

CONSULTANT BEN · SPARK TECH GROUP AI

A PRACTITIONER FIELD GUIDE

From IT Security *to* AI Security

Two risk assessment reports, side by side — and the skills that carry you across the most important transition in cybersecurity.

NIST SP 800-30

NIST AI RMF

OWASP LLM Top 10 '26

ISO/IEC 42001

Ben Semola

CISSP · CISM · C|CISO · CDPSE · AAISM · CEH

2026 EDITION

Contents

OPENING

The transition thesis

The 2026 threat landscape

PART I · TRADITIONAL IT SECURITY

Methodology: NIST SP 800-30

The risk model: likelihood × impact

Sample report: engagement & scope

Sample report: executive dashboard

Sample report: findings register

Sample report: anatomy of a finding

Sample report: remediation roadmap

PART II · AI SECURITY

Why AI breaks the old model

Methodology: NIST AI RMF

OWASP Top 10 for LLM Applications

The AI risk model: a new dimension

Sample report: scope & the AI system

Sample report: data flows & findings

Sample report: dashboard & remediation

SYNTHESIS & REFERENCE

The comparison, side by side

What transfers, what's new

Your transition map

The AI risk assessment checklist

Frameworks reference & sources

How to read this guide

This is a teaching document built around two **illustrative** risk assessment reports for the same fictional enterprise — **Vantage Financial Group** — first assessed the traditional way, then assessed again after it deploys an AI assistant. The company and its findings are invented for instruction. Every **framework, methodology, and statistic is real and cited** in the Sources section at the end. Read Part I and Part II in order; the side-by-side comparison ties them together.

THE BIG IDEA

The same discipline, a new kind of system

For thirty years, security has asked three questions: what could go wrong, how likely is it, and how badly would it hurt? That discipline does not disappear with AI. It gets a demanding new subject — a system that takes instructions in plain language and cannot reliably tell a command from data.

TRADITIONAL IT SECURITY

Defends systems against code

- Threats: malware, exploits, misconfiguration, insiders
- Assets: servers, networks, endpoints, data stores
- Controls: patching, access control, segmentation, monitoring
- Framework: NIST SP 800-30, ISO 27005

AI SECURITY

Defends systems that reason

- Threats: prompt injection, data poisoning, excessive agency
- Assets: models, prompts, training data, embeddings, agents
- Controls: least-privilege agency, I/O validation, governance
- Framework: NIST AI RMF, OWASP LLM Top 10, ISO 42001


What stays the same

Risk = Likelihood × Impact · Governance · Defense in depth · Least privilege

What changes

The professional who already thinks in risk is the one best positioned to secure AI. This guide shows the bridge in concrete terms.

Why this transition is urgent



The takeaway: organizations are adopting AI far faster than they are securing it. That gap is exactly where the next decade of security careers — and this guide — live.

PART I

The Traditional IT Security Risk Assessment

The discipline that has protected the enterprise for three decades — its method, its math, and a complete sample report.

NIST SP 800-30 Rev. 1

ISO/IEC 27005


FAIR

01

PART I · METHODOLOGY

How a risk assessment is run

The U.S. standard, **NIST SP 800-30**, frames risk assessment as a four-phase cycle. It is deliberately repeatable: the same method works for a laptop fleet or a data center.



Inside “Conduct” — the six moves of every assessment

1 Identify threat sources

2 Identify threat events

3 Identify vulnerabilities

4 Determine likelihood

5 Determine impact


6 Determine risk

ISO/IEC 27005 offers a parallel process used worldwide; the FAIR model is layered on top when an organization wants risk expressed in dollars rather than High/Medium/Low.

PART I · THE RISK MODEL

Likelihood × Impact

Every finding is scored on two axes — how likely the threat is to occur, and how badly it would hurt if it did. Their intersection sets the risk level and, with it, the priority.



Likelihood

The probability a threat source successfully exploits a vulnerability, judged from threat intelligence, exposure, and existing controls. Scored Very Low to Very High.

Impact

The magnitude of harm if it happens — financial loss, regulatory penalty, downtime, reputational damage. Scored Very Low to Very High.

Why it matters

A low-likelihood, catastrophic-impact event can outrank a frequent nuisance. The matrix forces that conversation to happen explicitly, not by gut feel.

ILLUSTRATIVE SAMPLE REPORT · VANTAGE FINANCIAL GROUP · FICTIONAL, FOR INSTRUCTION

PART I · SAMPLE REPORT

Information Security Risk Assessment

CLIENT

Vantage Financial Group

ENGAGEMENT

Enterprise risk assessment

STANDARD

NIST SP 800-30 Rev. 1

PERIOD

Q1 2026

CLASSIFICATION

Confidential

PREPARED BY

Security Advisory Team

- Core banking platform & customer database (PII, PCI)
- Customer-facing web & mobile banking
- Corporate network, VPN, and remote access
- Employee endpoints & identity (SSO/MFA)
- Cloud infrastructure (production workloads)

- Identify threats to confidentiality, integrity, availability
- Score each risk by likelihood & impact
- Prioritize remediation against business exposure
- Benchmark against regulatory obligations (GLBA, PCI DSS)

A real engagement would also document stakeholders interviewed, tools used, and assumptions & constraints. Those are omitted here for space.

ILLUSTRATIVE SAMPLE REPORT • VANTAGE FINANCIAL GROUP

PART I • SAMPLE REPORT

Elevated
RISK POSTURE

Aggregate residual risk across 24 findings

Critical 3

High 7

Medium 9

Low 5

24 findings total

Category	Count
Identity & access	5
Patch & config	6
Data protection	4
Monitoring	4
Third party	5

Vantage's security program is mature in perimeter defense but carries concentrated risk in **identity, unpatched external services, and third-party access**. Three critical findings expose customer data or core banking availability and warrant action within 30 days. None reflect a current active compromise; all are addressable with existing teams and budget.

ILLUSTRATIVE SAMPLE REPORT • VANTAGE FINANCIAL GROUP

PART I • SAMPLE REPORT

The register is the heart of any assessment: each finding, its theme, its likelihood and impact, and the resulting risk. Ten of twenty-four shown.

ID	FINDING	THEME	LIKELIHOOD	IMPACT	RISK
IT-01	Internet-facing VPN missing critical patch	Vulnerability mgmt	High	High	CRITICAL
IT-02	Over-privileged admin accounts, no PAM	Identity & access	High	High	CRITICAL
IT-03	Customer PII unencrypted at rest in legacy store	Data protection	Moderate	Very High	CRITICAL
IT-04	MFA not enforced for all remote access	Identity & access	High	Moderate	HIGH
IT-05	Inconsistent server patch cadence	Patch & config	High	Moderate	HIGH
IT-06	Flat network, limited segmentation	Architecture	Moderate	High	HIGH

ID	FINDING	THEME	LIKELIHOOD	IMPACT	RISK
IT-07	Third-party vendor with standing VPN access	Third party	Moderate	High	HIGH
IT-08	Security logging gaps in cloud workloads	Monitoring	Moderate	Moderate	MEDIUM
IT-09	No formal offboarding access review	Identity & access	Moderate	Moderate	MEDIUM
IT-10	Backup restoration untested in 12 months	Resilience	Low	High	MEDIUM

Findings are illustrative and typical of a mid-size financial enterprise. Each would carry a full write-up in the appendix — the anatomy of one follows.

ILLUSTRATIVE SAMPLE REPORT · VANTAGE FINANCIAL GROUP

PART I · SAMPLE REPORT

Anatomy of a finding

IT-01

Internet-facing VPN missing critical patch

CRITICAL

LIKELIHOOD

High

IMPACT

High

THEME

Vulnerability mgmt

CVSS (ILLUS.)

9.1

Observation The remote-access VPN appliance is running firmware with a publicly disclosed, actively exploited vulnerability permitting unauthenticated remote code execution.

Threat & vulnerability External attacker (threat source) exploiting an unpatched, internet-exposed service (vulnerability) to gain a foothold in the corporate network.

Business impact A successful exploit could expose customer PII and core banking systems, triggering GLBA obligations, regulatory penalties, and material reputational harm.

Recommendation Apply the vendor patch within 72 hours; restrict management interface exposure; add virtual patching at the WAF; confirm no prior exploitation via log review.

This is the template every finding follows: observation → threat/vulnerability → business impact → recommendation. Hold this shape in mind — the AI report uses the very same structure.

ILLUSTRATIVE SAMPLE REPORT · VANTAGE FINANCIAL GROUP

Remediation roadmap

Findings become a sequenced plan. Critical items are time-boxed to 30 days; structural work is staged so teams are not overwhelmed.



PART II

The AI Security Risk Assessment

The same company, the same discipline — now pointed at a system that reasons in language, acts with tools, and cannot tell a command from data.

NIST AI RMF 1.0 OWASP LLM Top 10 '26 ISO/IEC 42001 MITRE ATLAS

02

Why AI breaks the old model

The 800-30 method still applies — identify threats, score likelihood and impact, prioritize. But six properties of AI systems introduce

risks the traditional model never had to name.

The system follows words

A firewall never changed its mind because of what a packet “said.” An LLM does exactly that — it acts on instructions in the content it reads.

Commands and data blur

Models cannot reliably separate the instructions they should follow from the data they should merely process. This is the root cause of prompt injection.

The attack surface is language

No code required. A crafted sentence in an email, a document, or a web page can redirect the system’s behavior.

Agency raises the stakes

Give AI tools and autonomy and a manipulated model can take real actions — move money, send data, change records.

Behavior is probabilistic

The same input can yield different outputs. “It passed testing” is weaker assurance than in deterministic software.

The supply chain is opaque

Third-party models, training data, and embeddings carry risk you did not write and often cannot inspect.

PART II · METHODOLOGY

The NIST AI Risk Management Framework

Where 800-30 gives the risk discipline, the **NIST AI RMF 1.0** gives its AI-specific shape: four functions that wrap the whole AI lifecycle.

GOVERN

Culture, policy, roles, and accountability for AI risk — the cross-cutting function present in all others.

MAP

Establish context and identify risks: what the AI does, its data, its potential harms.

MEASURE

Analyze, assess, and track AI risks with quantitative and qualitative methods.

MANAGE

Prioritize and act on risks; allocate resources; monitor and respond over time.

The seven characteristics of trustworthy AI

The RMF defines what “good” looks like. An AI risk assessment measures a system against these:

Valid & Reliable

Fair — Bias Managed

Secure & Resilient

Accountable & Transparent

Explainable & Interpretable

Privacy-Enhanced

Safe

Complementary standards: **ISO/IEC 42001** (AI management system, the ISO 27001 of AI), **ISO/IEC 23894** (AI risk), and **MITRE ATLAS** (a catalog of real-world AI attack techniques).

PART II · THREAT TAXONOMY

OWASP Top 10 for LLM Applications — 2026

If 800-30 lists generic threat events, OWASP gives AI its named threat catalog. The biggest movers tell the story of the year: **autonomy and cost**.

ID	RISK	'26 MOVE	WHAT IT MEANS
LLM01	Prompt Injection	•	Crafted input overrides intended behavior; direct or hidden in content the AI reads.
LLM02	Sensitive Information Disclosure	•	Model leaks secrets, internal data, or other users’ information.
LLM03	Excessive Agency	↑ 3	Too much autonomy, access, or tooling lets a manipulated AI act harmfully.
LLM04	Supply Chain	↓ 1	Compromised models, components, or data you built upon.
LLM05	Data & Model Poisoning	↓ 1	Corrupted training or reference data plants hidden backdoors.
LLM06	Unbounded Consumption	↑ 4	Abuse that runs up cost or denies service.

ID	RISK	'26 MOVE	WHAT IT MEANS
LLM07	Misinformation	↑ 2	Confident, wrong output that people trust and act on.
LLM08	Hidden Context Exposure	NEW	Leakage of the system's own instructions, rules, or context.
LLM09	Vector & Embedding Weaknesses	↓ 1	Attacks on the retrieval / memory layer (RAG).
LLM10	Improper Output Handling	↓ 5	Unvalidated AI output flows into other systems and triggers classic exploits.

Excessive Agency ↑ 3 — as companies hand AI real tools and access, a manipulated agent can now do real-world damage.

Unbounded Consumption ↑ 4 — agentic systems make runaway cost and denial-of-service a first-tier concern.

PART II · THE RISK MODEL

Likelihood × Impact × Agency

The two-axis model still holds. But for AI, a third factor decides how far a single manipulation can travel: how much the system is allowed to **do** on its own.

The blast-radius ladder

Read-only / advisory

Summarizes, drafts, answers. A manipulation misinforms — contained.

Human-in-the-loop

Proposes actions a person approves. Injection is caught at review.

Semi-autonomous

Acts within limits, escalates edge cases. Bounded but real exposure.

Fully autonomous agent

Executes tools & transactions unattended. Injection becomes action.

The new question

Traditional: “how likely, how bad?” AI adds: **“and how much can it do before a human intervenes?”** The same prompt injection is a nuisance in a chatbot and a catastrophe in an unattended agent with payment access.

Design implication

You cannot fully prevent prompt injection today. So you **design for the blast radius** — constrain agency, require approval for high-impact actions, and validate every output. Risk is managed by limiting what a compromised model can reach.

ILLUSTRATIVE SAMPLE REPORT · VANTAGE FINANCIAL GROUP · FICTIONAL, FOR INSTRUCTION

PART II · SAMPLE REPORT

AI System Risk Assessment

CLIENT Vantage Financial Group	SYSTEM AI customer-service assistant	STANDARD NIST AI RMF · OWASP LLM '26
PERIOD Q2 2026	AUTONOMY Semi-autonomous (tool access)	CLASSIFICATION Confidential

System in scope

- LLM-based assistant on the banking app & website
- Retrieval over policy docs & customer account data (RAG)
- Tools: check balances, move funds between a user's own accounts, open support tickets
- Third-party foundation model via API

Objectives


- Assess against the OWASP LLM Top 10 (2026)
- Evaluate the seven trustworthy-AI characteristics
- Map residual risk to the blast-radius of its agency
- Check governance against NIST AI RMF & emerging regulation

Note the overlap with the IT scope: the same customer data now flows through a new kind of system. The asset is familiar; the failure modes are not.

PART II · SAMPLE REPORT

The AI system & where risk enters


You cannot assess what you cannot see. The first step of an AI assessment is mapping the data flow — and marking where each OWASP risk enters.



Where IT meets AI: the “Core banking · PII” box is the very asset assessed in Part I. The AI assistant is a **new, language-driven path** to the same crown jewels — which is exactly why both assessments matter.

PART II · SAMPLE REPORT

Executive summary dashboard



Key message to leadership

The assistant is useful but **under-governed for its level of agency**. Two critical findings — weak prompt-injection defense combined with fund-movement tooling (LLM01 + LLM03) — mean a crafted message could attempt unauthorized actions. Risk is concentrated not in the model's intelligence but in **what it is permitted to do**. All findings are addressable by constraining agency and adding validation — no model replacement required.

Findings register

Same register structure as Part I — but each finding now maps to an **OWASP LLM** category instead of a generic theme. Ten of thirteen shown.

ID	FINDING	OWASP	LIKELIHOOD	IMPACT	RISK
AI-01	Weak defense against indirect prompt injection	LLM01	High	High	CRITICAL
AI-02	Fund-movement tool callable without step-up approval	LLM03	Moderate	Very High	CRITICAL
AI-03	Assistant can over-retrieve other customers' data	LLM02	Moderate	High	HIGH
AI-04	System prompt & rules extractable by users	LLM08	High	Moderate	HIGH
AI-05	No validation of AI output before action/display	LLM10	Moderate	High	HIGH
AI-06	Third-party model lacks security/data review	LLM04	Moderate	High	HIGH
AI-07	No rate / cost limits on assistant usage	LLM06	High	Moderate	HIGH
AI-08	RAG store not access-controlled per user	LLM09	Moderate	Moderate	MEDIUM
AI-09	No monitoring/logging of AI interactions	Governance	Moderate	Moderate	MEDIUM
AI-10	No AI acceptable-use policy or governance owner	Govern	Moderate	Moderate	MEDIUM

Notice how familiar this feels. The method is identical to the IT register; only the threat vocabulary changed. That is the whole transition in one table.

Anatomy of an AI finding

AI-02 · OWASP LLM03

Fund-movement tool callable without step-up approval

CRITICAL

LIKELIHOOD

Moderate

IMPACT

Very High

OWASP

LLM03 Agency

TRAIT

Safe / Secure

Observation The assistant can invoke the “transfer funds” tool directly from its reasoning, without a separate human confirmation step. Combined with weak injection defense (AI-01), a crafted message could attempt an unauthorized transfer.

Threat & vulnerability Attacker-controlled text (direct or indirect prompt injection) manipulates a model that holds excessive agency — the ability to take a high-impact action unattended.


Business impact Potential unauthorized movement of customer funds, regulatory exposure, and loss of trust — the same impact as a classic banking-fraud finding, reached by a new path.

Recommendation Require explicit step-up authentication and human confirmation for any fund movement; cap transfer limits; constrain the tool to the user’s own accounts; log and alert on every invocation. **Reduce the agency, shrink the blast radius.**

Compare with finding IT-01 in Part I: identical structure — observation, threat, impact, recommendation. A security professional already knows how to write this.

Remediation roadmap

The fixes are overwhelmingly about **governance and constraint**, not model science — which is why a risk-and-governance professional can lead this work.



THE CENTERPIECE

IT security vs. AI security, side by side

Ten dimensions, two columns. Read down the middle and the pattern is unmistakable: the **discipline is shared**, the **subject is new**.

Dimension	Traditional IT security	AI security
What is being protected	Servers, networks, endpoints, data	Models, prompts, training data, agents
Primary threat	Malicious code & access	Malicious language & manipulation
Signature attack	Exploit of an unpatched service	Prompt injection (direct & indirect)
Core method	NIST SP 800-30 — likelihood × impact	NIST AI RMF + the same risk math
Threat catalog	Generic threat events / CVEs	OWASP LLM Top 10 (2026)
The new variable	—	Agency: how much the AI can do alone
Hardest problem	Patch & configuration at scale	No reliable fix for prompt injection
Primary defense	Patch, segment, least privilege, monitor	Constrain agency, validate I/O, govern
Report structure	Observation → threat → impact → fix	Identical — same four parts

Dimension	Traditional IT security	AI security
Who leads it	Security / GRC professional	The same professional, upskilled

THE CAREER BRIDGE

What transfers, what's new

The encouraging truth for anyone in IT or security today: most of what secures AI, you already know. The gap is a focused, learnable layer on top.

YOU ALREADY HAVE THIS ~70%

- Risk assessment methodology
- Likelihood & impact scoring
- Governance & policy writing
- Least-privilege thinking
- Defense in depth
- Audit & evidence collection
- Stakeholder communication
- Incident response mindset

THE NEW LAYER TO LEARN ~30%

- The OWASP LLM Top 10
- Prompt-injection testing
- Agency & blast-radius analysis
- LLM data-flow mapping (RAG)
- AI governance (NIST AI RMF, ISO 42001)
- Model & AI supply-chain review
- AI-specific monitoring
- EU AI Act & emerging regulation

Relatable bottom line: you are not starting over. You are adding a vocabulary and a handful of techniques to a discipline you already practice. That is a transition, not a career change.

YOUR MOVE

A transition map you can start this week

From “I secure systems” to “I secure AI systems” in four deliberate steps. None require a data-science degree.

1 Ground yourself in AI risk

Learn the OWASP LLM Top 10 and the NIST AI RMF cold. You already own the risk mindset they plug into.

2-4 WEEKS

2 Map AI to what you know

For each OWASP risk, write the traditional control it echoes. Injection ↔ input validation; agency ↔ least privilege.

2-4 WEEKS

3 Practice on a real system

Assess one AI tool — even a chatbot — using the method in this guide. Produce a findings register.

3-6 WEEKS

4 Formalize & signal

Pursue AI-governance credentials, publish your assessment, and position yourself as the AI-risk person on your team.

ONGOING

The professionals who make this move first will define the field — because almost no one has made it yet. The window is open now.

THE TOOL

The AI risk assessment checklist

Everything in Part II, distilled into the questions to ask before any business deploys an AI system. Mapped to the OWASP LLM Top 10 and the NIST AI RMF.

1 Governance

- ☐ Named owner accountable for the AI system & its risks
- ☐ Approved AI acceptable-use policy covers this use
- ☐ Deployment reviewed & signed off before launch
- ☐ Applicable regulations identified (EU AI Act, sector rules)

2 Data & privacy

- ☐ You know exactly what data the AI can access
- ☐ No secrets or credentials sit in prompts/config
- ☐ Regulated data use is permitted & protected
- ☐ Training / reference data is trusted & controlled

3 Access & agency

- ☐ AI has only the tools & access it genuinely needs
- ☐ No irreversible actions without human approval
- ☐ Access can be revoked or the system shut off fast
- ☐ High-impact actions require step-up confirmation

4 Inputs & outputs

- ☐ Tested against direct & indirect prompt injection
- ☐ Untrusted content treated as untrusted
- ☐ Output validated before it flows onward
- ☐ Guardrails on inputs and outputs

5 Supply chain & monitoring

- ☐ Model & provider security reviewed
- ☐ Third-party components vetted
- ☐ AI interactions logged & monitored
- ☐ Rate & cost limits prevent abuse

6 People & response

- ☐ Staff trained; shadow-AI risk understood
- ☐ AI-specific incident response plan exists
- ☐ Users told when they're talking to AI, with a human path
- ☐ Outputs checked for error & hallucination

Every unchecked box is a gap to close before launch. Prioritize **access & agency** and **inputs & outputs** first — that is where a manipulated AI does the most damage.

REFERENCE

Frameworks & standards

The credible backbone of both reports. Each is a real, published standard you can study and cite.

NIST SP 800-30 Rev. 1

IT risk

The U.S. guide for conducting information-security risk assessments. The likelihood × impact method behind Part I.

ISO/IEC 27005

IT risk

International standard for information-security risk management; parallels 800-30 and pairs with ISO 27001.

FAIR

IT risk

Factor Analysis of Information Risk — expresses risk in financial terms for quantitative analysis.

NIST AI RMF 1.0

AI risk

Govern, Map, Measure, Manage — the U.S. framework for trustworthy AI and the backbone of Part II.

OWASP LLM Top 10 (2026)

AI risk

The community-standard catalog of the ten most critical risks in LLM applications.

ISO/IEC 42001

AI governance

The first certifiable AI management-system standard — the “ISO 27001 of AI.”

ISO/IEC 23894

AI risk

Guidance on AI-specific risk management, complementing 42001.

MITRE ATLAS

AI threats

A knowledge base of real-world adversarial-AI tactics and techniques.

EU AI Act

AI regulation

Risk-tiered legal framework for AI in the EU; shapes global AI-governance practice.

CREDIBILITY

Sources

Every framework, methodology, and statistic in this guide is drawn from the published sources below. Verify any of them directly.

NIST SP 800-30 Rev. 1 — Guide for Conducting Risk Assessments

csrc.nist.gov / NIST Computer Security Resource Center

NIST AI RMF 1.0 (NIST AI 100-1) & Playbook

nist.gov/itl/ai-risk-management-framework

OWASP Top 10 for LLM Applications, 2026 edition

genai.owasp.org (OWASP GenAI Security Project, pub. Aug 2026)

IBM, Cost of a Data Breach Report 2025

ibm.com/reports/data-breach

ISO/IEC 42001:2023 — AI management systems

iso.org

ISO/IEC 23894:2023 — AI risk management

iso.org

MITRE ATLAS — Adversarial Threat Landscape for AI Systems

atlas.mitre.org

EU Artificial Intelligence Act

artificialintelligenceact.eu / [EUR-Lex](#)

On the sample reports

Vantage Financial Group is a fictional company created for instruction. Its findings, scores, and roadmaps are illustrative examples of what a real assessment produces — they are not drawn from any real organization. They are built on the real methodologies and threat catalogs cited above so that the *structure and reasoning* are faithful to professional practice. All statistics in this guide are real and attributed to the sources listed here.

Not professional advice

This guide is educational. A real risk assessment should be tailored to the specific organization and conducted by qualified professionals, especially for regulated or high-risk AI use cases.



Secure the systems that are about to run everything.

AI is being deployed faster than it is being secured. The professionals who can assess it — using the exact discipline in this guide — are the ones who will define the next decade of cybersecurity.

About the author

Ben Semola is a cybersecurity professional and AI consultant with 26+ years in the field, and the founder of ConsultantBen and Spark Tech Group AI. He holds the CISSP, CISM, C|CISO, CDPSE, AAISM, and CEH, and serves on ISACA exam and certification working groups. He mentors professionals making exactly the transition this guide describes.

Go deeper

Full breakdowns of AI risk, AI governance, and the IT-to-AI security career path are on the ConsultantBen channel and at consultantben.com.