

SCALING AI WITH CONFIDENCE

A Worked Executive Case in AI Transformation, Governance and Enterprise Value

Ajadi Ahmed

Chartered Safety & Health Practitioner (CMIOSH) - OSHCR Registered Consultant - BSc, PG Dip. Environment, Health & Safety.

Founder, IntelliSafely and VeriPrem

A practical, evidence-led approach to deciding where AI deserves investment, what must be true before it scales, and how leadership can govern consequential use without turning governance into a brake on innovation.

DISCLOSURE

Fiction Group plc is entirely fictional. Organisational characteristics, operating data, financial figures and projected benefits are illustrative assumptions created to demonstrate an executive decision-making approach. External regulatory, standards and research claims are separately sourced. This document is not formal advice.

BOARD DECISION

Executive recommendation

Approve a 24-month, evidence-led AI portfolio - not an enterprise-wide AI rollout.

Fiction Group should place capital behind a small number of deliberately different AI patterns, build the enterprise capabilities needed to govern them, and release further funding only when evidence shows that value, adoption, reliability and residual risk justify scale.

Opportunity	Current posture	Executive rationale
Enterprise Knowledge Assistant	PILOT	Lower-consequence productivity case; also builds reusable retrieval, evaluation and knowledge-governance capability.
Predictive Asset Intervention	VALIDATE DATA FIRST	Potentially material operational value, but only if historical data genuinely supports prediction.
AI-Assisted Field Risk Assessment	GOVERNED PILOT	Credible operational value; high consequence if wrong. Human safety authority remains explicit.
Customer-Service AI Agent	BOUNDED EXPERIMENT	Test assistive and reversible actions before increasing autonomy.
AI Recruitment Screening	DEFER WAVE 1	Legitimate problem, but current value case does not yet justify the higher assurance and rights-related burden.

Do not scale AI because a tool works. Scale it when value, readiness, control, adoption and evidence all justify further capital.

Enterprise dependency

- Federated AI operating model with central guardrails
- Enterprise AI inventory and named accountable owners
- Proportionate governance tiers and minimum evidence requirements
- Common benefits measurement and capital gates
- AI literacy and role-specific adoption capability

CONTEXT

Fiction Group: active with AI, not yet proven ready to scale

Dimension	Illustrative case assumption
Workforce	8,000 employees
Field workforce	c.3,000
Customer operations	c.1,000
Geography	Great Britain with limited EU operations
Activities	Field services, physical assets, customer operations and corporate functions
AI position	Multiple local experiments; limited enterprise coordination
Data	Strong reporting in some domains; fragmented operational data in others
Governance	Mature safety/privacy/IT controls; no common AI decision mechanism
Workforce	High curiosity; uneven AI fluency

A successful local pilot is evidence that a use case worked under particular conditions. It is not evidence that the organisation can reproduce the result across different data, users, workflows and risk environments.

The leadership question is not “Where can we use AI?” It is “Which problems deserve investment, under what conditions, and what evidence should be required before they scale?”

Board objectives

1. Improve productivity and service without treating workforce reduction as the sole source of value.
2. Use proprietary operational information more intelligently.
3. Prevent uncontrolled proliferation of generative and agentic AI.
4. Build an AI capability that can withstand board, workforce and regulatory scrutiny.

OPPORTUNITY PORTFOLIO

Five opportunities. Five different decisions.

Opportunity	Business problem	Likely AI pattern	Consequence if wrong
Knowledge Assistant	Time lost finding and reconciling approved internal information	GenAI + retrieval	Moderate
Predictive Asset Intervention	Maintenance effort not always directed to highest-risk assets	Predictive ML / analytics	Medium-High
Field Risk Assessment	Point-of-work assessments can be generic, slow or inconsistent	GenAI decision support	High
Customer-Service Agent	Routine demand consumes specialist capacity	GenAI / bounded agentic workflow	Medium-High
Recruitment Screening	High application volumes create manual screening burden	Classification / ranking	High

The 80/20 test

Before approving AI, ask whether process redesign, better data, workflow automation or existing software could deliver most of the value more cheaply and with less risk.

AI should earn its place by solving the residual problem that simpler interventions cannot address adequately.

Why this matters

- Prevents technology looking for a use case.
- Keeps business value separate from technical novelty.
- Makes “no AI needed” a valid value-creating decision.
- Forces leadership to define the consequence of error before enthusiasm drives scale.

DECISION FRAME

Five questions - kept separate on purpose

Question	What leadership is trying to establish
1. Value	Is the business outcome material enough to deserve capital and management attention?
2. AI suitability	Does AI add capability beyond simpler automation, analytics or process redesign?
3. Readiness	Can the organisation execute this use case safely, reliably and repeatedly now?
4. Consequence	What happens if the AI is wrong, over-trusted, biased, unavailable or misused?
5. Evidence strength	How much is demonstrated - and how much remains assumption?

Fiction Group deliberately avoids a single composite score. A use case can be valuable but not ready; technically strong but commercially weak; or attractive but too consequential to scale quickly.

Opportunity	Value	AI fit	Readiness	Consequence	Evidence	Current decision
Knowledge Assistant	High	Strong	Med-High	Moderate	Moderate	Pilot
Asset Intervention	High	Strong	Moderate	Med-High	Low-Med	Validate data
Field Risk Assessment	High	Strong	Moderate	High	Low-Med	Governed pilot
Customer Agent	Med-High	Strong	Moderate	Med-High	Moderate	Bounded pilot
Recruitment Screening	Moderate	Possible	Low-Med	High	Low	Defer

Source note: These ratings are fictional management judgements for the worked case - not externally measured facts.

READINESS

Readiness is not a percentage. It is a constraint.

Fiction Group assesses six interdependent dimensions: strategy, data, infrastructure, governance, talent and culture. Strength in one dimension does not cancel a threshold failure in another.

Use case	Binding constraint to test first
Knowledge Assistant	Authoritative knowledge ownership, permissions and source precedence
Predictive Asset Intervention	Historical data suitability, labels and intervention records
Field Risk Assessment	Controlled safety knowledge and safe human-AI interaction
Customer-Service Agent	Authority limits, escalation and reversible action design
Recruitment Screening	Governance, evidence and rights-related assurance burden

“Clean” data is not enough

Data readiness is workload specific. Reporting data may be perfectly adequate for dashboards while being unsuitable for prediction, retrieval or consequential AI decision support.

For retrieval / GenAI	For prediction
Authoritative sources	Representative history
Version control	Reliable target labels
Permissions	Stable definitions
Metadata and traceability	Sufficient observations
Freshness / precedence	Recorded interventions and outcomes

The useful output of a readiness assessment is a sequence: identify the binding constraint, fix it, then test whether the next constraint becomes limiting.

OPERATING MODEL

Federated execution. Central guardrails.

Layer	Accountability
Executive AI Investment Council	Portfolio priorities, capital, enterprise risk appetite, scale/stop decisions.
AI Governance & Assurance	Risk classification, evidence requirements, high-consequence review, policy, exceptions, monitoring standards.
Technology & Data	Approved architecture, identity, security, data services, evaluation and observability.
Business / Product Owners	Business outcomes, workflow redesign, adoption, domain controls and operational performance.

The governance forum governs the conditions. The business remains accountable for the outcome.

Decision rights

Decision	Decision authority / required concurrence
Problem is worth solving	Business executive
AI is an appropriate intervention	Business owner, with AI/technology concurrence
Data is fit for purpose	Data owner
Risk classification	AI Governance & Assurance
Domain risk is acceptable	Relevant domain owner
Pilot may commence	Executive sponsor + governance approval
AI may scale	Executive AI Investment Council
AI must be suspended	Named operational / risk owner

GOVERNANCE

Governance intensity should reflect consequence and autonomy

Fiction Group applies proportionate governance according to the potential impact of error and the degree of autonomy granted to the AI system. The internal tiers determine the level of approval, assurance, human oversight and monitoring required for each use case. Where external legal or regulatory classifications apply, those requirements take precedence.

Internal tier	Typical pattern	Minimum governance
Tier 1 - Standard	Low-impact drafting, summarisation or information retrieval	Local approval within enterprise controls; approved tools and data-handling requirements.
Tier 2 - Controlled	Internal decision support with limited consequence	Documented assessment, named accountable owner and defined monitoring.
Tier 3 - Enhanced	Safety, employment, financial or other consequential decision support	Formal impact and risk assessment, independent challenge, controlled pilot, explicit approval and defined human authority.
Tier 4 - Restricted autonomy	AI takes actions that may have a material effect on people, operations or assets	Executive or designated governance approval, strict permissions, intervention/override controls, auditability, fallback or rollback arrangements, and continuing assurance.

Governance standards and frameworks

- ISO/IEC 42001:2023 - specifies requirements for establishing, implementing, maintaining and continually improving an AI management system. Fiction Group uses it as the management-system foundation for enterprise AI governance. [5]
- ISO/IEC 42005:2025 - provides guidance for AI system impact assessment, including identifying, evaluating and documenting potential impacts throughout the AI system lifecycle. [6]
- NIST AI RMF 1.0 - provides a voluntary, non-sector-specific framework for managing AI risk, structured around Govern, Map, Measure and Manage. [7]
- NIST AI 600-1 - provides a Generative AI Profile and companion resource to AI RMF 1.0 for risks that are unique to, or exacerbated by, generative AI. [8]

INVESTMENT CASES

Three priorities - for three different reasons

1. Enterprise Knowledge Assistant - Pilot

Value hypothesis: reduce avoidable search and reconciliation effort while improving access to current, authorised information. The likely binding constraint is knowledge governance, not model capability.

- Controlled source corpus
- Permission-aware retrieval
- Visible source citations
- Conflict / ambiguity handling
- No silent invention of organisational policy

2. Predictive Asset Intervention - Validate data first

Value hypothesis: direct limited inspection and maintenance capacity towards assets with higher evidence-based intervention need. The first investment is a data-suitability test, not a model build.

- Are failure and intervention outcomes consistently recorded?
- Have definitions changed over time?
- Does inspection activity create sampling bias?
- Is there enough history for the required prediction horizon?

3. AI-Assisted Field Risk Assessment - Governed pilot

Value hypothesis: help competent operatives navigate controlled safety information and produce more relevant, timely site-specific assessments. AI supports access to relevant information and human judgement; it does not confer competence or determine that work is safe to proceed.

AI may retrieve, prompt, structure, suggest and cite. The competent operative observes, judges, escalates, stops work and decides whether work proceeds.

HIGH-CONSEQUENCE AI

Meaningful human authority - not a ceremonial “human in the loop”

HSE states that health and safety law applies regardless of the technology used, including workplace AI. It also states that AI uses affecting workplace health and safety require risk assessment and appropriate controls to reduce risk so far as is reasonably practicable, including consideration of cyber-security threats. [1]

What the pilot should test

- Correct retrieval of approved safety information
- Current version and clear provenance
- Material hazard omissions
- Inappropriate or unsupported control suggestions

- User recognition of uncertainty
- Challenge and escalation behaviour
- Independent hazard recognition before and after repeated AI use

HSE's 2026/27 work programme also includes research into cognitive offloading and cognitive resilience in AI-augmented safety decision-making. That does not prove a negative effect in Fiction Group; it does make human reliance a legitimate variable to test rather than assume away. [1]

A high "acceptance rate" is not automatically a success metric. It may indicate excellent recommendations - or automation bias. Acceptance should be assessed against independent expert review.

BOUNDARIES & RIGHTS

Autonomy should be earned. Consequential decisions should be evidenced.

Customer-Service Agent - start bounded

Autonomy level	Illustrative authority
0 - Retrieve	Find authorised information
1 - Recommend	Suggest answer / next action
2 - Draft	Prepare response; human sends
3 - Act within limits	Perform predefined reversible actions
4 - Bounded multi-step	Select and execute permitted steps with monitoring and escalation

Fiction Group starts at Levels 1-2. Movement to higher autonomy requires explicit allowed-action lists, transaction limits, identity controls, logs, exception handling and evidence that risk remains within tolerance.

Recruitment Screening - defer Wave 1

The business problem is legitimate, but Fiction Group has stronger first-wave opportunities. For UK deployment, the ICO's 2026 recruitment work recognises potential efficiency benefits from automated decision-making while emphasising transparency, fairness, meaningful human involvement and data-protection safeguards. [2][3] For EU operations, Annex III lists AI systems used for recruitment and candidate evaluation within the Article 6 high-risk classification framework, subject to the Article 6(3) derogation and the system's actual intended purpose. Under the consolidated text used for this case, the relevant Chapter III Sections 1-3 obligations for Annex III high-risk systems apply from 2 December 2027. Fiction Group should design any future EU deployment for those requirements now rather than retrofit controls later. [4]

The decision is not "AI recruitment is unacceptable." It is "other opportunities currently have a stronger value-to-assurance ratio."

PILOT DISCIPLINE

A pilot must prove more than technical performance

Gate	Core question	Illustrative evidence
Technical	Does the system perform sufficiently well?	Retrieval quality, prediction error, latency, permission integrity.
Operational	Does the redesigned workflow improve?	Cycle time, rework, throughput, handoffs.
Human	Do people use, challenge and escalate appropriately?	Adoption, overrides, correction, escalation, independent judgement.
Risk	Are controls effective and residual risk acceptable?	Incidents, unsafe outputs, fairness, privacy/security failures.

Fiction Group evidence maturity scale

Level	Meaning	Funding posture
E0	Hypothesis only	Explore cheaply
E1	Problem evidenced	Validate
E2	Feasibility evidenced	Controlled pilot
E3	Pilot result	Do not automatically scale
E4	Repeatability across contexts	Staged expansion
E5	Sustained production evidence	Sustain / optimise

This is Fiction Group's internal capital-allocation scale, not an external regulatory or standards classification. Capital follows evidence: a successful pilot is E3 - not permission to proceed directly to enterprise scale.

VALUE

Separate cash savings from capacity value

Value type	What counts
Hard financial	Cash cost removed or incremental revenue demonstrably generated.
Capacity	Time released for more productive work but not necessarily removed from payroll.
Operational	Cycle time, quality, throughput, rework, reliability.
Risk / control	Traceability, earlier intervention, assurance, lower control-failure exposure.
Strategic option	Reusable platforms, data, governance and workforce capability.

Illustrative Knowledge Assistant calculation

Assumption	Illustrative input
Eligible employees	1,500
Search / reconciliation time	25 minutes per week
Working weeks	46
Loaded labour rate	£38/hour
Gross addressable time	28,750 hours/year
Gross capacity value	£1.0925m/year

The £1.09m represents gross addressable capacity, not cash savings. The case applies an illustrative realisation factor before comparing realised capacity value with recurring operating cost.

Scenario	Capacity realised	Illustrative value	Less £180k recurring cost
Low	15%	£164k	-£16k
Base	30%	£328k	£148k
High	45%	£492k	£312k

Source note: All numbers on this page are Fiction Group assumptions/calculations, not external benchmarks or claimed client outcomes.

MEASUREMENT

What the Board should see

Lens	Board-level measures
Value	Benefits realised vs approved case; lifecycle cost; realised capacity; hard financial value.
Adoption	Workflow penetration; repeat use; competence; abandonment.
Reliability	Material error; corrections; false positives/negatives; source accuracy where relevant.
Risk	Incidents; escalations; unsafe or unfair outcomes; privacy/security exceptions.
Governance	Named owners; current reviews; unresolved findings; exceptions.

Lens	Board-level measures
Portfolio health	Pilots started; scaled; redesigned; stopped.

A mature AI portfolio should demonstrate the ability to stop weak initiatives, not only launch new ones.

Adoption is not login count

A technically strong tool can fail to create value if the surrounding workflow does not change. Fiction Group therefore measures whether work actually moves through the redesigned process - not simply whether employees opened the application.

- What work stops?
- What does AI now do?
- What does the human still own?
- What must be verified?
- What triggers escalation?
- What new monitoring data is created?

ROADMAP & DECISION

24 months: build discipline before scale

Phase	Focus	Exit condition
0-3 months	Inventory, decision framework, owners, risk tiers, baselines, literacy	Know what AI exists, who owns it, and which use cases deserve investment.
3-6 months	Prove three patterns: knowledge GenAI, predictive AI, high-consequence support	Technical, workflow, human and risk thresholds achieved.
6-12 months	Repeat in new team / context; revalidate benefits and controls	Value survives outside the initial pilot environment.
12-24 months	Industrialise selectively; platform reuse; supplier standards; retire weak systems	AI operates as a governed enterprise capability, not a pilot collection.

Board decisions requested

- Approve the first-wave portfolio and evidence-led funding model.
- Approve the federated operating model and AI Governance & Assurance capability.
- Require an AI Decision Record and named accountable owner for material AI systems.
- Release scale capital only when value, reliability, adoption, residual risk and accountable ownership meet agreed thresholds.

Scale is a consequence of evidence - not the objective of every pilot.

APPENDIX A

AI Decision Record - one-page control

Field	What must be recorded
Business problem	The outcome being improved - not the technology being purchased.
Accountable business owner	Named executive accountable for outcome and continued use.
AI / product owner	Operational owner for system performance and change.
People affected	Users and people materially affected by outputs/actions.
Data / knowledge	Sources, owners, permissions, provenance and critical limitations.
Allowed authority	What the AI may do.
Prohibited authority	What AI may not do.
Human authority	Who retains which decision and when escalation is mandatory.
Governance tier	Internal governance tier and rationale; external legal classifications assessed separately where applicable.

Field	What must be recorded
Evidence requirement	What must be proven before pilot / repeat / scale.
Monitoring	Value, reliability, adoption and risk KPIs.
Change triggers	Model/data/workflow changes that require reassessment.
Rollback / retirement	When the system is suspended or withdrawn.

Governance becomes real when decision rights, evidence requirements and stop conditions are explicit.

APPENDIX B

Assumptions register

ID	Assumption	Purpose	What a real engagement would require
A01	Workforce = 8,000	Scale context	Actual workforce / user segmentation
A02	Field workforce = 3,000	Safety case sizing	Actual operating population
A03	Knowledge users = 1,500	ROI illustration	Usage eligibility / task baseline
A04	25 min/week search effort	ROI illustration	Time study / survey / system evidence
A05	£38/hour loaded rate	ROI calculation	Finance-approved cost basis
A06	30% capacity realisation	Avoid 100% savings claim	Workflow evidence and benefits tracking
A07	£180k recurring cost	Sensitivity illustration	Supplier, platform, support and assurance TCO

Every financial number in the worked case should be traceable to this register. In a real engagement, assumptions would be replaced or progressively tightened with measured internal data.

APPENDIX C

AI risk register - starting taxonomy

Risk	Example manifestation	Primary control direction
Value	Technically successful system produces no measurable outcome	Baseline, benefits owner, stop condition
Data	Outdated, incomplete, biased or unpermissioned inputs	Ownership, provenance, testing, access control
Model	Hallucination, poor generalisation, drift	Evaluation, monitoring, fallback
Human factors	Automation bias, cognitive offloading, deskilling	Training, interface design, independent challenge
Operational	Integration failure or unavailable AI disrupts core process	Fallback, continuity, manual path
Fairness / rights	Materially different outcomes or opaque consequential decisions	Defined criteria, testing, transparency, recourse
Privacy	Inappropriate personal-data processing	DPIA / privacy governance as applicable
Security	Prompt injection, exfiltration, tool abuse	Security architecture and permissions
Supplier	Uncontrolled model/provider change	Contract, change notice, revalidation
Legal / regulatory	Deployment inconsistent with applicable obligations	Jurisdictional assessment and legal review
Autonomy	Agent acts outside intended authority	Tool permissions, limits, approvals, audit logs

APPENDIX D

Evidence register - verified source spine

Proposition	Primary authority	Status / limitation
GB H&S law applies to workplace AI	HSE policy statement, 12 Jun 2026	Regulatory position. Does not establish that a specific AI system is safe or compliant.
AI affecting workplace H&S requires risk assessment and controls	HSE policy statement	Regulatory position; exact controls remain use-case specific.

Proposition	Primary authority	Status / limitation
HSE is examining cognitive offloading	HSE 2026/27 research programme	Research direction, not evidence of a quantified effect in Fiction Group.
Automated recruitment can create efficiency and fairness/data risks	ICO, 31 Mar 2026	UK regulator expectations; case-specific legal assessment still required.
Recruitment/evaluation systems appear in Annex III	EU AI Act, Annex III	EU law where in scope; Article 6 rules, Article 6(3) derogation, intended purpose and the 2 Dec 2027 application date for relevant Annex III high-risk obligations matter.
AI management system structure	ISO/IEC 42001:2023	Recognised standard; not legislation.
AI system impact assessment	ISO/IEC 42005:2025	Recognised standard; not automatically mandatory.
AI risk management + GenAI profile	NIST AI RMF 1.0 / NIST AI 600-1	Voluntary framework + GenAI profile; not legislation or a legal compliance determination.

APPENDIX E

Use-case dossier template

Section	Questions
Problem	What is failing today? Who experiences it? What baseline proves it matters?
Value hypothesis	Which metric should improve and who owns it?
Why AI	What capability does AI add beyond simpler alternatives?
AI pattern	Generate, retrieve, predict, optimise or agentic action?
Data dependency	Which information is required, from whom, with what provenance and permission?
Affected people	Who uses the system and who is materially affected by it?
Consequence	What happens when the system is wrong, unavailable, unfair or over-trusted?
Human authority	Which decisions/actions remain human and when is escalation mandatory?
Readiness constraint	What is most likely to prevent reliable scale?
Pilot	Scope, users, duration, entry criteria and test conditions.
Exit thresholds	What evidence permits repeat / scale?
Stop condition	What evidence requires pause, redesign or retirement?

PRIMARY SOURCES

References and source notes

- [1] Health and Safety Executive (HSE). HSE's regulatory approach to Artificial Intelligence (AI). Policy statement, 12 June 2026. [HSE](#)
- [2] Information Commissioner's Office (ICO). Automated decisions can streamline the hiring process - with the right safeguards in place. 31 March 2026. [ICO](#)
- [3] Information Commissioner's Office (ICO). Recruitment rewired: an update on the ICO's work on the fair and responsible use of automation in recruitment. 2026. [ICO](#)
- [4] European Union. Regulation (EU) 2024/1689 (Artificial Intelligence Act), consolidated text current at 27 July 2026; see Articles 6 and 113 and Annex III. [EUR-Lex](#)
- [5] ISO/IEC. ISO/IEC 42001:2023 - Information technology - Artificial intelligence - Management system. December 2023. [ISO](#)
- [6] ISO/IEC. ISO/IEC 42005:2025 - Information technology - Artificial intelligence (AI) - AI system impact assessment. May 2025. [ISO](#)
- [7] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. January 2023. DOI 10.6028/NIST.AI.100-1. [NIST](#)
- [8] National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. July 2024. DOI 10.6028/NIST.AI.600-1. [NIST](#)