Trust Architecture Checklist

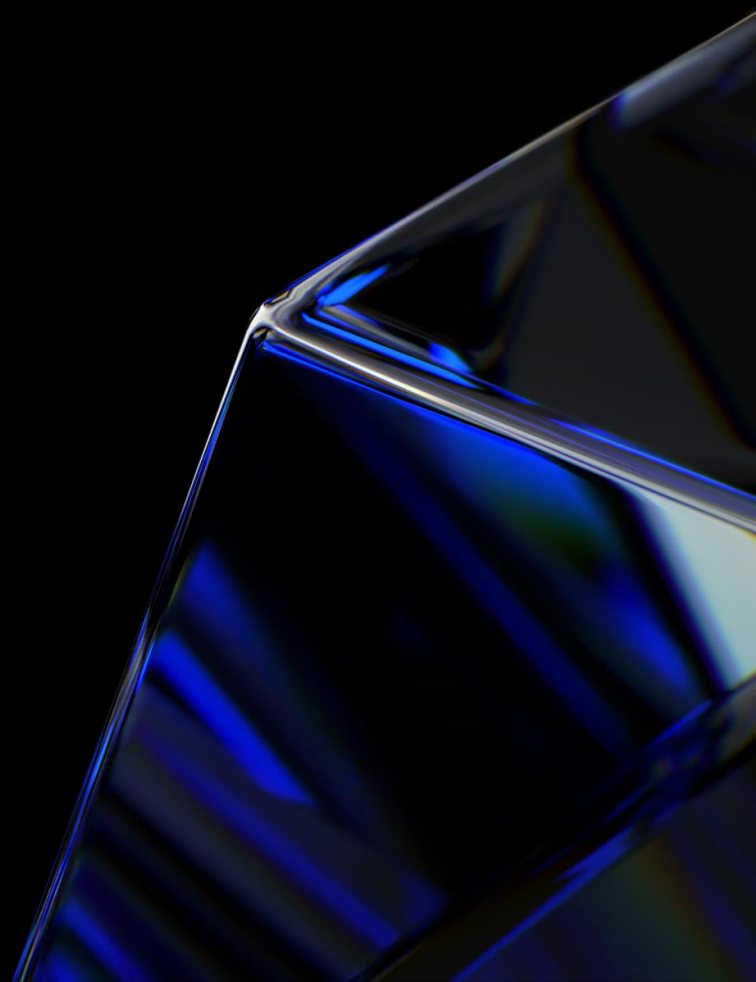

MERKLE

Trust Architecture Checklist

A Practical Guide to Scaling Ethical AI

Ethical AI is not a policy. It is an operating system property.

Use this checklist to assess whether your organisation is building AI on compliance or on trust architecture.

As you move through the checklist, consider the following signals related to each checklist section:

Red Flag: If 49% of leaders say they draw incorrect conclusions from data, have we validated that we are not in that 49%? (1)

Key Question: If a regulator or customer asked us why this outcome occurred, could we answer confidently? (2)

Reality Check: 84% of leaders say their data strategy requires overhaul before AI can succeed. Is ours future-ready? (3)

Signal: Trust erosion often happens before a human ever interacts with the customer. (4)

Final Check: If ethical AI governance disappeared tomorrow, would your organisation still build AI responsibly, or would decisions drift toward speed over trust? (5)

1. Data Foundation

☐ Do we have a clear inventory of the data feeding our AI systems?

☐ Is our data structured, consistent, and up to date?

☐ Are silos between marketing, product, CX, and operations limiting AI quality?

☐ Can we trace data lineage from input to outcome?

☐ Are we prioritising behavioural and voluntary signals over demographic shortcuts?

☐ Have we assessed bias risks in our segmentation models?

☐ Do we regularly audit for incorrect conclusions due to missing business context**?**

Trust Architecture Checklist
A Practical Guide to Scaling Ethical AI

2. AI Decision Logic

☐ Can we explain in plain language how our AI makes decisions?

☐ Do customers have clarity when interacting with AI-generated content or bots?

☐ Do we test recommendation and personalisation systems for fairness?

☐ Are exclusion or omission risks monitored?

☐ Is there human oversight for high-impact decisions?

☐ Do we monitor generative AI outputs for tone, hallucination, and brand alignment?

3. Operating Model

☐ Is ethical AI owned cross-functionally (data, tech, legal, privacy, marketing)?

☐ Do executives define clear AI values and non-negotiables?

☐ Is responsibility for AI outcomes clearly assigned?

☐ Do we have structured review processes before deployment?

☐ Are AI initiatives tested in sandbox environments before scale?

☐ Do we perform due diligence on AI vendors and partners?

☐ Is AI literacy embedded beyond technical teams?

4. Customer Experience

☐ Would a customer perceive our AI decisions as respectful and transparent?

☐ Are we avoiding demographic shortcuts & immutable characteristics?

☐ Are we designing segmentation based on behaviour and consent?

☐ Are we monitoring how AI assistants represent our brand in search and discovery?

☐ Is product information structured so AI systems interpret it accurately?

5. Governance & Growth

☐ Are guardrails clearly defined so teams can innovate safely?

☐ Do we know what is allowed, what requires oversight, and what is off-limits?

☐ Is speed supported by structure rather than slowed by escalation?

☐ Are sustainability & social impact of AI considered in deployment decisions?

☐ Do we treat ethics as a competitive differentiator, not a defensive necessity?

Trust Architecture Turns AI Risk Into Growth

Most organisations are scaling AI faster than their foundations.

Ethical AI is no longer about principles or isolated compliance checks. It is a structural capability that determines whether AI can scale across customer experience, commerce, and operations without creating risk.

Trust architecture addresses this at the source by aligning:

- Data foundations
- AI decision logic
- Operating models and accountability

In a focused session, Merkle supports you in assessing whether trust is structurally embedded in your AI ecosystem and where hidden risks may already exist.

All my best,

Clio Rossier
Senior CRM & Loyalty Consultant