

AI and Trade Secret Protection Checklist



Reasonable measures and precautions are essential to protecting confidential intellectual property. For trade secrets, they also bear directly on whether the law protects the information as a secret. As AI changes how information is accessed, processed, and shared, organizations need to reassess whether their existing protections remain adequate.

What is reasonable evolves with the risks, the available safeguards, and the circumstances of the organization. A measure suited to one business or system may be insufficient or unnecessary for another. Effective precautions should reflect the value and sensitivity of the information, how it is used, who can reach it, and the consequences of disclosure.

This checklist provides broad guidelines for selecting measures and developing practices tailored to your organization. Use it to identify gaps, assign responsibility, and document the protections put into practice. The four priority groups offer a starting point for that review; a specialized safeguard may deserve immediate attention when it addresses a significant risk to your business.

Priority 1 Essential foundations

1. Assess existing systems before deploying AI

Review the organization's existing permissions, shared folders, cloud configuration, and transfer controls before adding an AI assistant or agent. Correct excessive access and broad sharing links. An assistant operating with an employee's permissions may make sensitive information much easier to locate and aggregate, exposing weaknesses that existed before the AI deployment. Create the right environment with appropriate governance before deploying AI functions.

2. Identify the specific information worth protecting

Specify the formula, process, design, dataset, model, or other information whose disclosure would cause competitive harm. Identify and inventory where and how information is stored or saved. Record what secrecy contributes to its commercial value and distinguish it from information already available to the public. Start with the information that matters most to the business, then identify the people and systems that can reach it.

3. Assign information owners and handling rules

Each category of sensitive information should have an owner and a classification that determines permitted access, storage, sharing, and AI use. Apply those rules to copies, extracts, backups, and derived materials. Make the owner responsible for approving exceptions and reviewing whether the rules still fit the information and its use.

4. Map where sensitive information travels

Trace information through employees, repositories, vendors, AI models, connected applications, and external recipients. Include processing, temporary storage, backups, and generated outputs. Identify where information leaves a controlled environment and which protections apply at each destination. Use this map to select controls and document why they address the particular exposure.

5. Identify who makes and implements security decisions

Specify who approves AI uses, controls access to each repository, administers technical restrictions, and responds to alerts. Define the respective roles of legal, security, IT, and the business owner. Establish an escalation path for disagreements and exceptions so that approval responsibility does not disappear between departments or vendors.

6. Require authenticated individual access

Keep confidential repositories behind authenticated access and use identifiable accounts rather than shared employee logins. Apply company sign-on and multifactor authentication where appropriate. Access records should allow the organization to determine who used a system and whether that person had permission to reach the information involved.

7. Separate permission to read from permission to act

Limit employees and agents to the information their tasks require. Distinguish reading from downloading, exporting, modifying, deleting, or sending information. An assistant that summarizes documents may need reading access without authority to edit or transfer. Require the resource owner's approval before granting or expanding access.

8. Approve each AI tool and its intended use

Identify the permitted users, business purpose, information categories, and service configuration for each approved AI use. Maintain an accessible list of approved tools and their limits. Include AI features embedded in existing business software. Approval of a product should not automatically authorize every feature, connected service, or use of confidential information.

9. Give employees concrete examples of restricted AI inputs

Identify the information that unauthorized AI tools must not receive: nonpublic plans, financial data, customer or employee information, source code, models, configurations, logs, and security architecture. Include passwords and access keys. Address attachments, screenshots, metadata, and copied excerpts. Specify separately which sensitive information an approved enterprise or internal service may process under its authorized configuration.

10. Verify the promised controls are enabled for third-party services

Confirm that contractual protections and security features apply to the purchased account, selected features, and actual configuration. Check settings governing retention, access, external sharing, and model training where available. Keep evidence of the configuration and review changes. An enterprise subscription alone does not establish that every relevant safeguard is operating.

11. Enforce data loss prevention across applications and agents

Use data loss prevention, or DLP, controls to block or flag unauthorized copying, uploads, sharing, and unusual bulk downloads. Where supported, connect restrictions to information classifications and apply them across agents and applications. Combine these central controls with limits on each agent's access. Test the transfer routes the organization actually uses.

12. Check authorization when an agent requests an action

Require the receiving system to verify permission before executing an agent's request to read, send, change, or delete information. Enforce limits in the surrounding software and access controls. Instructions in a prompt can communicate the task, but should not be the sole barrier preventing an agent from exceeding its authority.

Priority 2 High priority operating controls

13. Give applications and agents separate limited credentials

Give each application or agent credentials limited to its assigned work, rather than sharing an employee's password or a broad administrator credential. Keep passwords, tokens, and keys out of prompts, shared documents, and ordinary logs. Use an approved credential-management mechanism and maintain a way to revoke an individual agent's access promptly.

14. Remove unused access and review permissions regularly

Disable or remove dormant accounts, unused connectors, and unnecessary credentials. Recheck permissions after job transfers, project changes, and the end of vendor engagements. Review outside collaborators and broadly shared links. Preserve records of who approved access and when it ended, even when the account itself is removed.

15. Provide approved devices and communication channels

Specify which equipment, accounts, applications, storage services, and communication channels may handle company information. Address personal email, messaging, cloud drives, and removable media expressly. Give employees a workable approved alternative for legitimate tasks. Keep company records in company systems and explain that a personal folder label does not authorize retaining work material.

16. Bring employee-built agents and departmental purchases into review

Apply approval requirements to employees building agents and to departments purchasing AI-enabled software independently. Require appropriate security and business review before those tools connect to company systems. Provide a practical approval path so employees can experiment within defined limits instead of creating unreviewed access to company information.

17. Review new AI features and applications

Require review before trials, research uses, new integrations, or expanded uses of an existing tool. Use public or synthetic material until confidential inputs receive approval. Revisit approval when functionality, data use, retention, or access changes, including when a software update introduces an AI feature or changes what an existing feature can reach.

18. Minimize the information supplied to each task

Limit an integration to the particular fields and content needed for its function. A tool creating contact records from email may need a sender's name, address, and signature without needing message bodies or attachments containing pricing or confidential advice. Review what the connection actually transmits rather than relying on the feature's description.

19. Define vendor rights in inputs and outputs

Address ownership and permitted use of submitted information and generated outputs, including generated code. Restrict use of confidential material for training, fine-tuning, product improvement, or other customers' benefit unless specifically authorized. Address human review and access, including who may inspect submitted information and for what purposes.

20. Address vendor retention and exit obligations

Review storage periods, deletion, backups, and onward disclosure. Specify incident notice, investigation cooperation, available logs, and return or deletion when the relationship ends. Determine how the company will obtain necessary records before access expires. Retain the operative agreement, amendments, and settings that applied when the information entered the service.

21. Require meaningful human approval for sensitive actions

Identify exports, external disclosures, permission changes, and other consequential actions requiring review before execution. Show the reviewer the actual information, intended recipient, and proposed action. Ensure that approval covers that specific action and that the system cannot execute a materially different action under the same approval.

22. Define confidentiality and permitted uses in agreements

Identify the protected information, permitted purpose, authorized recipients, and restrictions on onward disclosure. Address AI processing, training, distillation, and analysis expressly where relevant to the relationship. Secure the necessary assent and retain operative versions. Connect contractual return, deletion, and termination duties to a process that the business can actually carry out.

23. Train employees using situations they will encounter

Require practical training tied to the work of engineering, sales, legal, finance, and other relevant teams. Explain approved tools, prohibited inputs, connector permissions, and how to request an exception. Use concrete examples of accidental disclosure and employee-built agents. Retain completion records and refresh training when tools, policies, or responsibilities change.

24. Make offboarding cover AI access and retained material

Revoke employee and contractor accounts, agent credentials, connectors, and physical access at the appropriate time. Recover company material and address copies in personal accounts, devices, or AI workspaces. Review continuing access through shared links and third-party services. Preserve necessary business and investigation records while carrying out return and deletion obligations.

25. Prepare an incident response process

Assign legal, security, IT, and business responsibilities before an incident. Provide options to stop an agent, isolate a repository, restrict compromised credentials, and contact the provider. Identify the affected information, recipients, and continuing access. Preserve relevant files, prompts, outputs, logs, permissions, and settings while containing the incident. Assess notice obligations, recovery options, and urgent relief as the facts require.

26. Document implemented controls and approved exceptions

Retain policy versions, permissions, approvals, training records, configuration evidence, test results, and corrective work. Record who approved an exception, its purpose, duration, and safeguards. Keep evidence of relevant disclosures and accompanying restrictions. The record should show what the company actually did to protect the information, including how it addressed identified gaps.

Priority 3 Supporting safeguards and oversight

27. Use classification tools to help locate sensitive material

Consider automated scanning and labeling of document stores to identify confidential information and support technical handling restrictions. Review the results for missed or incorrect classifications. Labels should connect to actual access and transfer controls where the systems support them; attaching a label alone does not prevent disclosure.

28. Control installations and external devices

Restrict unauthorized USB storage and other external devices. Limit software installations to approved sources, including browser extensions, AI add-ins, and applications that synchronize files. Maintain an inventory, review the permissions these tools receive, and remove unused software. Document exceptions and test whether the restrictions remain effective.

29. Trace retrieval through external model processing

For retrieval-augmented generation, or RAG, identify which internal passages the system retrieves and where it sends them for processing. Confirm that access restrictions remain effective during retrieval and that any receiving model service has appropriate protections. An internal source document can still be disclosed externally when retrieved content becomes part of a model request.

30. Review connectors and plug-ins individually

Assess each connector's permissions, recipients, and transmitted information. Include connectors using the Model Context Protocol, or MCP, which lets AI systems interact with other tools and data sources. Check whether confidential content appears in search queries or tool requests. Approval of the underlying assistant does not establish that every connected service provides the same protections.

31. Match safeguards to the hosting arrangement

Identify whether the model runs internally, through a vendor, or through a combination of services, and allocate security responsibilities accordingly. For internal deployments, review the model's source, supporting software, vulnerabilities, and external connections. Local hosting can improve control, but an internally hosted model may still send information outside the organization through connected tools.

32. Review processing locations and subprocessors

Identify where information is processed and which subcontractors or subprocessors can receive it. Review available security documentation and the contractual restrictions covering each recipient. Address changes to the processing chain and ensure that confidentiality, permitted-use, and incident-response obligations extend to the entities that actually handle the information.

33. Ask managed service providers specific operational questions

For outsourced IT or AI administration, establish who configures controls, monitors alerts, and investigates incidents. Ask what activity is logged, how long records remain available, and how the company can obtain them. Address audit rights and cooperation. Require evidence of the controls actually implemented and identify responsibilities that remain with the company.

34. Limit information exchange between agents

Define which agents may communicate, what information they may exchange, and for what purpose. Prevent an agent from obtaining broader access by borrowing another agent's credentials or asking it to perform a restricted action. Apply access checks to delegated requests and specify who can authorize an exception.

35. Test for malicious instructions in external material

An agent may encounter instructions embedded in a website, email, document, or tool response that attempt to redirect its behavior. This is a prompt-injection risk. Test whether such material can induce unauthorized retrieval or disclosure. Restrict outbound destinations and available tools, especially where an agent combines confidential internal access with external browsing. See [OWASP AI Agent Security](#).

36. Test agent boundaries and provide a working shutdown

Test whether an agent can cross project boundaries, obtain broader authority, or disclose information through connected tools. Include misleading documents and requests passed through other agents. Provide a working way to stop execution and revoke relevant credentials. Recheck these controls after material changes to the agent or its environment.

37. Route alerts to people who can act

Assign responsibility for reviewing suspicious transfers, excessive querying, unusual agent activity, and departures from approved tasks. Define escalation thresholds and response coverage. Test whether alerts reach the right people and whether they can investigate or suspend the affected workflow. Record the event, investigation, and action taken.

38. Keep useful logs without creating an uncontrolled repository

Determine which records are needed to reconstruct access, tool requests, disclosures, and approvals, and retain them for a useful period. Restrict access because logs may themselves contain confidential information. Avoid routinely recording credentials or unnecessary sensitive content. Confirm that vendor-held records will remain obtainable and preserve relevant material when an incident or preservation duty arises.

39. Check permissions for using outside information

Review confidentiality obligations, licenses, and permitted-use restrictions before submitting another party's material to an AI service or using it to develop a system. Technical access does not establish permission for that new use. Obtain any required authorization and ensure that the receiving service's terms are compatible with the company's obligations.

40. Include the required DTSA immunity notice

Review confidentiality agreements with employees and individual contractors or consultants for the whistleblower-immunity notice required by the Defend Trade Secrets Act. The notice requirement applies to agreements entered into or updated after May 11, 2016, and permits the statutory policy cross-reference. Omission bars specified DTSA exemplary-damages and attorney-fee remedies against the individual who did not receive notice. See 18 U.S.C. § 1833(b).

41. Make mistakes easy to report promptly

Give employees a clear route to report an unintended upload, excessive access, or suspicious agent action. Explain whom to contact and what information to preserve. Instruct employees to coordinate with the response team before deleting chats, closing accounts, or attempting other cleanup that could destroy evidence needed to determine what happened.

42. Reassess protections and close identified gaps

Assign corrective tasks, owners, and deadlines after reviews or incidents. Use qualified security personnel and independent assessment where warranted to test implementation. Reassess when employees, clients, vendors, data sources, or system capabilities change. Confirm that remediation worked and document decisions to accept remaining risk. Prioritize the information whose disclosure would cause the greatest competitive harm.

Priority 4 Specialized and situational safeguards

43. Apply travel and location safeguards to the actual risk

Review travel destinations before departure and identify where additional controls are warranted. Consider designated travel devices containing only necessary information and limits on remote access. Obtain current legal review for sanctions-related restrictions. On return, revoke temporary access and have IT assess devices before ordinary reuse. Record approved exceptions and accompanying safeguards.

44. Control physical access to sensitive locations

Use badges, visitor procedures, and entry records for sensitive offices, laboratories, production areas, and engineering workspaces. Limit visitors and contractors to the areas needed for their work. Identify who authorized access and ensure that applicable confidentiality obligations are in place before sensitive information becomes accessible.

45. Control recording and observation on site

Set rules for photography, recordings, sample removal, and unattended screens. Explain them to employees, visitors, and contractors. Review exposure during plant tours, demonstrations, construction, and equipment service. Consider what scans, images, or samples could reveal when analyzed with AI, including information not obvious to a casual observer.

46. Control access to proprietary models and APIs

Identify users of application programming interfaces, or APIs, and define permitted uses of model outputs. Address bulk collection, model training, and circumvention expressly in applicable agreements. Select request, volume, and export limits suited to the service. Consider what repeated queries could reveal about proprietary behavior, protected datasets, or other confidential assets.

47. Detect and investigate systematic model extraction

Monitor repeated sampling, coordinated accounts, systematic coverage of inputs, and unusual volumes that may indicate extraction or unauthorized distillation. Use a defined process to investigate and, where appropriate, suspend access. Preserve relevant requests, responses, account history, and applicable terms when warranted. Query volume alone should trigger assessment rather than an automatic legal conclusion.

48. Assess what a finished product reveals

Ask engineering what a purchaser could learn through scanning, testing, teardown, or AI-assisted analysis. Distinguish observable product features from undisclosed manufacturing methods and process parameters. Revisit the assessment as analytical tools improve. Use the findings to decide which information can realistically remain confidential after a product reaches the market.

49. Limit disclosure of manufacturing and design information

Give suppliers and contractors only the drawings, samples, specifications, and process information needed for their work. Control prototypes, supporting design files, plant visits, and process records. Record the materials provided and the restrictions accompanying access. Consider whether different suppliers can perform their tasks without receiving the complete confidential process or design.

50. Choose protection before exposing product information

Evaluate confidentiality terms, enforceable restrictions on analysis or use, and patent strategy before disclosing sensitive designs or processes. Reassess reliance on trade secrecy when lawful examination of a finished product can reveal the information. Coordinate legal and engineering review so that commercial release and supplier disclosures reflect the protection strategy.

Yegparian PLLC

Yegparian PLLC represents companies and individuals in complex commercial disputes, including technology and trade secret litigation. To discuss how these precautions relate to your organization's information and practices, contact Varant Yegparian.

Varant Yegparian | Founding Partner

vyegparian@yegparianpllc.com | 512.357.8297

5900 Balcones Drive, Suite 28941, Austin, Texas 78731

yegparianpllc.com