MOTHER — Sovereign AI, Defence & Robotics Estate

A technical due-diligence dossier covering the sovereign model family and every platform it powers — with an honest maturity status, security & compliance posture, and a revenue framing anchored to the MSAI Scotland data-centre go-live.

SUBJECT
Media Stream AI (MSAI) / MOTHER

ANCHOR MILESTONE
MSAI Scotland DC — Go-Live 2026

CORE ASSET
MOTHER CORE — sovereign UK foundation model

PREPARED FOR
Technical & commercial diligence

BASIS
Verified engineering state + stated roadmap

STATUS KEY
LIVE TESTING PLANNED

01 · EXECUTIVE SUMMARY

MOTHER is the substrate, not a single product

MOTHER CORE is a sovereign, UK-trained foundation model. It is not one application — it is the reasoning engine underneath a family of platforms spanning consumer & enterprise AI, creative media, defence decision-support, robotics, and the MSAI Cloud infrastructure that runs them all on sovereign hardware.

The strategic thesis for diligence is simple: **every platform in this dossier is powered by the same core model**. Revenue attributed to “MOTHER AI” as a standalone product therefore captures only a fraction of the value the model creates, because it excludes the defence, creative, robotics and infrastructure revenue that the same core enables.

Section 13 develops this into the reason the KPMG line item is a conservative floor rather than a ceiling.

Diligence should weigh two things honestly: the estate is **broad and materially built** — a training-stage sovereign model, ~100 production & domain agents, a merged production-grade security harness, a live deployment pipeline — and it is **pre-general-availability**: the core model is still in training, several platforms are in testing, and the flagship deployment is anchored to the Scotland DC go-live in 2026. This document labels each accordingly and does not overstate.

> COMMERCIAL VALIDATION — FIRST MAJOR PARTNERSHIP
>
> **MOTHER AI powers the AUTM.ai platform** — MSAI's first major commercial partnership, and the first external, revenue-bearing deployment of the sovereign core. For diligence this is the single most important traction signal in the estate: it moves MOTHER from "capable in the lab" to "selected to run a third-party commercial product," validating the model, the agent stack and the sovereign-serving model in a real customer setting (see §11 and §12).

> REVIEWER'S NOTE ON METHOD
>
> Status labels reflect the *verified engineering state* at the time of writing. Forward-looking dates are estimates against the stated 2026 DC milestone, not commitments. Revenue figures in §13 are an explicit framework with a placeholder for the client-supplied KPMG value — no financial actuals are asserted or invented in this document.

CONTENTS

02 · AT A GLANCE

Capability status matrix

Read this before the detail. Each estate area, its current maturity, and an estimated time-to-live indexed to the Scotland DC go-live (2026).

ESTATE AREA	STATUS	EST. TIME TO FULLY LIVE	GATING DEPENDENCY
Commercial — AUTM.ai partnership (MOTHER-powered)	● FIRST PARTNER · LIVE INTEGRATION	In market	Scale & expansion (metrics client-supplied)
Sovereign model — MOTHER CORE v5	● TRAINING	Q3–Q4 2026	Training completion + real-data eval pass
AI platform — B2C chat & assistant	● BETA	On DC go-live	Core model GA + harness enforcing
AI platform — B2B agents & API	● TESTING	2026 H2	Agent GA, OAuth connectors, SLAs
Creative platforms (IntuiTV / T2V / audio)	● TESTING	2026 H2	Gated media routes + GPU capacity (DC)
Defence platforms (COP / OverWatch / EXO)	● PILOT	Phased, 2026→2027	Accreditation, DPIA, per-deployment sign-off
Security — MOTHER Agent Harness	● MERGED / LIVE	Enforcing: 2026 H1	Persistence + external pen-test
Interoperability — Governed MCP Host	● BUILT / TESTED	Live pool: on node	Runtime server pool at agent node / DC
Infrastructure — MSAI Cloud & portal	● OPERATIONAL (PRE-DC)	DC: 2026	Scotland DC build & commissioning

ESTATE AREA	STATUS	EST. TIME TO FULLY LIVE	GATING DEPENDENCY
Internal — deployment pipeline	● LIVE	In use	—
Robotics — models / platforms / SIM	● SIM / R&D	2027+	Head training, sim-to-real, Guardian certification

HOW TO READ “LIVE”

“Merged / Live” for Security means the code is merged to the production branch, unit-tested and building cleanly; full *enforcing* posture (persistence, external pen-test) is the remaining step. “Live” for the pipeline means it is in daily use. Everything else is honestly pre-GA.

03 · SOVEREIGN MODEL FAMILY

MOTHER CORE and its lineage

A sovereign, UK-trained decoder LLM designed to run on owned hardware with no dependency on foreign API providers — the differentiator for defence and regulated customers who cannot send data to third-party models.

MOTHER CORE v5 — the frontier line

● TRAINING IN PROGRESS ● GA ~2026 H2

- **Architecture:** ~6.9B-parameter decoder — `dim 3072 · 48 layers · GQA 24/6 heads · RoPE · RMSNorm · SwiGLU · vocab 50,258`. Full-parameter training (not adapter-only).
- **Training substrate:** NVIDIA GB10 (Grace-Blackwell) sovereign node, ~122 GB unified memory; balanced-category sampler with a live **priority boost on maths, code and security/cyber-defence**.
- **Corpus:** ~2.33M curated records after golden-dataset harvest; checkpointing with auto-prune and resumable state.
- **Current state:** actively training past chunk ~1,785 at a healthy loss band (~3.7–4.5, perplexity ≈ 40) — consistent with a genuinely-learning mid-training model.

Lineage & recovered assets — diligence discipline

A rigorous audit of archived weights was performed. It is included here because *the discipline is itself a diligence signal*: assets were promoted or excluded on measured evidence, not optimism.

ASSET	FINDING	DISPOSITION
v12_7b lineage	Same architecture as v5 (dim-3072/48L); serving & “best-of” checkpoints intact and mergeable	● RETAINED — FRONTIER CANDIDATES
epoch_3 (7B)	Recorded training loss 6.27 (ppl ≈ 528); 0% teacher-forcing accuracy — undertrained	● EXCLUDED FROM BLEND
TRM specialists	Recursive reasoning heads bound to a base model; not standalone generators	● ARCHIVED — RESEARCH

> WHY THIS MATTERS
>
> The recovered-asset review demonstrates a measured, eval-gated merge methodology (perplexity + generation tests, model-soup/TIES only among same-architecture, confirmed-strong weights) rather than uncritical model hoarding. Frontier growth is planned from the *measured-strong* v5/v12_7b set only.

Sovereign deployment property

Models are served on the GB10 node behind the MSAI gateway. The commercial and compliance significance: **customer data need never leave sovereign infrastructure** to be reasoned over — the property that underwrites the defence and regulated-enterprise proposition and much of §14–15.

04 · AI PLATFORMS — B2B & B2C

MOTHER Core application layer

● B2C BETA ● B2B TESTING

A Next.js 14 platform (motherai.uk) delivering reasoning + retrieval-augmented answers, a large catalogue of production task-agents, and an orchestration layer — the same core serving both consumer and enterprise surfaces.

B2C — Assistant & chat

Consumer reasoning, deep research, document building and retrieval-grounded answers via the sovereign core. Tiered accounts, billing and rate-limits in place.

B2B — Agent platform

~70 production task-agents (documents, code, Google & Microsoft suites, Slack/Notion/CRM connectors, résumé/CV) exposed via an authenticated agent service with an allow/ask/deny guardrail and human-approval queue.

Grounding & memory

Retrieval via the Onyx engine over the corpus + team documents; per-user persistent memory; live-web grounding gated per request.

Orchestration

Planner→executor→verifier→summarizer multi-agent loop for larger builds, with a bounded in-route tool harness.

Agent estate: 70 served production agents + 33 defence/robotics domain agents catalogued. Diligence note: the 33 domain agents are catalogued but flagged ● NOT YET TRAINED IN V3 — an honest indicator of remaining work, not a claim of completeness.

05 · CREATIVE PLATFORMS — B2B & B2C

IntuiTV & the generative-media stack

● TESTING ● GA ON DC GPU CAPACITY

A text-to-video and generative-media capability spanning consumer creation tools and enterprise content pipelines, assembled on a curated model stack and now being brought under the same harness & auth as every other surface.

CAPABILITY	MODEL STACK (STAGED)	STATUS
Text-to-video	HoloCine · LTX-2 · Wan 2.1 / 2.2 · CogVideoX · AnimateDiff	● TESTING
Image	SDXL	● TESTING
Audio / voice	MusicGen · XTTS-v2	● TESTING

CAPABILITY	MODEL STACK (STAGED)	STATUS
Gated delivery route	`/api/t2v — authenticated + harness-guarded`	● BUILT

DILIGENCE FLAG — RESOLVED IN BUILD

The T2V interface previously had no dedicated authenticated route and defaulted to the agent-service port. A first-class authenticated, injection-screened proxy route has now been built (see §7/§13). Full GA is gated on GPU capacity at the Scotland DC.

06 · DEFENCE PLATFORMS

MOTHER Defence — decision support, not weapons

● PILOT ● PHASED ACCREDITATION

An observe-and-advise common operating picture and ISR analytics layer. The posture is explicit and load-bearing for compliance: **Strike = 0** — the system refuses target selection, fire-control and kill-chains, and every action is human-in-the-loop.

Common Operating Picture

50+ live data feeds fused into a single operational picture (external open-data + sensor sources).

OverWatch

Large-scale surveillance & drone-ISR visualisation (153k-camera scale) with 3D globe and streaming.

DeepVision EXO

Detection, tracking, identification, scene-assessment and world-model weights served through the EXO inference chain over operational frames.

Domain reasoning agents

25 defence analyst agents (intel-fusion, threat-assessment, GEOINT/OSINT/SIGINT, maritime, air-defence, cyber-threat, counter-UAS advisory) — all under an ADVISE-only, non-kinetic mandate.

COMPLIANCE SENSITIVITY — DECLARED

Surveillance and defence processing carry the highest regulatory load in this estate (AI Act risk-tiering + GDPR special-category / DPIA obligations — §14–15). These platforms

are correctly held at **pilot** pending per-deployment accreditation and data-protection impact assessment; they are not represented as generally deployable.

07 · SECURITY

The MOTHER Agent Harness

● MERGED TO PRODUCTION ● 47 TESTS PASSING

A production-grade security & engineering layer wrapping every agent turn on four pillars — **Execution · Monitoring · Control · Recovery** — with a zero-trust spine and a continuous, self-learning internal red-team. This is the most mature asset in the estate: designed, built, tested and merged during diligence.

Control

Input prompt-injection / jailbreak / secret-elicitation / Strike-0 detection on user input *and* retrieved content; output secret & PII redaction; deny-by-default capability policy with human-in-the-loop on all send/post actions.

Monitoring

Tamper-evident, hash-chained audit of every decision with cryptographic verification.

Recovery

Circuit breakers, token-bucket rate-limiting, per-agent kill-switch and actor quarantine.

Autonomous security agents

Sentinel (anomaly detection), Self-Repair (learns a block signature from any red-team bypass and patches the live guard), and a measured Posture score.

Verifiable claim: a test demonstrates the closed self-repair loop — a novel attack the guard *allows* → red-team flags it → the agent learns a signature → the same attack is subsequently *blocked*. The layer maps to OWASP LLM Top-10 and was merged across three reviewed pull-requests.

HONEST SECURITY POSTURE

No part of this estate is described as “un-hackable”. What is delivered is defence-in-depth: minimal attack surface, deny-by-default control, contained blast radius, a tamper-evident trail, and a system that continuously tests and repairs itself — with an

external penetration test as the outstanding independent validation before any “fully hardened” claim.

08 · INFRASTRUCTURE SOFTWARE — MSAI CLOUD

Sovereign deployment fabric

● OPERATIONAL PRE-DC ● SCOTLAND DC 2026

MSAI Cloud is the orchestration layer that routes builds and services onto the right target — managed frontend hosting for the web tier, and the GB10 sovereign nodes for model & service workloads.

- **MSAI Cloud Portal:** receives deploy webhooks and routes each workload (frontend → managed hosting; services/models → GB10 node) — the control plane for sovereign delivery.
- **Compute:** GB10 (Grace-Blackwell) node(s) today; the **MSAI Scotland data centre (go-live 2026)** is the anchor that scales this from single-node to fleet and unlocks GPU capacity for creative + defence workloads.
- **Telemetry:** training and eval streams flow node → platform → dashboard for live operational visibility.

ANCHOR MILESTONE

The Scotland DC go-live is the single most important dependency across the estate: it converts a working single-node capability into a scalable sovereign cloud and is the gating event behind most of the 2026 dates in the matrix.

09 · INTERNAL SYSTEMS — SALES & DEPLOYMENT

Delivery & go-to-market plumbing

● DEPLOYMENT PIPELINE LIVE

- **Deployment pipeline (live):** commit → CI → MSAI Cloud Portal webhook → routed release (frontend / node). In daily use; it is how the platforms in this dossier ship.
- **Entitlements & billing:** tiered accounts, usage metering, rate-limits and payment integration on the AI platform.

- **Sales / deployment motion:** B2C self-serve on the platform; B2B via agent deployments, API access and (roadmap) white-label sovereign instances routed through MSAI Cloud.

10 · ROBOTICS — MODELS · PLATFORMS · SIM

Embodied MOTHER & the simulation estate

● SIM / R&D ● SIM-TO-REAL 2027+

A cognitive stack for humanoid and fleet robotics under the same non-kinetic, human-supervised safety posture, developed against a simulation-first test environment before any physical deployment.

Cognitive model

`mother-robotics-7b-n1` — an L0–L4 task-planning stack for the Fourier N1 humanoid, under Guardian sign-off; human-supervised, non-weapon.

SIM test environment

EXO simulation (`mother_exo_sim`, `/api/v1/exo/*`) — detect / track / assess / plan exercised in simulation before hardware.

Robotics agents

8 domain agents: telemetry, safety-guardian, mission-planner, fleet-coordinator, flight-policy advisor, maintenance-predictor, CAD-design assistant, swarm-strategist.

Safety posture

Strike = 0 across flight / driving / humanoid; observe-and-advise; Guardian certification gate before autonomy.

HONEST MATURITY

Robotics is the earliest-stage area: models & agents exist and are exercised in simulation, but the domain heads are not yet trained in the current line and sim-to-real + Guardian certification remain. This is correctly represented as R&D/SIM, not a shipping product.

11 · AGENT DEPLOYMENT, HARNESS & PROTOCOL INTEGRATION

MOTHER’s own agent deployment & interoperability

MOTHER runs its agents through its own harness (§7) and speaks the emerging agent-interoperability protocols, so it can both host agents and integrate into third-party agent ecosystems.

- **Own deployment:** agents are served from an authenticated agent service, governed by the harness (capability policy, guardrail allow/ask/deny, sandboxed execution, human-approval queue) — a uniform control plane across all ~100 agents.
- **Protocol integration:** agent-to-UI (AG-UI) protocol, Anthropic tool-use / MCP-style tool servers, and OAuth connector bundles for Google & Microsoft estates — the interoperability surface for embedding MOTHER agents into external workflows.
- **Self-defending agents:** the harness's security agents monitor and self-repair the agent fleet continuously (§7).

Governed MCP Host — interoperability, hardened

● GOVERNANCE + RUNTIME BUILT & TESTED ● LIVE POOL AT NODE/DC

MOTHER hosts external tool servers through the industry-standard **Model Context Protocol** (the same `mcpServers` config Claude Desktop and Cursor use), aggregating many servers into one tool namespace — but *governed*, not raw. Every tool exposure and call passes the harness: **deny-by-default** tool exposure (a server exposes nothing unless explicitly allowlisted), tool arguments **injection-screened** before dispatch, Strike-0 kinetic deny + HITL on mutating tools, a child-process **env allowlist** (a server never inherits platform secrets), and full audit.

Its commercial relevance: the standard config format is the **partner-integration surface** — the same governed MCP path that carries the AUTM.ai integration is the repeatable, secure template for the next partners. End-to-end the flow is `model → governance → transport → live server pool → tool`; the governance, transport and runtime are built and tested, with the live multi-server pool exercised on the agent node where the servers run.

AUTM.ai — first major commercial partnership

● LIVE INTEGRATION · FIRST COMMERCIAL PARTNER

MOTHER AI powers the AUTM.ai platform. This is MSAI's first major commercial partnership and the first external, revenue-bearing deployment of the sovereign core — the AUTM platform runs on MOTHER's reasoning and agent stack rather than a third-party frontier API.

Its diligence weight is disproportionate to its size for three reasons:

- **Product-market validation:** an external company selected MOTHER to run a commercial product — independent evidence that the model and agent harness are production-viable, not just internally impressive.
- **Sovereign-serving validation:** it proves the deploy-on-owned-infrastructure model works for a paying third party — the exact property that unlocks defence and regulated customers.
- **Repeatable motion:** the same integration surface (agent service + AG-UI / MCP / tool-use protocols + OAuth connectors) that carries AUTM.ai is the template for the next partners — the first reference deployment de-risks the pipeline behind it.

> INPUT TO COMPLETE
>
> Commercial specifics — AUTM.ai's vertical, user / revenue scale, contract term and go-live date — are the client's to supply and should be inserted here to quantify the partnership. This document states the verified fact (MOTHER powers AUTM.ai; first major commercial partnership) without asserting figures it cannot source.

12 · REVENUE FORECAST & THE CONSERVATISM CASE

Why the MOTHER line item is a floor, not a ceiling

The revenue question in this brief is best answered as a *methodology* plus a structural argument, because the responsible way to present a forecast in diligence is to show the mechanism — not to assert a number the reviewer cannot trace.

> INPUT REQUIRED — NOT INVENTED
>
> This document does **not** reproduce or estimate the specific KPMG forecast value, which is the client's figure to supply. Insert the KPMG "MOTHER" line-item revenue here → [`£X.X m — FYXX, per KPMG model`]. The analysis below explains why that figure most likely *understates* the model's economic contribution.

The structural argument (qualitative, robust)

A forecast that books revenue against "MOTHER AI" as a standalone product typically counts only **direct** model monetisation — consumer subscriptions and model/API usage. But MOTHER is the reasoning core inside every other line of business:

- **Defence** decision-support contracts run on MOTHER reasoning + DeepVision — high-value, but usually booked as “defence,” not “MOTHER.”
- **Creative / IntuiTV** SaaS is orchestrated by MOTHER agents — booked as “creative.”
- **Robotics** cognition is MOTHER-derived — booked as “robotics.”
- **Agent platform & connectors** (B2B) are MOTHER agents — often booked as “platform.”
- **MSAI Cloud** infrastructure revenue exists *because* sovereign MOTHER workloads need somewhere to run — booked as “infrastructure.”

So the conservative reading is deliberate double-counting avoidance: the KPMG “MOTHER” line is the **direct-monetisation floor**. The **enabled** revenue — the platforms that cannot exist without the core — is the true economic footprint, and it is distributed across the other line items.

This is already observable, not hypothetical. The AUTM.ai partnership is a live external platform running on MOTHER — commercial revenue that MOTHER creates but that is most naturally booked as a partnership/platform line, not as “MOTHER.” It is the first concrete instance of the multiplier: one core, powering revenue booked elsewhere. Each additional partner on the same integration surface compounds it.

Illustrative attribution frame — assumptions, not a forecast

> ILLUSTRATIVE ONLY
>
> The table below is a **framework with placeholder assumptions** to size the multiplier effect. Every cell is an assumption to be validated against the client model and market data — it is *not* a revenue projection and must not be read as one.

LINE OF BUSINESS	BOOKED AS	MOTHER DEPENDENCY	MOTHER-ATTRIBUTABLE SHARE (ASSUMPTION TO VALIDATE)
Direct model — B2C subs + API	MOTHER	100%	= KPMG line item
B2B agent platform & connectors	Platform	Core reasoning	high
Defence decision-support	Defence	Core + DeepVision	high
Creative / IntuiTV SaaS	Creative	Orchestration	medium–high
Robotics	Robotics	Cognition	medium

LINE OF BUSINESS	BOOKED AS	MOTHER DEPENDENCY	MOTHER-ATTRIBUTABLE SHARE (ASSUMPTION TO VALIDATE)
MSAI Cloud infrastructure	Infra	Demand driver	enabling

Conclusion for diligence: value MOTHER on the *platform it enables*, not the single line it is booked against. The KPMG MOTHER figure is the correct, conservative *direct-revenue* anchor; the investment thesis rests on the enabled-revenue multiplier across defence, creative, robotics and infrastructure — all running on one sovereign core.

13 · CYBER-SECURITY IMPLEMENTATION

Defence-in-depth across software and model deployment

Security is implemented as concrete, merged, tested controls across both the application software and the model-serving layer — not as policy statements.

Application & agent layer

- **Input defence:** multi-signal prompt-injection / jailbreak / secret-elicitation detection on user input and on all retrieved (RAG/web) content, which is treated as never-instructional.
- **Output defence:** secret & credential redaction (API keys, tokens, connection strings, private keys), Luhn-checked card redaction, PII handling, and system-prompt/scaffold leak detection on every output path.
- **Zero-trust policy:** deny-by-default capability grants per agent; surface→agent reachability; human-in-the-loop approval on all outbound send/post/draft actions; hard Strike-0 kinetic deny.
- **Edge:** centralised security headers + HSTS + CSP (report-only → enforcing), request-size backstop; the `Function()` dynamic-eval primitive removed in favour of a bounded parser.
- **Fail-closed auth:** the chat gate now refuses rather than degrading to permit-all when the auth backend is unconfigured in production.
- **Monitoring & recovery:** tamper-evident hash-chained audit; circuit breakers; rate-limits; kill-switch; quarantine.
- **Continuous assurance:** ever-learning internal red-team + autonomous self-repair; admin-gated posture & red-team endpoints usable as CI/deploy gates.

Model & infrastructure layer

- **Sovereignty:** models served on owned GB10 hardware; customer data reasoned over without egress to third-party model providers.
- **Secrets:** environment-only, never hardcoded; secret scopes never granted to agents; redaction if a secret ever reaches an output.
- **Sandboxing:** code-execution agents run in an isolated external sandbox, not in-process.
- **Adversarial training:** red-team corpora feed the model's own robustness training, complementing the runtime probe.
- **Supply chain / SDLC:** reviewed pull-requests, type-checked strict TypeScript, unit-tested security modules, protected CI/CD and infra files.

> OUTSTANDING TO REACH "FULLY HARDENED"
>
> Cross-instance persistence of audit & learned signatures, a scheduled red-team + posture dashboard, secrets-manager/KMS integration at DC, and an **independent external penetration test + threat-model sign-off**. These are scoped, not speculative.

14 · REGULATORY — UK & EU AI ACT

AI Act adherence posture

> NOT LEGAL ADVICE
>
> This section describes the *technical controls* that support compliance and MSAI's intended posture. It is not a legal opinion; formal classification and conformity assessment require qualified legal & notified-body review per deployment.

- **Risk-tiering by product, not blanket:** consumer/creative uses sit in the minimal/limited-risk band (transparency obligations); defence & biometric-surveillance uses are treated as the high-risk / specially-regulated tier and held at pilot pending accreditation and DPIA. (EU AI Act carve-outs for national-security/defence and the UK's sector-led framework are noted; scope confirmed per deployment.)
- **Transparency (limited-risk):** AI interactions are disclosed to users; AI-generated media is attributable — supporting the AI Act's transparency duties.
- **Human oversight (high-risk):** Strike = 0 and human-in-the-loop on every consequential action are architecturally enforced, directly supporting the human-oversight requirement.
- **Records & logging:** the tamper-evident audit chain provides the event logging high-risk systems must keep.

- **Robustness & security:** continuous red-teaming + adversarial training evidence the accuracy/robustness/cyber-security obligations (and GPAI systemic-risk expectations for capable models).
- **Technical documentation:** model architecture, training data provenance, eval methodology and this dossier form the basis of the technical file.
- **GPAI obligations:** as provider of a general-purpose model, MSAI maintains model documentation, training-data summaries and downstream-provider information; sovereign hosting simplifies data-governance evidence.

15 · DATA PROTECTION — GDPR

GDPR adherence posture

> NOT LEGAL ADVICE
>
> Technical & organisational measures supporting compliance are described; a DPO-led assessment and, for surveillance/defence, a formal DPIA are required and are treated as prerequisites, not options.

- **Data minimisation & purpose limitation:** retrieval-grounded answering scoped to the user's own data; per-user memory keyed to the user; secret/PII redaction on outputs.
- **Sovereignty & transfers:** UK/EU sovereign hosting on owned hardware materially reduces international-transfer exposure — no personal data sent to third-party model APIs.
- **Lawful basis & consent:** authenticated accounts with tiered entitlements; consent/contract basis for platform processing; web/RAG grounding gated per request.
- **Security of processing (Art. 32):** encryption in transit, secrets management, access control (RBAC), tamper-evident audit, and the harness's defence-in-depth.
- **Data-subject rights:** per-user memory and records are addressable for access/erasure; audit trail supports accountability (Art. 5(2)).
- **Special-category & DPIA:** OverWatch surveillance and biometric/vision processing are explicitly flagged as requiring a completed DPIA and, where applicable, appropriate-policy documents before live deployment — the reason those platforms are held at pilot.

16 · ROADMAP

Path to go-live — anchored to Scotland DC 2026

MILESTONE	WINDOW	UNLOCKS
Harness enforcing + external pen-test	2026 H1	“Fully hardened” security claim; regulated-customer readiness
MOTHER CORE v5 training complete + eval-passed	2026 H2	Model GA; frontier merge from v12_7b set
MSAI Scotland DC go-live	2026	Fleet compute; GPU for creative + defence; sovereign cloud GA
AI platform GA (B2C + B2B agents)	2026 H2, on DC	Commercial scale-out; SLAs
Creative / IntuiTV GA	2026 H2	Consumer + enterprise media revenue
Defence phased accreditation	2026 → 2027	Per-deployment operational use
Robotics sim-to-real + Guardian cert	2027+	Embodied deployment

---

Forward-looking statements (dates, enabled-revenue framing) are estimates against the 2026 Scotland DC milestone, not commitments. No financial actuals are asserted; the KPMG line-item value is a client input. Regulatory sections describe technical controls supporting compliance and are not legal advice. Security is defence-in-depth with continuous self-testing; no system is represented as un-hackable, and an external penetration test is the outstanding independent validation.