



Shadow AI Risk Sprint Kit

A general informational checklist and action planner to help security and IT leaders discuss potential unsanctioned AI use, oversharing risk, data protection considerations, and governance conversations with appropriate internal stakeholders.



Overview

The Shadow AI Risk Sprint Kit is a general informational framework designed to support structured conversations about potential unsanctioned AI use, data handling considerations, and governance approaches. This asset is intended to help teams discuss visibility, prioritization, and coordination of AI-related decisions across appropriate stakeholders.

It should be adapted to the organization's environment and reviewed against current Microsoft documentation, licensing, availability, and applicable legal, privacy, compliance, security, IT, procurement, and business requirements before use.

Contents

3

Shadow AI risk quiz

4

Five day sprint map

5

Action planner

6

Learn more

When to use this sprint

This sprint can be used as a facilitation tool when AI applications, copilots, agents, or extensions may be adopted at different speeds than internal review processes. It is intended to support discussion and prioritization and does not prescribe required actions.

Suggested sprint outcomes

- Develop a working inventory for discussion of AI tools and data touch-points.
- Identify and prioritize areas for further review.
- Explore potential data protection and access considerations for stakeholder review.
- Discuss a possible governance approach including ownership and cadence.

Disclaimer

This content is provided for informational purposes only. It does not constitute legal, regulatory, compliance, or technical advice and does not guarantee detection, protection, risk reduction, or compliance outcomes. Organizations are responsible for independently evaluating their obligations, risk tolerance, and technical configurations, including product capabilities, availability, and licensing.

Start here: quick shadow AI risk quiz

Use responses as directional inputs for discussion only. Multiple higher-concern responses may suggest topics for additional visibility, stakeholder review, or prioritization. Multiple moderate-concern responses may suggest areas where ownership, policy clarity, or evidence routines could be useful. Mostly lower-concern responses may indicate that the sprint could be used as a readiness discussion and optimization exercise.

This quiz is not a formal assessment, scoring methodology, audit, or compliance determination, and any next steps should reflect each organization's legal, regulatory, contractual, technical, and operational context.

Discussion question	Lower concern	Moderate concern	Higher Concern
Do you have a current view of AI apps, extensions, copilots, plugins, and agents in use?	☐ Inventory appears current.	☐ Partial visibility.	☐ No reliable view.
Can you see what sensitive data AI tools may access?	☐ Access appears mapped.	☐ Core tools only.	☐ Unknown or inconsistent.
Are approved AI tools and paths clear?	☐ Path appears clear.	☐ Guidance is uneven.	☐ Teams choose tools independently.
Are controls mapped to AI-related data paths?	☐ Controls appear mapped.	☐ Some workflows covered.	☐ Not mapped to AI use.
Is there evidence for AI governance decisions?	☐ Owner and cadence appear defined.	☐ Evidence is ad hoc.	☐ No consistent proof path.

Five day sprint map

Use responses as directional inputs for discussion only. Multiple higher-concern responses may suggest topics for additional visibility, stakeholder review, or prioritization. Multiple moderate-concern responses may suggest areas where ownership, policy clarity, or evidence routines could be useful. Mostly lower-concern responses may indicate that the sprint could be used as a readiness discussion and optimization exercise.

This quiz is not a formal assessment, scoring methodology, audit, or compliance determination, and any next steps should reflect each organization’s legal, regulatory, contractual, technical, and operational context.

Day	Focus	Actions	Outputs	Checklist
1	Discuss potential unsanctioned AI use.	Review signals, separate known from unknown tools, and capture business context.	Draft inventory, tool list, and follow-up owner.	☐ Sources identified ☐ Unknown tools noted ☐ Gaps captured
2	Discuss potential review areas.	Group use cases by sensitivity, access, sharing, automation, and compliance considerations.	Prioritized discussion inventory and stakeholder review list.	☐ Data paths noted ☐ Higher-concern tools identified ☐ Owners suggested
3	Protect data and access.	Discuss baseline controls, data protection scenarios, and access or endpoint considerations.	Control considerations, policy items, and monitoring areas.	☐ Controls discussed ☐ DLP or labels noted ☐ Monitoring gaps captured
4	Govern AI apps and agents.	Draft ownership, intake, exception paths, and evidence needs.	Roles, review path, and evidence register concept.	☐ Owners identified ☐ Exceptions captured ☐ Register started
5	Turn the sprint into a repeatable rhythm.	Confirm priorities, draft guidance, and propose a 30-day cadence.	Planning priorities, guidance outline, readout topics, and cadence.	☐ Priorities documented ☐ Guidance drafted ☐ Cadence proposed

30-day action planner map

Recommended leadership readout. Example topics include AI usage patterns, data considerations, prioritization of actions, stakeholder alignment, and near-term planning.

Priority	Owner	Action	Evidence / success signal	Due
1				
2				
3				
4				
5				
6				

Learn more from Microsoft

Explore Microsoft Security resources on securing and governing AI use, including guidance for data protection, identity and access, endpoint security, and responsible AI adoption.

Get ahead of risk, so you move securely forward with [Microsoft Security](#)
Explore Frontier Transformation—where intelligence meets trust with [Microsoft AI](#)
Secure AI and data security for productivity with [Microsoft 365](#)

