AI AS WEAPON AND SHIELD | CYBER NUTMEG

AI Threat Readiness Checklist

For Connecticut towns, school districts, and state agencies

Isidore "Izzy" Sobkowski
Executive Director, AskRAI
Executive Director, The Institute for Responsible AI
Senior Advisor for AI and Advanced Technology, Connecticut Education Network
For Discussion Purposes Only.
Do not rely upon the information contained herein. Always consult an appropriate professional.

How to use this checklist

AI has made familiar attacks (phishing, impersonation, payment fraud) cheaper, faster, and far more convincing. Most of the protection is process, not products, and most of it is free. Work through each section with the people named, tick what is already in place, and assign an owner and date to every gap.

Who should complete it: IT lead, finance or business manager, HR, communications, and an executive sponsor (superintendent, first selectman, town manager, or agency head). One meeting of about 90 minutes is enough for a first pass.

If you are starting from nothing: do Section 1 this week. It needs no budget and would have stopped the largest documented deepfake fraud to date.

If you already have controls: use the scorecard at the end to find gaps, then retest annually.

1. Stop the money (finance and procurement)

- ☐ **Callback rule.** Any change to vendor banking details or employee direct deposit is verified by phone using a number already on file, never one supplied in the request.
- ☐ **Two-person approval** for wires, ACH payments, and new payees. The second approver verifies independently rather than simply approving.
- ☐ **Verbal code phrase** known only to executives and finance staff, used to confirm urgent or unusual requests. Changed periodically.
- ☐ **No channel switching.** Official requests stay on official channels. Any request to move to personal text, WhatsApp, Signal, or Telegram is treated as suspicious.
- ☐ **Hold period** for the first payment to a new or changed account, where operations allow.
- ☐ These rules are **written into finance procedures**, not only into training.

© Copyright 2026, All Rights Reserved. The Institute for Responsible AI | www.theinstituteforresponsibleai.com

2. Protect identity and access (IT)

- ☐ **Phishing-resistant MFA** (security keys or passkeys) for administrators, finance staff, and executives. Text-message codes can be relayed in real time by modern phishing kits.
- ☐ **MFA on all email and remote access** for everyone else.
- ☐ **Email authentication:** SPF and DKIM configured, and DMARC at enforcement (quarantine or reject) so attackers cannot send mail that appears to come from your domain.
- ☐ **Separate admin accounts** from everyday accounts.
- ☐ **Alerts** for impossible travel, new mailbox forwarding rules, unusual login times, and mass downloads.

3. Hiring and onboarding (HR with IT)

- ☐ **Live video interview** with camera on, matched to a government-issued ID.
- ☐ **References verified** using contact information you find independently, not what the candidate provides.
- ☐ **Equipment shipped only to a verified address**; watch for requests to reship.
- ☐ **New remote accounts monitored closely** for their first 30 days, including remote IT contractors.

4. Leaders and reputation (deepfake response)

- ☐ **A named person** who can quickly confirm or deny whether audio or video of a leader is authentic.
- ☐ **Pre-written holding statement** for a suspected deepfake, approved in advance.
- ☐ **Contact list ready:** counsel, communications, local police, and platform reporting channels.
- ☐ **Schools:** student and staff deepfakes of one another are covered in conduct policy, with support for those targeted.
- ☐ **One tabletop exercise a year** using a deepfake scenario.

5. People and training

- ☐ **Training updated for AI.** The message: verify the request, not the messenger. New warning signs are urgency, secrecy, a channel switch, or a payment change, not typos or an odd voice.
- ☐ **Leaders state publicly** that no one will ever be disciplined for verifying a request from them.
- ☐ **One easy way to report** a suspicious message or call (a button, an email alias, or a help desk category).
- ☐ **Simulations include voice and text**, not just email.

© Copyright 2026, All Rights Reserved. The Institute for Responsible AI | www.theinstituteforresponsibleai.com

6. AI-assisted defense

- ☐ **Inventory the AI security features you already license** in email, identity, and endpoint tools; enable and tune them before buying anything new.
- ☐ **Ask five questions before buying** any AI security tool: What will it catch that we miss today? What is its false-positive rate in organizations like ours? What data leaves our network? Who acts on its alerts at 2 a.m.? Can we trial it on our own data first?
- ☐ **A named owner for alerts**, including after hours, or a managed service that covers them.
- ☐ **Quarterly review** of security settings, because vendors change defaults.
- ☐ **Check free services** available to public entities (for example, from CISA) and peer threat-sharing through CEN and regional councils of governments.

© Copyright 2026, All Rights Reserved. The Institute for Responsible AI | www.theinstituteforresponsibleai.com

7. If money moves: response card

Speed matters. In 2025 the FBI IC3 Recovery Asset Team reported freezing about $679 million in fraudulent transfers, and the chance of recovery drops sharply after the first day or two. Fill in the numbers below and post this page in your finance office.

Step	Action	Contact and number
1	Call your bank's fraud line and request a recall of the transfer	
2	File a complaint at ic3.gov with transaction details	
3	Call the FBI New Haven Field Office	
4	Notify your cyber insurer (check notice deadlines)	
5	Preserve evidence: emails, voicemails, call logs, transaction records	
6	Notify leadership, counsel, and IT	

8. Scorecard

Count the boxes ticked in each section.

Section	Ticked	Out of	Biggest gap	Owner	Target date
1. Stop the money		6			
2. Identity and access		5			
3. Hiring and onboarding		4			
4. Deepfake response		5			
5. People and training		4			
6. AI-assisted defense		5			
Total		**29**			

Priority order if you cannot do everything: Section 1 first (free, and it stops the costliest attacks), then phishing-resistant MFA and DMARC from Section 2, then the deepfake response plan in Section 4.

© Copyright 2026, All Rights Reserved. The Institute for Responsible AI | www.theinstituteforresponsibleai.com

Real cases behind this checklist

- **Deepfake video call fraud:** Arup's Hong Kong office lost about US$25.6 million in January 2024 after a video call in which every other participant was a deepfake (Hong Kong Police; Arup statement, May 2024).
- **AI voice impersonation of officials:** FBI public service announcements, May and December 2025.
- **Deepfake of a school principal:** Pikesville High School, Maryland, 2024; the creator was sentenced in April 2025.
- **Fake remote hire:** KnowBe4 disclosure, July 2024.
- **AI phishing effectiveness:** Heiding, Schneier, et al., human-subjects study; 54% click rate for fully AI-written spear phishing versus 12% for generic phishing.
- **AI in malware and intrusions:** Google Threat Intelligence Group and Anthropic reports, November 2025.
- **Losses:** FBI IC3 2025 Internet Crime Report (business email compromise about $3.05 billion).

This checklist is general guidance, not legal advice. Confirm current service offerings and legal obligations for your organization.

© Copyright 2026, All Rights Reserved. The Institute for Responsible AI | www.theinstituteforresponsibleai.com