklist of questions, or logic verification **on a fictional example**, and enter real data manually inside an approved corporate system.
- If the task cannot be done without confidential data and no suitable AI is approved, complete the work without AI and submit a request to management and IT (stating the required business outcome, data involved, and what the current solution lacks).
- If the issue recurs regularly, evaluate with your manager and IT another mode/service, a local deployment, a workflow change, or conventional automation (accounting for cost and data requirements: local hosting does not by itself make a solution better or compliant).

Full breakdown and sources

41. How do I ask technical specialists, legal, or InfoSec a substantive question?

Send technical specialists, legal, or InfoSec an operation description card (without attaching the confidential file itself in a shared chat):

- **Task:** what outcome is needed and why.
- **Data:** where it comes from and what it contains (personal data, client NDA restrictions, third-party copyrighted materials).
- **Tool:** service name, plan, account type, and operating mode.
- **Assistant actions:** what specifically it reads, and whether it can modify or send data.
- **Processing terms:** who receives the data, where and how long it is stored, and which external features are enabled *(indicate which of these terms you need help clarifying)*.
- **Output verification:** who will verify the reply and who will have access to it.
- **Question:** is the operation permissible and what conditions must be met to launch it?

Full breakdown and sources

42. What should a manager establish for their team first?

1. A list of approved combinations of "Task + Data Category + Service + Mode" with examples and a clear boundary for new approvals.
2. Corporate accounts and access rules: use multi-factor authentication (two-factor login verification) wherever supported, and make it mandatory in organizational policy.
3. Rules for retention, deletion, and connecting new data sources.

4. Requirements for human verification of replies and human confirmation of assistant actions.
5. A fast, clear channel for reporting incidents.
6. A designated owner responsible for reviewing rules (when changing providers, plans, contracts, training policies, retention periods, integrations, or after incidents).

Practical reference (from an international organization rather than a binding statute): World Intellectual Property Organization (WIPO) Generative AI Guidelines.

Full breakdown and sources

Checklist: Seven Checks Before Working with Data in AI

1. **Task:** what information is genuinely required for the result and whether a fictional example will suffice.
2. **Data rights:** whether the purpose and legal basis for personal data processing, confidentiality restrictions, client contracts, and copyright are covered; clarify anything uncertain with the responsible owner before sharing.
3. **Processing method:** permitted service, account type, operating mode, active integrations, and retention and deletion rules.
4. **Minimum material:** remove anything superfluous, **never paste passwords or keys**, inspect hidden layers (hidden sheets and rows, unaccepted tracked changes, Pivot Table caches, true PDF Redaction instead of black boxes), and preserve the meaning of the task.
5. **Access and actions:** only required folders and least privilege (read-only instead of automatic email sending or record modification).
6. **Output verification:** who will check facts, calculations, legal rules, and code; with whom the reply may be shared; which actions require human confirmation.
7. **Incident response:** keep contact details for InfoSec and your Data Protection Officer at hand in case of a mistaken submission or key compromise.

If a material security condition is unknown, clarify it before sending sensitive data, and complete the general portion of the task without confidential information in an approved tool.
Extended Mind · extended-mind.ai