



CFA Risk Event

Kris Wulteputte, CRO

October 1st, 2026



Ground Control

AI-First Risk Management


At bunq, we pioneer an AI-first approach to keep innovating our integrated risk - compliance - legal approach.

bunq

What are the most common AI uses today?

What are the most common ways your team uses AI today?

% of 50 team lead responses · multiple answers allowed



AI USE CASES in Non-Financial Risk Management



Horizon Scanning

news, law & regulation developments, external incidents, reputation, third parties, emerging risks



Knowledge (RAG)

library of applicable law and regulation, in-house policy framework, regulatory letters and inspections, audits



Policy Work

review against new regulation and findings, update, drafting support



Internal Support

support agents -> risk analysis, AI risk assessment, DPIA, LIA; teaching e.g. privacy, sanctions, etc



QA & QC

board and senior management updates, regulatory & audit responses, terms & conditions, contracts, etc



Risk/Compliance Analysis

risk and compliance assessments e.g. product compliance, marketing review, strategy review, etc



Monitoring, Reporting

dashboards and reports [incident clustering & trend analysis, TPRM,...etc



Reporting Automation

board decks, internal reporting, incident reporting, etc



Phonebook

where to find owners, particular skills etc

Example - Non-Financial Risk - Intelligence & Monitoring

Regulatory Developments

Continuity of service - Q3 2026

Q2 2026

Low



Q3 2026

Low

SIGNAL COUNTS

Incidents	0
Open CAPs	1
RAS KPI Breaches	0
Risk Radar	0
Control Test Findings	0

ASSESSMENT

Residual Risk	Low
Direction	→ same
Confidence	moderate
Month in Quarter	3/3

ANALYSIS NARRATIVE

Regulatory Developments shows a quiet quarter with no incidents, KPI/KRI breaches, or open Risk Radar items. One open Internal Audit CAP (medium risk rating, due 2026-09-30, not overdue) is the only active signal. The Regulatory Watch control is in place.

Analysis Doc

GitLab Issue

Ground Control Bot is our autonomous risk operations agent that continuously scans, evaluates, and reports across all projects in real-time 🐧

Automated Continuous Data Collection:

Puts our **Monitoring & Reporting** framework into real-time practice by scanning all projects to aggregate:

- Incidents & control testing execution
- Audit findings & pentest vulnerabilities
- Risk Appetite KPI & KRI threshold breaches

Preventive Risk Analysis & Decision Support:

Shifts management from reactive oversight to **Preventive Risk Management**:

- Generates context-aware residual risk insights for driver owners
- Streamlines **Reporting Automation** for the Board & Risk Committee



Example - Governance - AI Risk Preliminary Checks

model-register#231 [🔗](#)

📄 Monitoring third party systems for status and availability

• [redacted] added **AI Risk pre-check** label 2 weeks ago

RiskOffice

 **AI Risk Agent** Maintainer 😊 ↩ ✎ ⋮

@project_10438_bot_7205455c009a115f6cf0b...
2 weeks ago

🔍 **Risk Classification:** AIActRisk: MinimalRisk

📌 **Assessment Note:** This system uses an LLM-based agent to scrape and parse third-party status pages, then sends Slack notifications and hourly summary updates to internal staff. While it qualifies as an AI system due to the LLM component inferring status from web page content, it does not trigger any high-risk use case — it monitors external service availability, not natural persons. It does not generate content presented to end-users, does not interact with customers as a chatbot, and does not produce synthetic media requiring transparency disclosures. The outputs are internal operational alerts with no impact on individuals' rights, creditworthiness, employment, or access to services. No prohibited, high-risk, or limited-risk transparency triggers apply.

• **AI Risk Agent** added **AIActRisk** **MinimalRisk** label 2 weeks ago

⚡ Optimized for Innovation & Accuracy

We designed the pre-check form to capture **only essential info** giving AI agents clear, grounded context to make reliable calls.


1. Anyone in bunq wants to launch a new AI feature.
2. Fills out a streamlined AI pre-check issue with zero unnecessary noise.
3. **AI Risk Agent** evaluates the idea against:
 - EU AI Act Guidelines
 - AI Risk Management Framework
4. **Instant Verdict:** Gives clear notes & decides immediately if a full risk review is needed!



bunq

Six Steps towards a Risk Specialist's AI Agent

Essential building blocks configured by domain risk experts for reliable performance



AI Set-up & Governance Principles

Set-up Principles

- **Ownership-based:** Risk team members build AI agents to support or automate their existing ownerships, rather than standalone AI projects.
- **Built by owners:** Domain experts build and refine their own agents ("everyone an engineer"), prioritizing where the impact is largest. Headcount is kept lean. No centralized agent development.
- **AI Roll-out:** Owners are responsible for their agents. Progress is measured via AIE2E methodology, prioritizing high-impact tasks. Redesign, not automate.
- **Governance Rail:** Required before going live: (1) Log in GSheet inventory → (2) Pass AI risk pre-check → (3) Register in GitLab.
- **Internal & 2nd-Line Support:** No autonomous customer-impacting decisions; a human owner always confirms output.

Governance Triggers

Trigger	Examples
User Impact	AI influencing support, claims, pricing, creditworthiness, or decisions affecting users.
Employee Impact	AI used for recruitment, performance review, monitoring, or task allocation for employee.
Regulated / Audit	AI drafting or reviewing material for compliance, legal, audit, or formal decision records.
Sensitive Data	AI utilizing personal, financial, confidential, or production data.
Operational Dep.	AI embedded in live high-volume workflows or core operations.
Fairness & Trust	AI that could create biased, inconsistent, or unexplained outcomes.

Risk Rule: If AI changes what a user gets, what a colleague decides, or what bunq records as evidence, bring Risk Management in early.

Practical Benefits & Drawbacks

Practical Benefits



- **Speed & 24/7 Access:** Lower error rate, instant and round-the-clock access to a vast knowledge base.
- **Volume & Scalability:** Delivers significantly higher throughput than a human team can maintain.
- **Consistency & Auditability:** Versioned, date-stamped scoring built on standardized, defensible workflows.
- **Scales Senior Expertise:** Extends 2nd-line expert capacity across jurisdictions, freeing time for judgment.
- **Regulatory Readiness:** Faster, better-documented responses for horizon scanning, gap analysis, and audit prep.

Practical Drawbacks & Risks



- **Agent Misalignment:** Poorly defined agents may not execute tasks strictly as expected.
- **Planning Overhead:** Preparation, implementation, testing take time - requires ongoing validation & strict governance.
- **Non-deterministic:** applies to the guardrails too.
- **Rapid Drift:** Effective solutions from three months ago may yield different results today.
- **Output Quality & Verbosity:** verbose or sub-optimal unverified analysis reaching critical decision folders. Analyst laziness or inexperience.
- **Core Imperative:** Balance innovation with continuous fine-tuning and critical oversight.

Fundamental Issues

Hallucinations & Errors

Missing background or context leading to poor conclusions.

Response: Human-in-the-loop, RAG grounding, constrained workflows & full reviewability.

Data Protection

Sensitive corporate, personal, customer data & credentials exposure.

Response: Approved tools only, input sanitization, pseudonymization & self-hosting.

Access Scope & Blast Radius

Unintended API actions on app/systems and broad external data egress.

Response: Least privilege / RBAC scoping, endpoint testing & human oversight.

Agentic Autonomy & Loops

Autonomous actions occurring at scale faster than humans can catch.

Response: Hard kill switch mechanisms and strict governance rules.

Third-Party & Concentration Risk (DORA)

Dependence on external model providers, GPU quota limits & availability errors.

Response: Self-hosting infrastructure, status monitoring & resilience workstreams.

Model Governance & Explainability

Regulatory compliance.

Response: Acceptable Use Policy; 3-step governance: register / AI risk check / formal inventory & version control: triage and red lines

Over-Reliance & Skill Erosion

Human domain expertise risks degradation as active practice diminishes and AI agents replace specialist functions. Cognitive offloading. Judgment erosion. Confirmation bias.

Response: Mandatory skill-retention frameworks, active human-in-the-loop oversight & continuous eval versioning. Deliberate friction and adversarial challenge. AI critiques.

Technical Challenges, Organization & HR

Technical Challenges

- **Rapid Tool Evolution:** Fast shifts in LLMs; avoiding lock-in.
- **Connector Bottlenecks:** Building & *securing* MCP connectors.
- **Integration Plumbing:** Permissions & orchestration across platforms.
- **Cost Management:** Tool usage costs scaling enterprise-wide; need active dashboarding and efficiency monitoring.
- **Quality & Upkeep:** Output QC, versioning & maintaining agents on lean headcount.

Organizational Challenges

- **Uneven Adoption:** Fast-moving pioneer teams vs. lagging business units.
- **"Everyone an Engineer":** Asking domain experts to build raises capacity & capability questions.
- **Grassroots Governance:** Keeping registration & pre-checks in step with rapid building.
- **Process Integration:** Embedding agents into core workflows rather than running side-tools.

Knowledge & AI Literacy

- **Targeted Hiring:** Recruiting AI natives and curious, continuous learners.
- **FrAIdays:** Dedicated recurring time for employees to build, test, and learn.
- **Knowledge Exchange:** Internal best practice exchange, central AI hub. Champions.
- **Adoption KPIs:** Setting clear metrics and targets to drive engagement.
- **Mandatory Training:** EU AI Act compliance, AUP review cycles & GDPR M10 modules.
- **Build-Your-Own Assistants:** Hands-on agent ownership serves as the primary learning model.

Conclusions

- AI-first gives the Risk Management Function wings.
AI increases **speed, quality** and **scalability**.
- AI is a **tool**, not a replacement. The **expertise** of your Risk Professionals is your **source** material.
And **data**.
- Risk Professionals **fine-tune** the agents continuously and critically, and keep **ownership**, keep the agents in check -> **brevity**, less is more!
- It's a lot of (rewarding) **work**. Don't be afraid to **try**.
- Keep **thinking** humans in the loop ;)
- Start from a **vision**, embrace the opportunity, find a way.

