Acquisition Brief (EN) - AISafetyCase.com

Strategic domain for inspection-ready AI safety cases (draft v2026-02)

Asset offered

- **Domain name:** aisafetycase.com (.com, exact-match)
- **Nature:** descriptive digital asset, reserved as a neutral, vendor-independent banner for the emerging category “AI Safety Case”, i.e., a reviewable, inspection-ready argument supported by evidence, designed to support high-stakes deployment decisions for frontier and high-risk AI.
- **Not included:**
    - no certification, no regulatory status, no accreditation, no official label,
    - no audit, consulting, legal, compliance, safety engineering, security or assurance service,
    - no software, datasets, indices, proprietary methodology, or operational platform,
    - no claim of compliance, safety, security, performance, or “guaranteed trust”.

Contacts (suggested)

- **Site:** www.aisafetycase.com
- **Email:** contact@aisafetycase.com
- **LinkedIn:** www.linkedin.com/company/aisafetycase (if applicable)

This document - who is it for, why

This brief is intended for a C-suite / Board decision committee:

- CEO, CFO, COO, CRO, CAE (Chief Audit Executive), CISO, CTO, CIO, Heads of Risk / Safety / Assurance / Compliance,

- Procurement leadership (enterprise and public sector), audit & assurance leadership (internal and independent),
- AI governance, model risk management, safety engineering, cyber risk and resilience teams,
- General Counsel / Compliance, Corporate Development, M&A, Partnerships, standards and industry initiatives.

Purpose: assess whether **aisafetycase.com** should be secured as a category-grade banner for an institutional initiative centered on **reviewable AI safety cases**: structured claims, arguments, and evidence that can withstand governance review, procurement scrutiny, independent audit, insurer underwriting, and high-stakes deployment oversight.

This document is informational only. It is not legal, compliance, audit, security, financial, technical, or investment advice.

Disclaimers (must remain identical across site and documents)
“AISafetyCase.com is an independent, informational resource. It is not affiliated with any government entity, standards body, certification authority, or commercial provider.”
“Nothing on this site constitutes legal, compliance, audit, or security advice. Consult qualified professionals and primary sources.”
“The domain AISafetyCase.com may be available for institutional partnership or acquisition by qualified entities.”

---

1. Decision in one page

What it is
AISafetyCase.com is a category-grade .com designed to name a structural governance artifact for high-stakes AI: an **AI Safety Case**, i.e., a **reviewable, evidence-backed argument** that a system’s key risks are understood, mitigations are in place, and residual risks are explicitly stated for a defined operational context.

Category definition (short)
An **AI Safety Case** is a structured set of **claims, arguments, and evidence** that supports an explicit **go / no-go deployment decision** for an AI system in a given operational context.

Key attributes (non-technical)

- **Decision-grade:** built for deployment gating and governance review, not for marketing documentation.

- **Reviewable structure:** claims decomposed into auditable sub-claims and argument strategies.
- **Evidence-backed:** evaluations, red-teaming, controls, monitoring, incident handling, mitigations.
- **Operational-context bound:** explicit scope, assumptions, and constraints.
- **Residual risk statement:** what is evidenced, what is assumed, and what remains uncertain.
- **Updatable over time:** compatible with “dynamic safety case” thinking for post-deployment changes.

Why it matters now (signals, non-exhaustive)

- Frontier AI safety governance is increasingly described in terms of **safety case reviews** and pre-deployment decision gates.
- Public institutions and labs publish **templates and safety-case patterns** (including structured argumentation and inability arguments).
- Governance regimes converge toward **technical documentation, evidence quality, and continuous oversight**, making safety cases a natural, portable “inspection-ready” format.

---

2. What it is / what it is not

2.1 Natural scope (examples)

- Frontier AI and other high-stakes AI deployments where safety scrutiny is expected (critical infrastructure, defence and national security, finance, health, public sector).
- Enterprise procurement where deployment requires **reviewable evidence** and governance sign-off.
- Independent review, audit, assurance, and insurer underwriting where **structured evidence** reduces ambiguity.
- Cross-organisational settings where third-party AI must be evaluated in a repeatable, defensible way.

2.2 What it is not

- Not an audit firm, not a certification authority, not a regulator, not a standards body.

- Not a promise of compliance, safety, security, or performance.
- Not a commercial tool, platform, dataset, index, methodology, or service layer unless a future owner builds one independently.
- Not an endorsement of any vendor, lab, or institution.

---

3. Buyer set (who can rationally own it)

Frontier AI labs and safety institutes

- Entities formalising safety governance into reviewable artifacts and templates.

Audit, assurance, and risk governance

- Firms industrialising AI governance review, evidence workflows, and third-party inspection.

Insurance and reinsurance

- Underwriters, brokers, reinsurers, and risk modelers seeking standardised evidence bundles to price and cover AI risk.

GRC and model governance platforms

- Platforms extending into evidence packs, control testing, audit trails, and governance workflows.

Public sector, defence, and high-stakes alliances

- Multi-stakeholder initiatives that need a neutral banner for definitions, templates, and shared review language.

Typical sponsors
CRO, CAE, CISO, CTO, Head of AI Governance / Model Risk, General Counsel / Compliance leadership, VP Platform, Corporate Development.

---

4. Deployment options (examples, non-prescriptive)

A. Reference hub (public, neutral)
Definitions, glossary, curated primary references, and clear explanations of AI safety cases as decision artifacts.

B. Template and patterns library
Safety case templates, CAE-style patterns, “inability argument” structures, evidence taxonomy, and review checklists.

C. Governance review kit
A pragmatic "reviewable dossier" format, review workflow primitives, and role-based expectations (risk, legal, safety, security, procurement).

D. Dynamic safety case framing
Guidance on maintaining safety cases over time: versioning, incident updates, model changes, monitoring evidence, and post-deployment evolution.

Related category assets (optional, seller portfolio signal)

- aiassurancecase.com
- AISystemicRisk.com
- ModelSovereignty.com
- SyntheticAudit.com
- AuditableCompute.com
- SignedResponse.com
- ComputeIntegrity.com

---

5. Acquisition process (domain name only)

Typical institutional flow: **NDA → strategic discussion → formal offer → escrow → domain transfer**.
Unless explicitly agreed otherwise, the transaction covers only the **aisafetycase.com** domain name as an intangible digital asset. No software, datasets, indices, consulting, lobbying, infrastructure, licence, or service layer is included.

Initial contact for serious enquiries: **contact@aisafetycase.com**

Primary references (curated)

- UK AI Security Institute (AISI): safety cases and safety case templates (including inability arguments)
- Google DeepMind: Frontier Safety Framework and safety case review framing
- Anthropic Alignment: safety case components (ASL-4)
- arXiv: "Safety cases for frontier AI"
- arXiv: "Dynamic safety cases for frontier AI"
- UK Ministry of Defence (ASEMS / Def Stan 00-56): safety case definition lineage

- EU AI Act resources: technical documentation framing (Article 11 and Annex IV)
- NIST AI RMF 1.0
- ISO/IEC 42001: AI management systems