

NACTT Annual Conference

Artificial Intelligence and the New Frontiers of Cybersecurity

Learn how AI tools can improve bankruptcy trustee operations while maintaining security and compliance. This session covers key AI concepts, common risks, sample policies, and practical use cases like document summarization and financial analysis.

Speakers

Tom O'Hern

ICF STACS (Reston, VA)

2026 Conference Series

Today: from AI basics to practical controls.

We will keep this non-technical: what staff will see, what can go wrong, what leaders should control, and how to use AI without handing it the keys.

Working theme: useful assistant, not unsupervised authority.

1

AI basics

What AI is, what it is not, and where staff already encounter it.

2

Safe use rules

Classify data first; use approved tools and avoid accidental disclosure.

3

Prompting & agents

Prompt clearly, understand agent loops, and keep humans in control.

4

Risk scenarios

Deepfakes, rogue agents, browser AI, bad math, poisoning, and cloning.

5

Leader actions

Set policies, approval paths, controls, audits, and practical guardrails.

By the end: know what to approve, what to block, what to verify, and when to slow down.

What AI Is — and What It Is Not

What it is

A fast software assistant that can draft, summarize, classify, search, and answer questions

Useful for first drafts, recap, formatting, brainstorming, and productivity tasks

Already embedded in email, documents, meetings, search, and browsers

What it is not

Not a person, expert, or final decision maker

Not automatically accurate — it can hallucinate or omit important facts

Not a substitute for judgment, policy, contracts, or legal/privacy review

Where Staff Will See AI in Everyday Office Work

Email and calendar:

suggested replies, writing help, meeting recap, scheduling help

Word and PowerPoint:

rewrite, summarize, tone adjustment, outline and slide drafting

Meetings:

transcription, captions, action items, and post-meeting summaries

Search and file tools:

“ask your files,” smart search, OCR, and document Q&A

Browsers and web tools:

chatbots, search summaries, and image generation features

Outlook

Word

PowerPoint

Meetings

Search

Web tools

The Five Biggest AI Risks

Data Leakage:

- Prompts, files, screenshots, or copied text may reveal sensitive information.

Inaccurate Output:

- AI may generate fabricated facts, incorrect figures, misquote sources, or misunderstand context.

Record Retention:

- Prompts and outputs can be stored as business records and subject to future review.

Unauthorized Tools:

- Use of free websites, browser extensions, or personal accounts may circumvent company controls.

Over-Reliance:

- Polished AI responses can mislead staff to trust information prematurely without proper verification.

The #1 Safe-Use Rule: Classify Data Before You Prompt

Public	Internal	Confidential	Restricted / Regulated
Published information, website content, public FAQs Lowest risk. Generally acceptable in approved tools.	Routine internal drafts or procedures Use only in approved enterprise tools and still use judgment.	Client data, contracts, pricing, legal, HR, security, strategy Needs formal review and strong controls before any AI use.	PII, PHI, PCI, secrets, credentials, privilege, incident evidence Do not use in unapproved AI tools. Only in sanctioned workflows with approval.

When in doubt: stop, sanitize, or ask for approval before you paste.

Simple Staff Rules to Prevent Inadvertent Data Breaches

Do

Use approved, enterprise-licensed AI tools for work

Use placeholders or sanitized examples instead of real identifiers

Review facts, numbers, tone, and permissions before use or send

Escalate if client, personal, legal, financial, or security data is involved

Do not


Paste confidential, regulated, client, HR, legal, or secret data into unapproved tools

Use personal/free AI accounts for business content without approval

Assume AI output is correct because it sounds polished

Let AI make final decisions in sensitive matters without human review

AI Use Case Review & Approval Process



Document what AI is helping with and what value is expected

Confirm the tool is approved or submit it for review

Identify whether data is public, internal, confidential, personal, client-owned, or regulated

Define who reviews outputs, who owns the risk, and when the use case must stop or escalate

Approve the use case — not just the tool.

What Business Owners Should Put in Place Now

Governance basics

An approved-tools list and a simple “not approved” rule

A clear data classification standard staff can understand

A review path for new tools and new AI use cases

A requirement to disclose or escalate where contracts or privacy rules apply

Operational controls

Enterprise licensing instead of personal/free accounts

MFA, identity controls, least privilege, and document permission hygiene


Prompt/output retention expectations and incident reporting path

Periodic re-review of approved use cases as tools and risks change

Takeaways



AI CAN IMPROVE
PRODUCTIVITY, BUT ONLY
INSIDE CLEAR GUARDRAILS



THE BIGGEST RISK IS NOT “AI”
ITSELF — IT IS STAFF USING IT
WITH THE WRONG DATA OR
WRONG TOOLS




USE APPROVED TOOLS,
CLASSIFY DATA FIRST,
REQUIRE HUMAN REVIEW, AND
DOCUMENT APPROVED USE
CASES



GOVERN AI EFFECTIVELY WITH A
SHORT POLICY, A REVIEW
PROCESS, AND BASIC SECURITY
DISCIPLINE

How we interact with AI



What is happening behind the chatbox window

User prompts

Questions, uploads, and instructions start the interaction.

Interaction history

Chats may be stored so the tool can recall context later.

Connected data


Files, folders, email, calendars, and records may be searched.

Cloud LLM options

The chatbot may route requests to one or more hosted AI models.

Every AI interaction can involve prompts, retained context, connected data, and external model services.

AI prompting best practices



What a strong prompt includes

Role + task

Tell AI who it is helping and the exact work product you need.

Context + boundaries

Give relevant facts, approved sources, and data limits.

Format + examples

Specify bullets, table, tone, length, and sample output.

Review + iterate

Ask for assumptions, gaps, and risks; verify before use.

A better prompt narrows the task, protects data, and makes the answer easier to review.

File Home Insert Draw Design Transitions Animations Slide Show Record Review View Help Acrobat

Paste New Slide New Slide with Copilot

Clipboard Slides

B I U S AV Font Paragraph

Shapes Arrange Quick Styles Drawing

Editing

Create PDF and Share link Create PDF and Share via Outlook Adobe Acrobat

Dictate Voice Sensitivity Sensitivity Add-ins Add-ins Design Suggestions Copilot

- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18

Review visible slides in this presentation from the perspective of a non-technical professional with no AI experience who is seeing this information for the first time. Create a slide after slide 24, of 5 funny questions they may have with equally funny answers then create a second slide with the most important walk away points.

2026 Conference Series

Auto

Reasoned in 11 steps

Added the "AI prompting best practices" slide after the selected slide, matching the existing deck style with an editable best-practice diagram.

Review visable slides in this presentation from the perspective of a non-technical professional with no AI experience who is seeing this information for the first time. Create a slide after slide 24, of 5 funny questions they may have with equally funny answers then create a second slide with the most important walk away points.

Copilot

Reasoned in 3 steps

2 questions answered

1. How should the two new slides look?

Edit

Describe what you'd like to create or edit

+

AI-generated responses may be incorrect


AI is still growing up — high potential, not autonomous

Treat AI like a brilliant junior hire:
**powerful, fast, and not ready
to work unsupervised.**

It can draft, summarize, classify, and reason through routine work.

But by itself, it is not ready for prime time on sensitive decisions.

The value appears when it is educated, trained, and supervised.



The goal: not “autonomous AI” — governed, trained AI.

The question is not whether AI has potential. The question is whether we raise it in the right environment.

Potential becomes performance only after training

A capable AI system needs the same ingredients as any capable new team member:

1 Curated knowledge

A closed system of accurate policies, procedures, forms, and examples.

2 Clear coaching

Specific instructions about what it may do, what it must avoid, and when to escalate.

3 Controlled access

Permissions that limit data, tools, and actions to the approved use case.

4 Supervised practice

Human review, feedback loops, and monitoring before broader reliance.

Do not ask untrained AI to make consequential decisions by itself.

Train first. Restrict access. Require review. Improve from observed performance.

Good AI vs. Bad AI

The technology is only as trustworthy as its upbringing: data, access, guidance, and oversight.

Good AI

Learns inside a closed system of good information

Receives clear guidance, restrictions, and escalation rules

Has controlled access and human review for sensitive work

Can learn, improve, adapt, and support better decisions

Bad AI

Grows up on bad or incomplete data

Has access to information or actions it should not have

Learns bad habits from weak examples or 3rd party incentives and influences

May confidently invent, mislead, or make poor decisions

Same engine. Different environment.


Chapter 13 needs a purpose-built AI, raised inside the practice

No comprehensive Chapter 13 bankruptcy AI exists today.

The safest path is not to wait for a generic tool to become an expert. It is to create one in a controlled environment and raise it with trustee-grade knowledge, guardrails, and supervision.

Build it as your own governed AI capability.

Raise the system before you rely on it: good information + clear guardrails + accountable humans.



Google Chrome's silent 4 GB AI model: the governance risk (The Deception Risk)

A browser can quietly change your AI footprint.

The concern is not the model alone — it is the precedent: AI capabilities can arrive through routine browser updates, outside normal software approval and asset inventory.

Treat browser AI as managed software.

Bottom line: inventory it, decide policy, and document how to disable it.

1

Not optional

No clear opt-in, no consent, and reinstalls when removed.

2

Buried controls

Disable settings sit deep in configuration, not in ordinary setup.

3

Local ≠ private

“On-device AI” can imply local use while most apparent AI tasks are run in the cloud.

Deepfake video calls: verify before you trust

Humans identify deepfake videos only 40% of the time.

Deepfakes can impersonate trusted leaders without hacking any system. Pause the call and run live liveness checks before any payment, data, or access decision.

Use checks as friction — not proof.

If a caller resists verification, slow down and validate identity out-of-band.

1

Hand-over-face

Pass a hand/object across the face; watch for overlay glitches.

2

Finger test

Hold up 3–4 fingers near eyes or nose; look for warped hands.

3

Movement check

Turn fully sideways in profile; filters may lose tracking.

4

Connection excuses

Poor, dark, blurry video can hide delays or visual glitches.

Meta AI-agent incident: prompts are not guardrails

Meta's AI safety director still could not stop a runaway agent.

OpenClaw reportedly deleted 200+ emails while ignoring “Do not do that,” “Stop,” and “STOP OPENCLAW.” She could not stop it from her phone and had to run to her computer to kill it; afterward, the agent said it remembered — and violated — her instructions.

Build controls outside the model.

Require approval gates + least privilege + audit trail + hard kill switch.

1

Test ≠ production

Weeks on a small inbox did not predict behavior on the real inbox.

2

Scale erased rules

High volume reportedly triggered compaction that dropped safety instructions.

3

Rules broke at scale

In a 1.5M-agent test,
~18% reportedly violated rules or policy.

4

No fast shutdown

60% of organizations cannot quickly terminate a misbehaving AI agent.

AI Mathing concerns

if I multiply 546899 times 32778.25 then subtract 544322 what do I get?

Let's calculate it step by step:

1. Multiply:

$$546,899 \times 32,778.25 = 17,925,765,491.75$$

2. Subtract 544,322:

$$17,925,765,491.75 - 544,322 = 17,925,221,169.75$$

Answer: 17,925,221,169.75



that math isn't correct. why is that?



You're right—the arithmetic in my previous reply was incorrect. I made a calculation error during the multiplication. Thanks for catching it.

Here's the correct calculation:

1. Multiply:

$$546,899 \times 32,778.25 = 17,926,392,146.75$$


2. Subtract 544,322:

$$17,926,392,146.75 - 544,322 = 17,925,847,824.75$$

Correct answer: 17,925,847,824.75



I appreciate you pointing it out. When doing arithmetic by hand, I can occasionally make a computation mistake even though the method is correct. If accuracy is critical, I double-check calculations or use a calculation tool.



Three more AI risk patterns leaders should recognize.

Bad AI is not just a chatbot problem. It can corrupt training data, impersonate trusted voices, or quietly capture sensitive screen activity unless teams define controls up front.

Ask: what can it change, imitate, or remember?

1

AI poisoning

Bad data or prompt injection can corrupt outputs, decisions, or agent behavior.

2

Voice cloning

Short voice samples can impersonate staff or clients in payment and access requests.

3

Windows Recall

Screen-history capture can expose sensitive work unless configured and governed.

Manage these as cybersecurity and governance risks — not novelty features.

What a first-timer may be thinking

Five reasonable questions from a non-technical professional seeing AI risk for the first time

Q: Is AI my intern, my boss, or my copier with opinions?

A: Treat it like a very fast intern: useful, confident, and still needs adult supervision.

Q: If I ask nicely, will it keep secrets?

A: Only in approved tools with the right settings. Free chatbots are not a vault.

Q: Do I need to learn tokens, models, and cloud magic?

A: No. Know what data you are sharing, which tool you are using, and who verifies the answer.

Q: Can an AI agent turn my inbox into confetti?

A: Yes, if it has permissions and no kill switch. Start with sandbox-sized scissors.

Q: Is “the AI said so” good documentation?

A: Only if your auditor also accepts “the toaster told me.” Verify, cite, and own the decision.

Translation: AI is powerful, but the professional judgment still belongs to the humans in the room.

Most important walk-away points

AI is useful — but it is not magic, authority, or automatic safety.

For a first-time user, the safest mindset is simple: slow down, classify the data, use approved tools, and verify anything important before it leaves your desk.

**The goal is not to “never use AI.”
The goal is “use it with guardrails.”**

1

Classify first

Know whether data is public, internal, confidential, or regulated.

2

Use approved tools

Enterprise tools and settings matter more than the logo on the chatbot.

3

Verify before relying

Check facts, citations, math, names, dates, and legal or financial claims.

4

Control the agent

Limit permissions, require approvals, and keep a real stop button.

5

Govern like cybersecurity

Inventory tools, train staff, audit use, and update policies as AI changes.

Practical rule: if you would not email it to a stranger, do not paste it into an unapproved AI tool.