CYBERNOMICS RESEARCH REPORT

Your Employees Are Already Using AI

A practical SMB leader's guide to human AI risk, data exposure, agent boundaries, and safe adoption.

Based on a Cybernomics episode featuring Dr. Nikki Robinson and Jennifer Baca.

EXECUTIVE SUMMARY

Your AI risk is probably already in the workflow

AI is no longer a future project for small and medium businesses. Employees are already using it to write emails, summarize meetings, build spreadsheets, answer customers, create reports, write code, analyze tickets, and speed up admin work.

The risk is not that people are bad at AI. The risk is that normal people are using powerful tools inside unclear workflows. That is where sensitive data, bad answers, unapproved tools, and over-permissioned agents create business exposure.

Bottom line: do not start with a giant AI policy. Start with the five places where people touch AI most often - data, prompts, outputs, approvals, and mistakes.

88%	97%	63%
Organizations using AI in at least one business function, according to McKinsey 2025.	Organizations with AI-related security incidents that lacked proper AI access controls, according to IBM 2025.	Breached organizations that lacked or were still developing AI governance policies, according to IBM 2025.

Who should read this

CEO / Owner

Understand where AI creates business risk before customers, employees, or regulators discover it first.

COO

Control AI inside daily workflows without turning every task into a committee meeting.

CFO

Reduce AI tool sprawl, wasted spend, surprise renewals, and unmanaged vendor risk.

IT / Security Leader

Give employees a safe path without becoming the AI police.

CYBERNOMICS FRAMEWORK

The Five Human Touchpoints of AI Risk

Most SMB AI risk shows up in five simple places. These are the points where employees, managers, agents, vendors, data, and decisions collide.

If you control these five touchpoints, AI becomes easier to use safely without slowing the business down.

1. Data

What information people paste, upload, summarize, attach, or expose to AI tools.

2. Prompts

What people ask AI to do, including whether the request gives away confidential context.

3. Outputs

What AI gives back, including confident but wrong answers, hidden assumptions, or old information.

4. Approvals

What humans allow AI output or agents to affect: customers, money, legal matters, HR, security, or records.

5. Mistakes

What happens when an employee sees something wrong, strange, risky, or embarrassing. This is where no-blame reporting matters.

Turn the framework into action.
Use the Cybernomics AI Human Risk Score, then schedule an AI Readiness Review at Cybernomics.io to decide what to fix first.
cybernomics.io/

FROM THE CYBERNOMICS EPISODE

Do not blame the human. Design the work.

The Cybernomics conversation with Dr. Nikki Robinson and Jennifer Baca challenged a common security idea: "humans are the problem." That phrase is too simple. It also leads to the wrong fix.

The usual fix is more training. Training matters, but it is not enough. If the tool is confusing, the default setting is risky, or the employee is measured only on speed, training will not save the business.

Human factors means designing tools and work around how people actually behave. People get tired. People rush. People miss details. People avoid speaking up if they think they will be blamed. That does not make them bad employees. It makes them human.

Jennifer Baca raised the executive concern many leaders already have: if AI makes a mistake and the company loses money, who is responsible? That question is the right starting point. AI risk is not just a user problem. It is a leadership problem.

Leadership lesson: if AI touches real business work, someone must own tool approval, data rules, review points, agent permissions, and incident reporting.

The car safety analogy

We do not rely on drivers to remember every safety rule every second. Cars have seat belts, brake lights, mirrors, alerts, lane warnings, and airbags. AI needs the same thinking. Do not just tell employees to "be careful." Build work so careful is the default.

BUSINESS IMPACT

What AI risk looks like in a real SMB

SMB leaders do not need fear. They need concrete examples. These are the kinds of AI risks that often appear before there is a formal AI program.

Scenario	What could happen	Business impact	First control
Sales pastes a contract into a public AI tool	Customer terms, pricing, or private negotiation details may be exposed.	Lost trust, legal exposure, vendor risk.	Approved tool list and do-not-paste rule.
HR uses AI to summarize employee issues	Sensitive employee information may be processed, stored, or reused without review.	Privacy risk, employee relations risk, bad documentation.	HR review gate and approved data handling.
Manager sends AI-written customer response	The answer may sound polished but contain false facts, wrong promises, or outdated guidance.	Client frustration, refunds, contract disputes.	Human review for customer-facing output.
Agent can update records without approval	A helpful automation may delete, change, send, or approve something it should only draft.	Operational disruption, financial loss, audit trail gaps.	Read, draft, approve, execute permissions.
Every department buys a different AI tool	Different contracts, retention rules, and data practices spread across the business.	Waste, confusion, compliance gaps.	AI tool intake and renewal review.

The pattern is simple: AI risk is usually not one big dramatic event. It is a chain of small unclear choices that no one owns until something goes wrong.

PRIORITY RISKS

The seven AI risks SMBs should watch first

These are the risks most likely to show up before a company has a formal AI program. They are simple, common, and expensive when ignored.

Risk	What it looks like	Simple control
1. Shadow AI	Employees use personal AI accounts because approved tools are missing, unclear, or harder to use.	Publish an approved AI list. Give people a safe tool first.
2. Data leakage	Customer data, employee data, contracts, passwords, code, or financial details are pasted into the wrong tool.	Create a do-not-paste list. Turn on data loss controls where available.
3. Bad answers	The AI sounds confident but is wrong, outdated, or made up.	Require human review for high-impact work. Ask for sources when facts matter.
4. AI overconfidence	People stop thinking because the AI answer looks polished.	Teach employees that polished does not mean true.
5. Agent overreach	An AI agent can read, write, send, delete, approve, or change records without clear permission.	Separate read, draft, approve, and execute rights.
6. Tool sprawl	Every department buys a different AI tool with different contracts and data rules.	Create one intake process for new AI tools and renewals.
7. Silent failures	People do not report mistakes because they fear blame or embarrassment.	Create a no-blame AI incident channel and fix issues fast.

Why it matters: AI adoption is spreading faster than security and governance. SMBs should close the visibility gap before AI becomes invisible infrastructure.

PRACTICAL CONTROLS

Controls that work without slowing the business down

SMBs do not need a giant AI bureaucracy. They need a short list of rules that are easy to understand, easy to follow, and easy to enforce.

1. Start with tools you already pay for

Before buying ten new AI tools, check the AI features inside Microsoft 365, Google Workspace, Slack, CRM, help desk, and security platforms.

2. Create an approved AI tool list

Make the safe path easier than the risky path. Tell employees what they can use, not just what they cannot use.

3. Write a plain data rule

Example: do not paste customer records, employee files, contracts, passwords, source code, or financial details into unapproved AI tools.

4. Use review levels

Low-risk work can move fast. Customer-facing, legal, financial, HR, security, and regulated work needs review.

5. Limit AI agents

Let agents read and draft first. Add approval before they send, delete, update, buy, or change anything important.

6. Track mistakes

A small AI mistake reported early is a win. A small AI mistake hidden for weeks becomes an incident.

7. Review vendor terms

Confirm whether your data is used for training, how long logs are kept, where data is stored, and who can access it.

Cybernomics rule: do not remove people from the loop. Put people in the right loop. Humans should review important decisions, not babysit every small AI task.

AGENT BOUNDARIES

Use a simple AI agent permission model

AI agents are different from chatbots because they can sometimes take action. That means SMB leaders need permission levels that are easy to understand.

Level	What the agent can do	Risk level	Required control
1. Read	View or summarize information for the user.	Low to medium	Limit data access. Log what the agent accessed.
2. Draft	Create emails, notes, tickets, reports, or recommendations.	Medium	Require human review before use.
3. Act with approval	Prepare an action, then ask a human to approve before sending, deleting, updating, or purchasing.	Medium to high	Approval workflow, owner, audit trail.
4. Act autonomously	Complete actions without human approval inside defined limits.	High	Narrow scope, monitoring, rollback, incident plan, executive ownership.

A safe first step

Start agents at read and draft. Add approve and execute only after the workflow has an owner, a review process, a rollback plan, and a way to detect mistakes.

Need help setting agent boundaries?
Cybernomics can help map where AI should read, draft, approve, and execute across your business workflows.
cybernomics.io/

90-DAY PLAN

A simple plan for business leaders

This plan is for an SMB that wants to use AI safely without spending months in meetings.

Time	What to do	Output
Days 1-15	Find where AI is already being used. Ask each department what tools they use, what data they enter, and what work AI helps with.	AI tool inventory and top 10 use cases.
Days 16-30	Pick approved tools. Block or discourage risky tools. Write a one-page AI use rule.	Approved AI list and simple policy.
Days 31-45	Map risky workflows. Focus on customer data, employee data, money movement, legal work, security work, and customer messages.	Risk map for high-impact workflows.
Days 46-60	Add review gates. Decide when AI output needs a manager, subject expert, or security review.	Review matrix by work type.
Days 61-75	Set agent boundaries. Decide what AI can read, draft, recommend, approve, and execute.	Agent permission model.
Days 76-90	Create a feedback loop. Track mistakes, near misses, useful prompts, time saved, and fixes made.	Monthly AI risk and value scorecard.

Do not wait for the perfect AI policy. Start with visibility, ownership, data rules, review gates, and a no-blame mistake channel.

ONE-PAGE RULE

A one-page AI rule your team can understand

Use this as a starting point. Keep the rule short enough that employees actually read it.

Company AI Use Rule

Use approved AI tools. Do not paste sensitive data into unapproved AI. Review AI work before it affects customers, money, legal matters, employee decisions, security, or regulated work. Report AI mistakes quickly. No blame for honest mistakes reported fast.

Five questions employees should ask	Five questions executives should ask
1. Am I using an approved tool? 2. Am I pasting sensitive data? 3. Could this answer hurt a customer, employee, or the company if it is wrong? 4. Does a human need to review this before it is used? 5. Do I know how to report a mistake or weird AI behavior?	1. Where is AI already being used? 2. Who approves AI tools and contracts? 3. What data is off limits? 4. Which AI outputs require human review? 5. How do we learn from AI mistakes without blaming people?

Want an editable one-page AI policy template?
Download the Cybernomics worksheet pack and adapt it for your business.
cybernomics.io/

SCORECARD

Cybernomics AI Human Risk Score

Score each item from 1 to 5. A score of 1 means "not started." A score of 5 means "working and reviewed." Total your score, then schedule an AI Readiness Review at Cybernomics.io to interpret the results.

Question	Score 1-5	Owner
We know which AI tools employees are using.		
We have an approved AI tool list.		
Employees know what data not to paste into AI.		
High-impact AI outputs require human review.		
AI agents have clear permission limits.		
We review AI vendor terms before turning tools on.		
Employees can report AI mistakes without fear.		
We track AI value, risk, and incidents every month.		

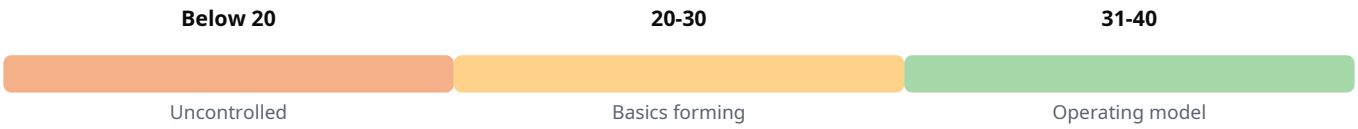

Goal: visibility and control. The score is a conversation starter, not a compliance certification.

Schedule your review.
Visit Cybernomics.io to schedule an AI Readiness Review and turn this score into a practical priority list for your business.
cybernomics.io/

WORKSHEET

AI risk register worksheet

Use the worksheet to track the main risks tied to your company's AI tools, AI workflows, and AI-enabled business processes. Keep it simple. One row should describe one risk in plain English.

Quick scoring guide: likelihood = how likely the risk is to happen. Impact = how much damage it would cause if it happens. Use a 1-5 scale.

Examples of AI risk: staff entering sensitive data into a public model, hallucinated outputs being treated as facts, too many AI tools being used without oversight, or unclear ownership when an AI workflow makes a mistake.

Bring this to your AI Readiness Review.
Bring this worksheet to your AI Readiness Review. Cybernomics can help turn it into a prioritized action plan, tool inventory, and simple AI policy starter.
cybernomics.io/

Tip: Review this register monthly. Start with the 5-10 risks that matter most. Update owners and next actions after each review.

AI RISK REGISTER WORKSHEET

Main register

Risk ID	AI Tool / Use Case	Risk Description	Likelihood (1-5)	Impact (1-5)	Current Controls	Owner	Next Action	Due Date / Status

Optional status labels: Open, In Progress, Accepted, Mitigated, Closed. Good owners: department lead, IT lead, security lead, operations lead, or tool owner.

Additional entries

Keep adding risks as new AI tools, vendors, and workflows enter the business. If the list gets long, group risks by department, tool, or use case.

Risk ID	AI Tool / Use Case	Risk Description	Likelihood (1-5)	Impact (1-5)	Current Controls	Owner	Next Action	Due Date / Status

NEXT STEP

How Cybernomics helps SMBs

Cybernomics helps SMB leaders move from AI curiosity to AI control. The goal is not to scare the business away from AI. The goal is to make AI useful, safe, and governable.

Find where AI is already being used

Interview departments, review tool usage, and identify shadow AI before it becomes a blind spot.

Identify risky workflows

Map where AI touches customer data, employee data, finance, legal, security, and customer communications.

Build simple rules employees can follow

Create a plain-language AI use rule, approved tool list, data boundaries, and review levels.

Create an AI risk register

Turn scattered concerns into owned risks, controls, next actions, and review dates.

Set human review points

Put people in the right loop for high-impact decisions without slowing low-risk work.

Prepare for safer AI adoption

Align tools, vendors, workflows, people, and leadership ownership before agents spread across the business.

Schedule an AI Readiness Review.
Use the report, scorecard, and worksheet to start the conversation. Cybernomics can help turn the results into a practical 30-60-90 day AI readiness action plan.
cybernomics.io/

Final word

AI will not wait for SMBs to become experts. It is already inside daily tools and daily work. The companies that do best will not be the ones with the longest AI policy. They will be the ones that make safe AI use simple.

SOURCES AND NOTES

Research basis

Primary source: Cybernomics episode featuring Dr. Nikki Robinson and Jennifer Baca discussing human factors, cybersecurity, psychological safety, AI use, agentic workflows, tool sprawl, and SMB adoption.

Source	How it was used
NIST AI Risk Management Framework	NIST describes the AI RMF as voluntary guidance intended to improve the ability to incorporate trustworthiness considerations into AI design, development, use, and evaluation.
McKinsey, The State of AI: Global Survey 2025	McKinsey reports that 88% of surveyed organizations use AI in at least one business function.
IBM, Cost of a Data Breach Report 2025	IBM reports a global average breach cost of about $4.4 million, and highlights the AI oversight gap around access controls and governance.
Stanford HAI, AI Index Report 2026	Stanford HAI reports that responsible AI practices are not keeping pace with AI capability and that documented AI incidents increased from 233 in 2024 to 362 in 2025.
Verizon, 2026 Data Breach Investigations Report	Public reporting on the 2026 DBIR highlights increased vulnerability exploitation, ransomware, third-party risk, and AI-enabled attack techniques.

Cybernomics note: This report is written for executive education and planning. It is not legal advice, audit advice, or a replacement for a formal security assessment.

Ready to move from AI risk to AI readiness?
Schedule an AI Readiness Review at Cybernomics.io and turn this report into your next 30-60-90 day action plan.
cybernomics.io/