NAVIGATING UNCERTAINTY

# AI Infrastructure Risk Audit

Demo 03 - Does an AI growth story have the compute, data and deployment base to support it?

| REPORT TYPE | DOMAIN | DEMO ENTITY |
| --- | --- | --- |
| Single Question Review | AI infrastructure | ModelForge AI, a synthetic model provider with high adoption and concentrated compute dependency. |
| **STATUS** | **PREPARED FOR** | **ADVICE STATUS** |
| Synthetic GBP 800 demonstration report | Public website demo | Not legal, financial, valuation or investment advice |

This demo uses synthetic values to show the fixed-scope review style: evidence, limits, risk and next checks.

AUDIT SUMMARY

# 1. Executive dashboard

ModelForge has strong usage and credible technical momentum, but scaling confidence is limited by compute concentration, cloud region exposure and incomplete deployment evidence. Product demand and infrastructure capacity are different signals.

**Client decision: treat the company as high-growth but infrastructure-constrained until compute dependency and deployment maturity are validated.**

| High | High | Medium | Med |
|---|---|---|---|
| Growth signal | Compute concentration | Deployment maturity | Regulatory drift |
| Users and API calls rising. | Few providers dominate. | Enterprise readiness partial. | Governance burden rising. |

| Section | What it shows |
|---|---|
| 1 Executive dashboard | Headline conclusion, demo scores and practical implication. |
| 2 Question and pipeline | Bounded GBP 800 scope and the repeatable review cycle. |
| 3 Data views | Market, usage, developer, infrastructure and geography evidence. |
| 4 LSS and metrics | How features become local structures and comparable risk signals. |
| 5 Drift and risk | Concordance, divergence, temporal movement and heatmap. |
| 6 Recommendations | Immediate actions, monitoring triggers and next checks. |
| 7 Limits | Assumptions, exclusions and terms. |

SCOPE

# 2. Question and pipeline

## Core audit question: Does the AI growth story have the compute, data and infrastructure base required to scale?

- Synthetic review of one AI company/model-provider.
- Focuses on usage, developer momentum, compute dependency, infrastructure exposure and deployment risk.
- Excludes model safety audit, legal compliance opinion, valuation and access to private cloud contracts.

## Framework cycle

| 01 Entity | 02 Features | 03 View vectors | 04 LSS | 05 Global metrics | 06 Temporal | 07 Dashboard |
|---|---|---|---|---|---|---|

| Pipeline stage | What happens | Example output |
|---|---|---|
| Preprocess entity | Align company, product, model, repo, cloud provider and region identifiers. | Clean entity profile |
| Features | Extract usage, revenue proxy, release, compute, data-pipeline and geography signals. | Comparable signal set |
| View vectors | Build market, usage, developer, infrastructure, geography and risk vectors. | Market / usage / developer / infrastructure / geography views |
| LSS | Compare the company against model labs, SaaS AI tools and infrastructure-scale peers. | Nearest peers, local density and structural position |
| Global metrics | Score adoption quality, compute concentration, deployment maturity and view disagreement. | Concordance, clustering and divergence |
| Temporal | Track usage stress, release cadence, cloud exposure and governance drift over time. | Drift, persistence and transition signal |
| Dashboard | Return capacity risks, escalation triggers and validation actions. | Risk score, decision gates and next checks |

EVIDENCE REGISTER

# 3. Data views

The audit does not rely on one headline metric. Each view contributes a different test of the claim.

| View | Evidence used | Why it matters | Example metrics |
|---|---|---|---|
| Financial value | Funding, revenue proxy, customer mix, valuation narrative | Tests external expectations. | Growth rank, revenue proxy, demand quality |
| Usage activity | Users, API calls, enterprise adoption, retention | Tests real product demand. | API growth, usage spikes, retention proxy |
| Developer/innovation | Model releases, papers, repos, issues, benchmarks | Tests technical momentum and maturity. | Release cadence, contributor concentration |
| Infrastructure | GPU access, cloud providers, regions, data pipelines | Tests scaling capacity and bottlenecks. | Compute concentration, cloud exposure |
| Geography | Cloud region footprint, data-centre exposure, jurisdiction | Tests spatial and regulatory concentration. | Region concentration, regulatory drift |

## Quality controls

| Control | Demo check | Purpose |
|---|---|---|
| Identifier alignment | Entities, sites, providers and regions matched across views. | Prevents false divergence. |
| Fixed windowing | Time windows are set before comparison. | Prevents cherry-picking. |
| Missingness labels | Unknown/private data is marked as caveat. | Prevents false certainty. |
| Source hierarchy | Direct, proxy and inferred evidence separated. | Keeps weak evidence from becoming the headline. |

STRUCTURE

# 4. LSS and metric layer

Features are turned into view vectors. View vectors form a local state space around the entity. The review then asks whether the entity sits where the public narrative says it should sit.

| Metric group | Example metrics | What it answers |
|---|---|---|
| Representation | PCA, UMAP, explained variance, trustworthiness | Can the evidence views be compared without losing structure? |
| Local state space | kNN graph, local density, nearest neighbours, centroid distance | Which assets, sites or systems are structurally close? |
| Taxonomy | Hierarchical clustering, spectral clustering, silhouette | Does the observed grouping match the stated label? |
| Concordance | ARI, NMI, Jaccard overlap, Cramer's V | Do views agree or contradict each other? |
| Anomaly | LOF, robust z-score, view-disagreement entropy | Where is the entity unusual or unstable? |

## Domain interpretation

| Finding | Interpretation | Decision value |
|---|---|---|
| Growth-capacity gap | Usage grows faster than verified compute capacity. | Separate demand from scalability. |
| Provider concentration | Cloud/GPU dependence is concentrated in few providers. | Map exposure before scaling claims. |
| Deployment uncertainty | Enterprise usage exists but operational maturity is partial. | Validate production readiness. |

RISK TRANSLATION

# 5. Concordance, drift and risk

Concordance shows whether views agree. Drift shows whether structure is changing. Risk turns both into decision priorities.

| Comparison | Relationship | Interpretation | Status |
|---|---|---|---|
| Market vs usage | Agreement | Growth narrative is supported by demand. | Stable |
| Usage vs compute | Divergence | Demand may exceed infrastructure comfort. | Flag |
| Research vs product | Partial agreement | Technical activity supports adoption but may overstate readiness. | Watch |
| Cloud vs risk | Agreement | Provider concentration increases operational and compliance risk. | Monitor |

| Risk | Severity | Likelihood | Why it matters | Control |
|---|---|---|---|---|
| Compute bottleneck | High | Medium | Demand may exceed capacity. | Capacity plan |
| Cloud concentration | High | Medium | Few providers create operational risk. | Provider map |
| Deployment overreach | Medium | Medium | Demo usage may not equal enterprise reliability. | Production evidence |
| Regulatory drift | Medium | Medium | Governance obligations may rise with scale. | Quarterly review |

NEXT CHECKS

# 6. Recommendations and monitoring

| Priority | Action | Reason | Cadence |
|---|---|---|---|
| Immediate | Create compute dependency map. | Main hidden risk. | Now |
| Immediate | Separate usage growth from capacity growth. | Prevents overconfidence. | Now |
| 30 days | Track cloud region and provider concentration. | Infrastructure risk is spatial. | Monthly |
| Quarterly | Review governance and regulatory exposure. | External constraints may drift. | Quarterly |

| Indicator | Cadence | Escalation trigger | Meaning |
|---|---|---|---|
| API usage growth | Weekly | Usage spikes beyond capacity plan | Demand stress |
| Compute availability | Monthly | Capacity fails to expand | Bottleneck risk |
| Cloud concentration | Monthly | Provider share rises | Operational risk |
| Governance tracker | Quarterly | New rule affects model/provider | Compliance risk |

# 7. Limitations and terms

- All entities, scores and values are synthetic demo examples.
- A real paid review requires client context, source validation and agreed scope.
- This is not legal, financial, investment, valuation, tax, security, engineering or regulatory advice.
- The GBP 800 style is one bounded question, one report and one clarification round unless agreed otherwise.

## Commercial value: the client is paying for a traceable evidence trail, not a generic opinion.

Glossary: View = evidence layer. LSS = local state space. Concordance = agreement between evidence views. Drift = movement in structure, behaviour or risk over time.