

SHADOW AI

IN PUBLIC SECTOR

Governing the Unofficial Use of Generative AI

Mike Bell, CISSP, MLIS
Senior Cybersecurity Advisor



**BUILDING CYBERSAFE
COMMUNITIES**



Quick Audience Poll

Show of hands — how many of you...

Poll Question 1



Have used ChatGPT or another Generative AI tool?

Poll Question 2



Have used AI for work-related tasks?

Poll Question 3



Believe employees are already using AI?

Poll Question 4



Have an organizational Generative AI use policy?

Discussion: How much AI use is occurring without leadership visibility?

What Is Shadow AI?



Use of AI tools outside approved organizational governance.

COMMON EXAMPLES

- Personal ChatGPT accounts
- AI being used for work tasks
- AI browser assistants
- AI email rewriting tools
- AI meeting transcription services
- Embedded AI features in other apps


Why Does It Happen?

Most Shadow AI originates from employees trying to save time and work more efficiently.

They aren't acting with malicious intent — they're solving real problems with available tools.

The gap isn't intent. It's governance.

Where is AI actually getting its information from?



Why Shadow AI Matters

AI adoption is moving faster than governance

77%

Of employees are sharing sensitive organizational data through AI tools

(eSecurityPlanet.com, 10/9/2025, Ken Underhill)

68%

Of employees use AI tools without the organization's knowledge

(Watchguard Technologies, 2026 report)

3.8

Average number of GenAI pastes/day that include sensitive organizational data

(eSecurityPlanet.com, 10/9/2025, Ken Underhill)

Employees are already using AI to:

- Summarize documents
- Rewrite emails
- Analyze spreadsheets
- Create reports
- Generate meeting notes



Leadership may not know what tools are being used or where information is stored.

This is a governance gap — not a technology gap.

Does your organization manage AI use?



Most organizations are way behind the curve on managing AI use at work.

THE CURRENT AI PRODUCTS:

- Recent tests show AI favors guessing vs. admitting “I Don’t Know”
- AI training data is riddled with inaccuracies (GIGO, Garbage In, Garbage Out)
- LLM’s are designed to please and validate, think “Digital Yes Man”
- Human language is complicated and AI doesn’t understand sarcasm, intention, or unspoken assumptions

Some recent statistics:

AI tools are deployed at 73% of organizations, yet only 7% have real-time governance.

63% describe their AI governance as reactive or still developing.

39% have already had an AI-related “near miss” involving data exposure.

Let's Examine 3 AI Work Scenarios

- Editing emails using AI
- Using AI to create a summary of a large data set
- Use of automated AI tools built inside applications

Scenario 1: “Can You Clean This Up For Me?”

*A manager forwards an email thread to their administrative assistant, asking:
“Before this goes to HR, can you clean it up a bit for me?”*

Trying to save time, the admin:

- *Copies the entire email chain (including prior comments and manager notes)*
- *Pastes it into a public AI tool using their personal account*
- *Requests: “Summarize this thread and rewrite this to sound more professional”*

What risks do we see here?

Let's examine the possible risks:

- What **sensitive details** were likely included in the email thread?
- What happens to this content once it's pasted into a **public AI tool**?
- Whose account was used — **personal** or work-managed?
- Why does that **matter**?
- Could the AI **store or reuse** any part of this email?
- Would the manager or HR **know this was run through AI**?
- If HR, Legal, or the employee later requests this thread, **where does the AI copy live**?

Scenario 2: Summarizing Data

- A public works supervisor sends an employee a work-order export and says: *“We need totals for the year, can you categorize this by work type, hours, and cost for the annual report?”*
- Because it is a huge spreadsheet, the employee:
 - Uploads the entire spreadsheet into a public AI tool using their personal account
 - Asks the AI: *“Summarize by work type, total hours, and total cost by year”*

The spreadsheet contains:

- | | |
|--|---|
| <ul style="list-style-type: none">• Customer Names• Service Addresses• Phone Numbers• Utility Account Numbers | <ul style="list-style-type: none">• Work-Order Notes• Hours• Materials Used• Job Costs |
|--|---|

Do we have a problem?

Is this data confidential?

Scenario 2: What can go wrong?

- Which data elements in this spreadsheet are non-public or sensitive?
- What happens to this customer information once it's uploaded to a public AI tool?
- Does using a personal account change how the data is stored or retained?
- Could the uploaded file be logged, reused, or retained by the AI provider?
- Would the supervisor, or IT, have any idea the data left your systems?
- How would this impact open records, auditability, or citizen trust if discovered later?

Scenario 3: The Hidden Data Leak

A department head and two staff join a video call with an outside project partner to discuss project issues including budget overruns, missed deliverables, and schedule delays.

The external partner hosts the meeting using a free video conferencing platform that includes automatic AI note-taking. As the meeting starts, a small banner displays: “AI Notes Enabled — Transcription Active.” Your staff assume it is just part of the partner’s system and continue the call.

- In the background, the AI tool:
 - **Records and transcribes the entire meeting**
 - Store the recording and transcripts in the external partner’s account
 - Generates an AI-created summary of the discussion
 - Retains the audio/transcript data under terms that were not reviewed or approved

What risks do we see here?

Let's examine the possible risks:

- The external party controls the recordings, your staff have **no way to delete, restrict, or audit** the transcript. They might be retained indefinitely by the outside platform.
- The partner's AI platform terms may allow:
 - **Reuse of the transcript** for AI training
 - Sharing or internal analysis
 - Long-term retention
- Open Records Act (Texas Public Information Act) implications:
 - Your organization may be legally responsible for a record it does not possess.
 - Requests could expose the organization to compliance gaps.
- Recordings may contain contractual information that should remain under your custody.

What Do These Scenarios Have In Common?



None involved malicious intent.

Employees were simply trying to save time, increase productivity, improve communication, and complete routine work.

The Real Issue

Lack of visibility and governance — not employee wrongdoing — is the root challenge facing public sector organizations today. When tools evolve faster than policy, the gap becomes a liability.

'Invisible adoption' is often the greatest AI challenge facing organizations today.

Five Practical AI Guardrails

Goal: Enable responsible use — not prohibit innovation

1

Establish a Generative AI Use Policy

Define what tools are permitted, under what conditions, and by whom.

2

Use Organization-Managed Accounts

Organizational accounts keep data under institutional control and audit trails.

3

Avoid Uploading Sensitive Information

Train staff to recognize data that should never be shared with external AI tools.

4

Understand Vendor Settings & Defaults

Review data retention, sharing, and training opt-out settings before deployment.

5

Verify AI-Generated Content Before Use

AI outputs can be inaccurate. Require human review for any official communications.

Three Actions to Take Tomorrow

Start now — governance cannot wait for formal adoption

1

Action #1: Identify Where AI Is Already Being Used

Survey employees, check software procurement records, and ask IT to review network traffic for AI services. You cannot govern what you cannot see.

2

Action #2: Provide Clear Employee Guidance

Issue interim guidance on approved tools, prohibited actions, and what to do when uncertain. Employees want direction — they are not trying to cause harm.

3

Action #3: Engage Leadership, IT, HR & Legal

Governance requires cross-functional ownership. Convene a working group to develop formal policy, acceptable use standards, and vendor review criteria.

Questions?

Thank You

Ryan Burns, MLIS
Director of Cyber Services

Mike Bell, CISSP, MLIS
Senior Cybersecurity Advisor

Email: cybersquad@tmlirp.org



**BUILDING CYBERSAFE
COMMUNITIES**



"Invisible adoption is often the greatest AI challenge facing organizations today."