

Brunch, Business & AI Breaches

AI Is the New Cybersecurity Risk



YOUR HOST

Meet Ikram Massabini.

25 years building secure technology businesses in Western New York.

Applied AI Standards Alliance



CEO
MVP Network Consulting

25 YEARS
of IT experience

EXPERT
IT security consultant

HIPAA
consultant & auditor

AI COACH
for business leaders

DEGREE
BS, Electronic Engineering

GET IN TOUCH

ikram@mvpworks.com

716.630.1701 • mvpworks.com

1485 Niagara Street, Buffalo, NY 14213

MVP Network Consulting
Technology That Works.

THE QUESTION EVERY OWNER MUST ANSWER

Are you using AI intentionally?

Or is AI happening TO your business — through whatever tool your employees happen to be using on personal accounts today?

INTENTIONAL

- Approved tools
- Written policy
- Trained employees
- Full visibility
- ROI tracked

ACCIDENTAL

- Shadow AI
- Personal accounts
- Data exposure
- No insight
- No accountability

TODAY We move your business from the right column to the left — in six months.

WHY WE'RE HERE

AI is already inside your business.

Whether you approved it or not.

78%

of knowledge workers are using AI on the job — and most haven't told their employer which tools.

Free tools

ChatGPT, Gemini, Claude — on personal accounts, in browser tabs you can't see.

Real data

Customer info. Financials. Contracts. Strategy docs. Pasted into prompts.

Zero visibility

No policy. No log. No way to know what left the building today.

THE RISK If you don't see what AI tools your team uses, you don't know where your data lives.

WHAT IT COSTS YOU

Four ways unsanctioned AI hurts your business.

01

Loss of IP

Your proprietary know-how walks out the door — into someone else's model.

02

Confidential data exposed

Customer info, financials, and contracts pasted into prompts you can't recall.

03

Inconsistent outputs

Different tools, different rules. No guardrails. No audit trail.

04

Data leaves with employees

Personal AI accounts go with them — including everything they fed in.

BOTTOM LINE Unsecured AI usually means your data is the product.

MEANING

What is AI, really?

It's not as "intelligent" as the marketing makes it sound.

DEFINITION

AI is the simulation of human intelligence by machines — the ability to learn from data, reason within rules, and improve over time through feedback.

It predicts patterns. It doesn't "understand" — it guesses the most likely next word.

It learned from us

Every model is trained on human text, images, and decisions.

It needs a human

AI works best with you in the loop — reviewing, correcting, deciding.

It remembers what you feed it

Every prompt you give a free tool can become part of the next model. Your IP can become someone else's answer.

REMEMBER Better inputs in, better outputs out. Garbage in, garbage out still applies.

The Cyber Risks You Don't See Coming

01 — Shadow AI

Employees using unsanctioned AI tools outside of IT visibility or governance.

02 — Overshared Data

Sensitive business information entered into into AI tools without policy or guardrails. guardrails.

03 — Excessive Permissions

AI can surface files and data the user was was technically allowed to access — but but never should have been.

04 — AI-Powered Impersonation Impersonation

Attacks that mimic real people — voices, voices, faces, and writing styles — with alarming accuracy.

05 — Uncontrolled Integrations

Third-party AI tools connected to your core platforms create data flows nobody approved.

The biggest AI security problem may not be the AI itself. It may be what the AI can access.

Risk #1 — Shadow AI

UNGOVERNED AI USE

Your employees are probably already using AI tools today — and that's not a criticism. It's a reality that requires a response.



Uploading to Summarize

An employee uploads a financial spreadsheet to ChatGPT to get a get a quick summary before a board call.



Pasting Email Threads

A manager pastes a sensitive contract negotiation thread into an AI tool to draft a response.



AI Meeting Transcription

A team uses an AI transcription tool for a confidential strategy session strategy session without realizing the audio is stored externally. externally.



Browser Extensions

An employee installs an AI writing extension that reads every webpage and document they open.

> Helpful? Absolutely. Governed? Maybe not.

• SHADOW AI

Employees put company secrets into the chatbot

Samsung Electronics • March–April 2023 landmark case

DATA LEAKAGE

3 CASES

REPORTED WITHIN WEEKS



WHAT HAPPENED

Employees reportedly entered source code and internal meeting material into ChatGPT while trying to debug and summarize work.

THE EXISTING CRACK

No approved workflow, classification boundary or clear rule prevented confidential data from being pasted into a public AI service.

AI AMPLIFIED

Convenience + data reuse risk

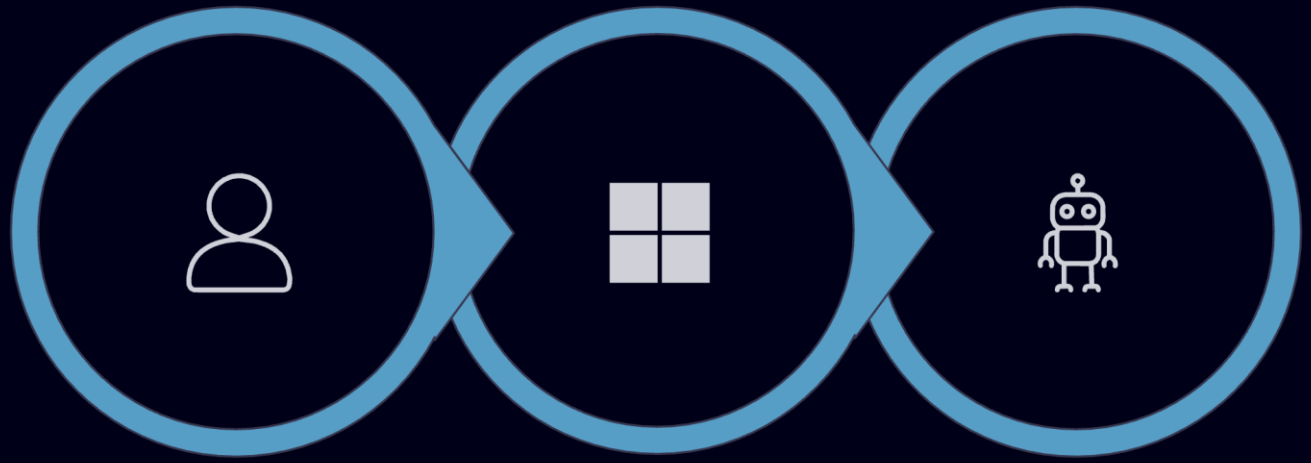
CONTROL THAT BREAKS THE CHAIN

Approved enterprise tools, policy, DLP and training

Sources: Bloomberg (2 May 2023); TechCrunch (2 May 2023), citing Samsung's internal response.

Risk #2 — AI Can See What the User Can See

THE PERMISSIONS EXPOSURE PROBLEM



Yesterday's Permission Problem Is Problem Is Today's AI Problem

Before AI

A forgotten file sits three folders deep in an in an old SharePoint site. Almost no one one ever finds it. The overpermissioning is overpermissioning is harmless in practice practice — invisible, dormant, forgotten. forgotten.

With AI

An employee asks Copilot: *"What do we pay our executives?"*
AI searches everything that employee has permission to access — instantly, thoroughly, and without judgment.

> AI makes information easier to find. That includes information people shouldn't have shouldn't have been able to find in the first place.



Risk #3 — The Phishing Email Got an Upgrade

AI-ENHANCED SOCIAL ENGINEERING

Old Phishing

- Bad grammar and misspellings
- Generic, impersonal language
- Obvious suspicious links
- Poor branding and formatting
- Implausible scenarios

AI-Enabled Attack

- Flawless grammar and tone
- Company-specific context and detail
- Executive names and real relationships
- Natural conversational flow
- Personalized, time-pressured requests

> The clues we've trained employees to look for are disappearing. The old detection rules no longer apply.

The Attacker Doesn't Need to Hack You First

OPEN-SOURCE INTELLIGENCE + AI

Publicly Available Information

- LinkedIn profiles and connections
- Company website and About page
- Press releases and announcements
- Social media posts and events
- Employee bios and org charts
- Job postings (revealing tech stack)
- Vendor and partner announcements

An attacker feeds this into AI and generates a highly personalized attack in minutes.

The Message That Arrives

> "I'm heading into the Buffalo Club meeting. Can you send me the updated banking information before 10? I'm losing cell service."

Fictional example — but the names, the location, the urgency, the timing, and the ask are all drawn from information that was already public.



Risk #4 — When Seeing and Hearing Isn't Believing

VOICE CLONING · DEEPFAKE VIDEO · EXECUTIVE IMPERSONATION

This is no longer theoretical. The technology to clone a voice or generate a convincing video of someone of someone is widely available, inexpensive, and improving rapidly.

1

CFO Receives a Teams Call

The request comes through a familiar, trusted channel.

2

CEO Appears on Video

The face, the voice, the mannerisms — everything looks and sounds right.

3

Urgent Financial Action Requested Requested

Wire transfer. Account change. Credential reset. The request feels normal and time-sensitive.

4

What Happens Next?

Without a verification process, the answer depends entirely on what the CFO decides in that moment.




SLASH
FREE GIFT
862
377

• DEEPFAKE VIDEO FRAUD

The video call where everyone was fake

Hong Kong finance employee • February 2024

DEEPFAKE EXECUTIVE FRAUD

HK $200M

15 TRANSFERS AUTHORIZED



WHAT HAPPENED

A finance worker joined a group video call with people who looked and sounded like the CFO and colleagues—then sent the money.

THE EXISTING CRACK

A high-value payment could proceed based on apparent executive presence, without a separate trusted-channel check.

AI AMPLIFIED

Visual + voice authority

CONTROL THAT BREAKS THE CHAIN

Dual approval + callback to a known number

Sources: Hong Kong Police Force CyberDefender (4 Feb 2024); Arup confirmation reported by Financial Times (17 May 2024).

Technology Can't Be Your Only Defense

Your process becomes a security control.



High-Risk Request

Exit Original Channel

Call Known Contact

Second-Person Person Approval

Approve & Execute

- Wire Transfers
- Banking Changes
- Payroll Changes
- Credential Resets
- Vendor Payment Changes

Risk #5 — AI Agents Change the Equation

FROM ANSWERING TO ACTING

AI Assistant

Responds to a question. Surfaces information. Drafts content. **A human human still takes every action.**

AI Agent

Can take actions on behalf of a user — autonomously, across systems, at scale. **The human may not be in the loop for each step.**

What an Agent Might Do

- Read and respond to email
- Create and share documents
- Contact a customer or vendor
- Update CRM records
- Access and modify files
- Trigger automated workflows

> The risk profile changes significantly when AI moves from answering questions to taking action. Start with least-privilege permissions and expand deliberately. expand deliberately.

AI isn't the problem.

Uncontrolled AI is.

The goal isn't to stop your organization from using AI. The goal is to understand **who is using who is using it, what it can access, where your data is going, what actions it can take, and take, and how you would respond if something went wrong.**

That's not restriction. That's leadership.



FIX THE FOUNDATIONS FIRST AI sits on top of your security — not next to it.

ATTACK SURFACE

AI doesn't create most security problems. It amplifies them.

Every weakness you have today gets bigger when an agent has access to it.

Weak passwords

Faster to crack. Easier to test at scale.

Missing MFA

AI-driven fatigue floods until someone clicks.

Excess permissions

Agents inherit overprivileged accounts.

No monitoring

AI can act for hours before anyone notices.

Stale offboarding

Departed staff's AI access remains active.

Unpatched systems

Attackers automate exploit discovery.

FIX THE FOUNDATIONS FIRST AI sits on top of your security — not next to it.

MY AI PREDICTION

Every business will have many AI agents — just like every business has a website.

Customer service

AI agents handling first-touch, 24/7.

Sales

Agents that research prospects and draft outreach overnight.

Operations

Agents triaging tickets, routing invoices, owning onboarding.

The winners won't be the businesses that adopted first — they'll be the ones that adopted right.

THE PATH FORWARD



So… what should businesses actually do?

Everything you just heard is why the roadmap starts with KNOW and SECURE — you automate on a foundation, or you automate your problems.

Six phases. Real outcomes at every step. And a trusted partner walking with you the whole way.

THE 6-STEPS ROADMAP

Six Steps. Six phases. From shadow AI to competitive edge.

BY Step 6: 5–10 hrs reclaimed per employee / week • 60–80% faster client response • onboarding in half the time

M1 KNOW

AI Readiness Assessment. Inventory tools, exposed data, permissions.

M2 SECURE

Build the foundation. Approved tools, locked-down data, signed AUP.

M3 TRAIN

Launch with your data. Team certified on RISEN. First retrieval agent.

M4 CONNECT

Integrate approved line-of-business applications with AI.

M5 AUTOMATE

Task agents replace repetitive work — with human review.

M6 GOVERN & SCALE

Monitor, audit, refine. AI is how your business runs.

HOW IT'S BUILT Step 1–3 are your 90-day quick start. Step 4–6 scale and govern what works.

Step 1 • KNOW

Before anything else: the AI Readiness Assessment.

It's easy to buy an AI license and switch it on — without knowing what data it can reach. Find out where you stand first.

01 Software adoption

Are employees using what you already pay for? Active users, license fit, feature patterns.

02 Security

Identity, access, and threat posture. MFA, access policies, endpoint protection.

03 Technical readiness

Is your environment configured for AI? Licensing, configuration, network readiness.

04 Data governance

Where is your data, and who can see it? Labels, sharing, permissions, DLP, retention.

BY END OF STEP 1 You know every AI tool in use, what data is exposed, and exactly where to start.

Step 2 • SECURE

AI is hungry for data. Open the right door.

"Just give AI everything" is tempting — and dangerous. The safe rule: AI sees only what each person is allowed to see.

1

Sensitivity labels

Classify data so AI knows what it can use — and what to leave alone.

2

Least privilege

AI inherits each user's access. Clean up over-shared sites first.

3

DLP policies

Block sensitive content from leaving — or entering AI prompts.

4

Audit logs

Every prompt, every action — searchable when you need it.

5

Conditional access

Tie AI use to device, location, and identity posture.

BY END OF STEP 2 Approved tools chosen, data access scoped and logged, and your AI use policy signed by every employee.

Step 3 • TRAIN

Launch AI with your data — and train everyone.

Five pillars turn approved AI into daily habit.

1

Data

Connect AI chat to your files, email, and collaboration tools — where permissions allow.

2

Tools

Give AI access to the applications your team already works in.

3

Policy

Deliver clear, plain-language rules for responsible use and compliance.

4

Training

AI basics plus the RISEN framework for better prompting — every employee certified.

5

Monitor

Watch onboarding, adoption, and early wins. Coach where it's slow.

BY END OF STEP 3 Every employee is certified and using approved AI daily — and your first retrieval agent is live.

HOW YOUR TEAM TALKS TO AI

The RISEN framework.

Five inputs that turn AI from a curiosity into a daily tool. Every employee learns this on Day 1.

R

Role

Tell AI who to be.

"Act as a CFO reviewing a proposal."

I

Instructions

Say what you need done.

"Summarize the top three financial risks."

S

Steps

Show how to get there.

"First skim, then list risks, then rank."

E

End goal

Define what success is.

"A 3-bullet summary I can paste in Slack."

N

Narrowing

Set constraints.

"Plain English. Under 100 words. No jargon."

WHY IT MATTERS Better prompts in, better work out. Prompting is the new business skill.

WHAT IT PROS ACTUALLY WANT TO SEE

ChatGPT vs. Claude vs. Copilot.

IT leaders don't compare chatbots on vibes. They compare control, data boundaries, and integration. Here's the short version.

ChatGPT

OpenAI

Fastest feature cycle, strong reasoning + agents.

IT wants: admin console, SSO, no-training-on-data in Enterprise.

Copilot

Microsoft 365

Lives inside Teams, Outlook, Word, Excel.

IT wants: it inherits your existing permissions — good and risky.

Claude

Anthropic

Long documents, careful answers, safety focus.

IT wants: strong policy controls, low hallucination on long files.

Gemini

Google

Deep Workspace and Drive integration.

IT wants: unified Google admin and DLP controls.

Open models

Llama, Mistral

Run in your own cloud or on-prem.

IT wants: full data residency — but you own the security.

THE REAL ANSWER The model matters less than the guardrails. Pick the one that fits your stack — then govern it.

Step 4 • CONNECT

Extend AI into the apps that run your business.

Beyond the productivity suite: HubSpot, Sage, Epic, Infor — the systems your work actually lives in.

01

Identify

Pick the line-of-business apps where AI would save the most time.

02

Confirm

Check each against its API, vendor support, licensing, and security requirements.

03

Clean & tag

Get the data in those systems clean and labeled before AI touches it.

04

Build & test

Develop, validate, and document each integration before rollout.

BY END OF STEP 4 AI works inside your core business systems — not just in a chat window.

Step 5 • AUTOMATE

Agents and workflows that do the work.

Trusted answers and automated processes — using only approved company data.

01

Evaluate

Pick high-value, repetitive workflows with the people who own them.

02

Build & pilot

Build, refine, and roll out agents using proven patterns and your integrations.

03

Map triggers

Connect workflow triggers so agents act when events happen — not on demand.

04

Measure

Validate outcomes with usage data before scaling each automation.

BY END OF MONTH 5 Your first task agents run real work — invoices, tickets, follow-ups — with human review.

APPLY IT TO REAL WORK

Take any process and ask: "Can I AI that?"

Example: a weekly LinkedIn post — 8 steps, 40 minutes today.

TODAY → 40 MIN

Research	Compile	Edit copy	Review	Edit again	Design	Save	Post
10 min	5 min	5 min	Human	3 min	20 min	1 min	1 min

WITH AI → 5 MIN (–87%)

Research	Compile	Edit copy	Review	Edit again	Design	Save	Post
AI	AI	AI	Human	AI	AI image	AI	AI

5 of 8 steps offloaded to AI. ~35 minutes saved per post — every week. Now multiply across every process you have.

START HERE Start with existing processes — not new ones. You already know what good output looks like.

HOW WE EVALUATE EVERY AI USE CASE

The BXT framework.

Business + eXperience + Technology. Score every AI idea 1–5 on each.

B Business

Does this drive a real outcome?

- Quantifiable time or revenue impact
- Aligned to a clear business priority
- An owner accountable for results

X eXperience

Is it usable by real humans?

- Easy to trigger and review
- Output your team trusts
- No added friction to the day

T Technology

Will it work safely and reliably?

- Data is available and clean
- Permissions are appropriate
- Monitorable and reversible

THE RULE If anything scores below 3, fix it before you build.

STEP 6 • GOVERN & SCALE

What good AI governance looks like.

Eight pieces. Missing any one? You're not governing — you're hoping.

1 Executive owner

Someone above the team accountable for AI.

2 Approved tools

One short list everyone knows.

3 Written policy

Plain-language AUP everyone signs.

4 Security review

Every tool vetted before it's allowed in.

5 Training

Every employee certified before access.

6 Access controls

Permissions tied to role, reviewed quarterly.

7 Monitoring

Every prompt and agent action logged.

8 Regular review

Quarterly audits, annual policy refresh.

BY END OF STEP 6 AI is governed, measured, and compounding — and you're finding the next opportunities.

THE STAKES

There's too much to lose to get this wrong.

Customers

Trust is hard-won, easily lost.

Your data

Your IP — out of your hands.

Reputation

One leak makes the news.

Compliance

HIPAA, PCI, state privacy laws.

Cyber insurance

No governance → no coverage.

Your business

Years of work in the balance.

GOOD NEWS Every one of these risks is manageable — with governance in place before scale.

HERE'S WHAT YOU CAN DO

STEP 1 starts here. Book your AI Readiness Assessment.

WE WILL NOT

- ✕ Install anything on your network
- ✕ Touch your sensitive data
- ✕ Require a complicated setup
- ✕ Disrupt your business

WE WILL

- ✓ Review your current AI exposure
- ✓ Identify security gaps and risky patterns
- ✓ Prioritize your most impactful next steps
- ✓ Walk through practical AI use cases for you
- ✓ Give you a clear, written path forward

ikram@mvpworks.com • 716.630.1701 • mvpworks.com

ONE LAST THOUGHT

AI is already inside your business.

You came in with AI happening TO your business. You leave knowing how to make it intentional.

QUESTIONS? CONTACT IKRAM.

ikram@mvpworks.com • 716.630.1701 • mvpworks.com



MVP Network Consulting
Technology That Works.

Is Your Business AI-Ready?

Before you adopt more AI, know where you stand.

AI is already in the workplace. MVP helps businesses move forward with AI safely, strategically and with the right guardrails in place.

- Identify AI & security gaps
- Create smarter policies & guardrails
- Find practical AI opportunities



SCAN ME

SCHEDULE YOUR AI & CYBERSECURITY DISCOVERY CALL

mvpworks.com/ai-readiness

For Western New York Businesses