EXECUTIVE BRIEF / 2026

AGENTIC AI RISK

// THE RISK YOU ARE INTRODUCING INTO YOUR BUSINESS

Why protecting the AI assistants on staff desktops does not address the AI now acting inside your operations — approving transactions, moving data, and shipping code at machine speed, with no human watching each step. Written for the C-suite.

// WHAT THIS BRIEF COVERS

01	02	03	04	05
THE SHIFT YOU'VE MADE	TWO VERY DIFFERENT RISKS	WHY IT REACHES THE BOARD	THE COVERAGE GAP	THE NEWORDER APPROACH

CYBER SECURITY THAT STOPS HACKERS

// THE SHIFT YOU HAVE ALREADY MADE

01 THE SHIFT YOU HAVE ALREADY MADE

You adopted AI to improve how the business runs, not only how staff write email. That AI is no longer just a chatbot on a desktop.

It is now an agent wired into your processes: approving transactions, updating records, moving data, writing and shipping code. It holds credentials, calls your systems, and acts on its own — at machine speed, with no human watching each step.

> **THIS IS WHERE YOUR REAL EXPOSURE NOW SITS — AND IT IS GROWING FAST.**

40%

of enterprise applications are expected to run task-specific AI agents by the end of 2026 — up from under 5% a year earlier.

SOURCE — GARTNER, 2025

// THE DISTINCTION THAT MATTERS

02 TWO VERY DIFFERENT RISKS

Most organisations are governing the AI they can see. The AI that can act is a different problem entirely.

THE VISIBLE RISK

DESKTOP AI

- ChatGPT, Copilot, Gemini on a desktop.
- A person prompts, the AI answers.
- Worst case, a staff member leaks a secret.
- Governable with policy, training, and data controls.
- You are most likely already managing it.

THE RISK YOU ARE INTRODUCING

AGENTIC AI

- Agents embedded in finance, operations, support, and code.
- They act — they do not just answer.
- Hijacked, one moves money, changes records, or ships code.
- Invisible to most controls, unpredictable, and usually ungoverned.
- **This is the exposure you are adding right now.**

> ONE ANSWERS. THE OTHER ACTS. ONLY ONE CAN BE HIJACKED INTO MOVING YOUR MONEY.

// A BOARD-LEVEL ISSUE, NOT AN IT DETAIL

03 WHY THIS REACHES THE BOARD

01 THE CONSEQUENCE IS OPERATIONAL, NOT INFORMATIONAL

The blast radius is not a leaked file — it is everything the agent is allowed to touch.

02 ACCOUNTABILITY HAS MOVED TO YOU

Regulators now expect rapid disclosure of material incidents and hold executives answerable for the gap.

03 THE EXPOSURE IS ALREADY LIVE AND UNMANAGED

Most organisations have already deployed agents — with no controls, no policy, and no visibility wrapped around them.

// THE NUMBERS BEHIND THE EXPOSURE

20%	+$670K	97%	63%	79%
of breaches now feature shadow AI	added to the average breach cost	of AI breaches hit orgs with no AI access controls	have no AI governance policy	have already adopted agents

SOURCES — GARTNER 2025; IBM COST OF A DATA BREACH 2025; ENTERPRISE AI ADOPTION SURVEYS 2025

// THE COVERAGE GAP

04 WHY SECURING DESKTOP AI DOES NOT COVER THIS

Policy, training, and data-loss controls govern what a person types into an assistant. They do not see the agent your team wired into a workflow last month, they cannot test how it behaves when an attacker manipulates it, and they cannot stop it from acting in real time.

Every agent is a non-human identity you never onboarded and never offboarded. Classic application testing assumes predictable software — but agents are non-deterministic, and they open entirely new attack classes.

// NEW ATTACK CLASSES AGENTS OPEN

PROMPT INJECTION

Hostile instructions smuggled into the content an agent reads.

AGENT HIJACKING

An attacker redirects the agent to act on their behalf.

TOOL ABUSE

The agent's own system access turned into the weapon.

DATA EXFILTRATION

Sensitive data drawn out through the agent's reach.

> **SECURING THE CHATBOT LEAVES THE PART THAT CAN ACT UNTOUCHED.**

// THE NEWORDER APPROACH

05 HOW NEWORDER ADDRESSES IT

We secure agentic AI across its whole life, as a managed, human-led outcome. Three moves — continuous, not one time.

01

DISCOVER

Inventory every AI agent, model, and tool across your environment — including the shadow agents nobody registered.

CONTINUOUS THREAT EXPOSURE MANAGEMENT

02

TEST

Red-team each agent before and after deployment against real adversary techniques, mapped to the OWASP LLM and Agentic Top 10.

ADVERSARY PATH ENGINEERING

03

PROTECT

Block prompt injection and agent hijacking at runtime, with human operators validating what the tooling flags.

HUMAN VALIDATION AT SCALE

Posture is mapped to NIST AI RMF, OWASP, and MITRE ATLAS, so your governance and reporting align to recognised standards. We engineer the adversary's path to your agents — then close it.

OUTCOME OVER OUTPUT

// THE NEXT STEP

FIND YOUR AGENTIC AI BEFORE AN ATTACKER DOES

Get your 30-minute Agentic AI Exposure Snapshot. We map every agent in your stack that holds credentials or acts autonomously, and rank your top hijack paths.

30 MIN	EVERY AGENT	TOP PATHS
NO-OBLIGATION DISCUSSION	MAPPED ACROSS YOUR STACK	HIJACK ROUTES RANKED

CAN YOUR BUSINESS AFFORD TO BE HACKED?

BOOK YOUR AGENTIC AI EXPOSURE SNAPSHOT

EMAIL INFO@NEWORDERGROUP.NET · WEB WWW.NEWORDERGROUP.NET

GET STARTED →