

GEM CORTEX AI

– PRODUCT AND COMMERCIAL SPECIFICATION

The core principle of GEM Cortex AI: You are not buying an AI model or a chat tool. You are buying a measurably better process, secure access to AI, and the ability to manage quality, costs and operational accountability.




GEM Cortex AI Executive Summary

GEM Cortex AI is a product family and implementation framework for organisations that want to integrate AI into real processes, data and systems while retaining control over identities, costs, quality and accountability. It does not start by asking which chatbot to buy. It starts with a task, an owner accountable for the outcome, a source of truth and a secure point at which AI hands work over to a person or an enterprise system.

- ✓ Integrates AI into existing applications, document flows, knowledge sources and workflows.
- ✓ Measures the accepted business outcome, critical errors, service levels and unit cost.
- ✓ Selects between a precise path without an LLM, an answer based on corporate sources, a general-purpose model and a controlled tool.
- ✓ Enables organisations to start with a narrowly scoped pilot and extend shared governance without pretending that every pilot automatically constitutes a product.
- ✓ Unifies corporate identity, projects, models, access and cost accountability.

Core Positioning

Governed AI for real-world operations. GEM Cortex AI connects models with enterprise data, identity and processes so that clients can use, measure and develop AI securely, without depending on a single user interface or use case.



1. Why GEM Cortex AI Is Being Developed


In most organisations, the main problem is not a lack of access to models. Models are available both as end-user services and through APIs. The problem arises when moving into production: internal data has access permissions and owners, applications need stable interfaces, finance teams want to attribute costs to outcomes, and security teams need to know where data is being sent and who can initiate an action.

- ✓ **Fragmented use:** Teams purchase different licences and APIs without a common owner. Cortex provides an inventory, identities, a provider/project registry and consistent rules.
- ✓ **A pilot without a process:** The demo responds, but it is unclear who approves the output and where it is recorded. Cortex provides workflows, checkpoints, integrations and acceptance.
- ✓ **Data without context:** The chatbot does not know corporate roles, current revisions or the source of truth. Cortex provides permission-aware Knowledge, metadata and citations.
- ✓ **AI costs without value:** The invoice shows tokens or licences, rather than completed tasks. Cortex provides AI FinOps, cost allocation and a value ledger.
- ✓ **Dependence on a single vendor:** The UI, data, prompts and workflows become locked into one service. Cortex provides adapters, exports and deliberately managed lock-in.
- ✓ **Concerns about autonomy:** The organisation fears that an agent may perform an unintended action. Cortex provides minimal scopes, previews, human approval, idempotency and audit.

1.1 | What Cortex Is – and What It Is Not

Cortex combines repeatable product capabilities, integration components and GEM services for design, implementation and operation. Each deployment is configured according to the use case. Some components can be delivered separately; they provide the greatest value when they share identities, policies, telemetry and an evaluation process.

- **It is not a single proprietary GEM foundation model.** Cortex can use multiple commercial or local models depending on the task and data profile.
- **It is not a replacement for all of the client's AI tools.** It can complement them with shared management or support processes that an off-the-shelf SaaS licence does not cover.
- **It is not a universal autonomous agent.** The level of autonomy is selected according to the impact and can remain strictly in proposal-only mode.
- **It does not automatically certify compliance.** It provides evidence and controls, but accountability and legal assessment remain with the relevant roles.



2. *Ideal Client and Opportunity Qualification*

The ideal Cortex client has more than a general ambition to use AI. It has a specific process or management problem, an owner accountable for the outcome, and a willingness to provide access to representative data and the current baseline.

- ✓ **Enterprises and regulated organisations:** *Secure Workspace, Knowledge, Document Intelligence, Control.*
- ✓ **Public administration and public service organisations:** *Classification of submissions, document workflows, Knowledge in a restricted environment.*
- ✓ **IT and shared services organisations:** *Access & Spend, Service Desk, Code Discovery, internal Knowledge.*
- ✓ **Document-intensive businesses:** *Document Intelligence, RAG with citations, Flow and integrations.*
- ✓ **Organisations beyond their first wave of AI pilots:** *Control, integrating successful use cases into a platform and discontinuing weak ones.*

2.1 | *Indicators of a Suitable First Use Case*

- *The task is recurring, has a recognisable input and output, and currently consumes a measurable amount of time or creates risk.*
- *There is a process owner and a qualified reviewer who can determine what constitutes a correct outcome.*
- *It is possible to start in read-only or draft-only mode, or within a limited scope, without irreversible consequences.*
- *A representative historical dataset is available and can be used legally and securely for evaluation.*
- *The output can be incorporated into an existing application or workflow instead of creating another isolated inbox.*


2.2 | *When to Stop or Reframe an Opportunity*

- *The client requires guaranteed one hundred per cent accuracy for generative output without a checkpoint.*
- *The process owner, source of truth and method for identifying a correct outcome are unknown.*
- *The sole objective is to try the technology, with no value hypothesis or decision following the pilot.*
- *An existing feature in an already licensed SaaS service meets the need adequately, securely and at a lower cost.*
- *The requested data or automated action lacks an approved purpose, approved access or an accountable owner.*

3. The GEM Cortex AI Product Family

The names below describe capability blocks that can be combined into solutions. Not every block needs to be a separately licensed product in every project. What matters more to the client is the problem the block solves, the systems it connects and the evidence it provides.

- ✓ **Cortex Core:** *Orchestration of intent, runtime path, model and tenant context. A single governed entry point instead of application code tightly coupled to one model.*
- ✓ **Cortex Access:** *SSO, roles, workload identity, provider accounts, projects and lifecycle. Fewer shared keys, orphaned access rights and instances of unclear ownership.*
- ✓ **Cortex Knowledge:** *Connectors, ingestion, ACL-aware retrieval, citations and index lifecycle. Answers based on approved sources with traceable supporting evidence.*
- ✓ **Cortex Flow:** *Workflows, tools, approvals, human tasks and process state. An AI draft safely becomes a step in an existing process.*
- ✓ **Cortex Control:** *Policies, an AI registry, telemetry, evaluations, budgets and FinOps. Shared governance of quality, risk and costs across use cases.*
- ✓ **Cortex Workspace:** *A user channel for approved chat, knowledge and work tasks. A usable corporate alternative to ungoverned AI use.*
- ✓ **Cortex Document Intelligence:** *OCR, classification, extraction, normalisation and a structured draft. Less transcription and faster document processes while retaining control.*



4. Offering Packages and Entry Products

The packages provide a clear starting point for the sales conversation. Each has a buying trigger, an initial concrete deliverable and deliberate boundaries. The final scope and price are established after qualifying the data, integrations, volume and operating model.

- ✓ **Access & Spend Baseline:** *The CIO/CISO/CFO lacks visibility into the owners of AI accounts, projects and costs. Initial deliverable: a read-only inventory, gap map and 90-day baseline. Transparency before automation.*
- ✓ **Secure AI Workspace:** *The company wants to offer employees approved AI with SSO. Initial deliverable: a workspace for one group and three tasks. Governed adoption and audit.*
- ✓ **Knowledge Assistant:** *People search for procedures and do not know which version is current. Initial deliverable: a curated knowledge domain, citations and ACL tests. Faster identification of the source.*
- ✓ **Document Intelligence:** *The team manually sorts or transcribes documents. Initial deliverable: one class, an output schema and a draft. Less routine manual work.*
- ✓ **RFP Analyzer:** *The bid/procurement team reviews extensive terms and conditions. Initial deliverable: a matrix of requirements, sources, owners and risks. A lower risk of oversight.*
- ✓ **Code Discovery & Modernization:** *The legacy system has hidden dependencies. Initial deliverable: a read-only module map and test backlog. Reduces uncertainty before migration.*

5. *When Cortex Complements or Replaces Other Approaches*

Cortex is not an argument against off-the-shelf AI services. An appropriate design can use their models, interfaces and licences. The value of Cortex increases where the client needs to connect its own process, multiple providers, a shared identity, measurable quality or cost control.

- ✓ **SaaS assistant / copilot:** Strong for standard functionality. Cortex makes sense for custom workflows, provider-neutral control and shared FinOps. We recommend Cortex to complement governance, rather than as a replacement without justification.
- ✓ **Direct model API:** Strong for a single team and a bounded application. Cortex makes sense for multiple applications that require a gateway and a control layer.
- ✓ **Custom chatbot:** Strong for a simple communication channel. Cortex makes sense where ACLs, citations, integrations and operational ownership are required.
- ✓ **RPA/workflow:** Strong for deterministic steps. Cortex makes sense for unstructured inputs requiring language comprehension.
- ✓ **Single-vendor agent platform:** Strong for rich functionality in a specific cloud. Cortex makes sense for hybrid deployments.

A fair recommendation: If a standard feature in an already licensed service meets the need, data handling requirements and economic requirements, there is no need to build a custom Cortex use case. GEM can assist with selection, integration and governance.

6. *Value for Each Buying Role*

The same platform has a different significance for each role. Sales communication should start with the problem that the role faces and only then explain the technology.


- ✓ **CEO / executive management:** AI investment becomes a governed portfolio of use cases with owners and measurable outcomes.
Question: Which three processes offer the greatest value?
- ✓ **CIO / CTO:** Shared integration and operational capabilities instead of a series of isolated prototypes.
Question: What can be reused from the first use case?
- ✓ **CISO / DPO:** Identities, data profile, scopes, audit and controlled actions.
Question: Which data and tools are permitted for the given purpose?
- ✓ **CFO / procurement:** Cost by owner and outcome, with a comparison of licences, APIs and operations.
Question: Do the savings represent capacity, cash or reduced risk?
- ✓ **Process owner:** Faster task completion and less rework without loss of decision-making authority.
Question: Where should AI prepare a draft, and where is human confirmation required?

7. Value Economics and TCO

The Cortex economic model separates technical consumption, process impact and realised value. The calculation starts with a baseline and ends with how the organisation actually uses changes in capacity, throughput, quality or risk profile.

Working equation: Realised annual value = value captured from freed capacity + contribution from increased throughput + eliminated external cost + expected reduction in loss – recurring TCO.

- ✓ **Baseline:** Volume, time per case, error rate, rework, external cost and current SLAs. Without a baseline, the effect of AI cannot be distinguished from ordinary process change.
- ✓ **Adoption and eligibility:** How many cases are suitable for the solution and how many users actually follow the procedure. Laboratory accuracy does not automatically translate to the full volume.
- ✓ **Acceptance of the output:** The proportion of drafts that pass without substantive correction. A better measure than the number of generated responses.
- ✓ **Capacity realisation:** Redirecting time towards higher volume or reducing overtime. A minute saved is not, in itself, a cash saving.
- ✓ **Risk:** The probability and impact of an error before and after controls. Do not present this as a guaranteed saving; state the scenario and sensitivity.



Unit metrics: Cost per accepted draft, rather than per token; cost per document with correctly extracted data; cost per active user who completed a task; the difference between internal telemetry and the invoice.

8. Commercial, Implementation and Operating Model

The commercial model follows the reduction of uncertainty. The client is not buying an undefined transformation; it is buying a decision supported by data.

- ✓ **Qualification workshop:** Problem, sponsor, data, risk, alternative. Output: Opportunity brief.
- ✓ **Discovery & Pilot Design:** Baseline, solution outline, evaluation, security, TCO. Output: Pilot Charter.
- ✓ **Pilot:** Limited implementation, historical data, draft/read-only. Output: Go / adjust / stop report.
- ✓ **Limited production:** Real identities and volume, integrations, support. Output: Production readiness.
- ✓ **Scale & Managed Service:** Additional teams/use cases, capacity, FinOps, SLA.

Pricing components: One-off implementation (discovery, design, testing), platform service (Core/Access/Control), consumption (model/API, storage, compute), operations and support, ongoing quality and governance, change requests.

9. Access & Spend – A Standalone Entry Offering

A suitable starting point even for a client that is not yet ready to build its own AI application. It creates a shared data foundation for the CIO, CISO, finance and procurement. It reveals which costs relate to user licences and which to API consumption, who owns the projects, and where the link to a business outcome is missing.


- ✓ **Baseline:** Read-only collection of directory, provider project, usage and billing data. An inventory of owners, access rights and unallocated costs.
- ✓ **Control Pilot:** JML process, recertification, tags, budgets and showback within a limited scope. Initial governed operations.
- ✓ **Governed Scale:** Multiple providers/teams, model rules, chargeback and a value ledger.

Realistic boundaries: Cortex does not consolidate personal passwords or circumvent licensing terms. Automated account creation/removal is possible only where the provider offers supported API/SCIM capabilities. A hard budget applies only to consumption routed through the Cortex gateway.

10. Message House and Sales Narrative

Core message: GEM Cortex AI – governed AI for real-world operations. We turn a model into a securely integrated service with corporate identity, authorised sources, measurable quality, cost control and a clear point at which a human makes the decision.

- ✓ **Pillar 1 – Integration:** AI enters the existing workflow, API and data layer. The output becomes a controlled step in the process.
- ✓ **Pillar 2 – Control:** Identities, policies, citations, approvals, audit and restricted tools. The organisation knows who used what.
- ✓ **Pillar 3 – Choice:** Direct/no-LLM, RAG, a general-purpose model, a local runtime or a tool, depending on the task.
- ✓ **Pillar 4 – Economics:** Consumption linked to a use case and an accepted outcome. The CIO and finance decide on scaling.



10.1 | Thirty-Second Pitch

GEM Cortex AI helps businesses and public organisations move AI from isolated experiments into governed operations. It connects models with corporate identity, authorised knowledge and existing processes. Depending on the task, it selects a precise path without an LLM, an answer with citations or a controlled workflow step. At the same time, it measures quality, cost and acceptance of the output. We start with a narrowly scoped pilot, a clear baseline and a decision to go, adjust or stop.



10.2 | Two-Minute Pitch

Today, access to a model alone is not a competitive advantage. The challenge is to ensure that AI is permitted to see only the appropriate data, uses a suitable tool, delivers an output that enters a real process, and enables the organisation to demonstrate both quality and economic performance. GEM Cortex AI addresses this gap with a shared layer of Core, Access, Knowledge, Flow and Control. The client does not have to rebuild identities, audit, model adapters and cost measurement for every use case. In the first project, we deliver only the components needed for the specific outcome; successful components can be reused in other domains. If the hypothesis is not confirmed in the pilot, the project has a predefined stopping point.

11. Qualification Questions and Objection Handling


11.1 | Common Objections

- **“We already have Microsoft Copilot or another enterprise assistant.”** That may be the right choice for standard tasks. We will assess the gaps: custom processes, additional models, permission-aware knowledge and shared FinOps. If there is no gap, we will not recommend Cortex.
- **“All we need is to call the model’s API.”** For a single bounded prototype, yes. Cortex becomes cost-effective when you need to share identities, evaluation, audit and budgets across multiple applications.
- **“We want the single best model.”** The best model depends on the language, task and data handling requirements. Registries and evaluations make it possible to select a default model and assess a change safely.
- **“We need 100% accuracy.”** Generative output cannot provide such a general guarantee. Critical fields can be constrained by a schema and validation; decisions involving risk remain with a human.
- **“We require zero data retention.”** We will assess the specific product, region and contractual arrangements. ZDR cannot be inferred from the provider’s name.
- **“Everything must run on-premises.”** This is possible with suitable models, but it requires hardware, MLOps and licence verification. We will compare both hybrid and local TCO.
- **“MCP will solve the integrations.”** MCP can standardise the contract for resources and tools, but it does not replace data mapping, permissions or security policies.
- **“The pilot worked, so why is it not in production immediately?”** A pilot tests a hypothesis. Production requires real identities, SLOs, incident handling, recovery and contractual accountability.
- **“How do you guarantee ROI?”** The measurement process and delivered capabilities can be guaranteed, but the organisation’s future use of capacity cannot.

12. *Procurement and Evidence Pack*

An enterprise client does not buy functionality alone. It needs evidence that the solution can be securely taken over, operated, controlled and, if necessary, exited. Presales prepares a structured evidence pack proportionate to the stage of the project.

- ✓ **Product and scope:** Use case, in/out of scope, channels, data, models, tools.
- ✓ **Architecture:** Logical and deployment diagrams, runtime paths.
- ✓ **Data and privacy:** Data flow, classification, region, retention, deletion.
- ✓ **Security:** Threat model, identities, testing and incident interface.
- ✓ **Quality:** Evaluation dataset, metrics, critical errors and release gate.
- ✓ **Operations:** SLO/SLA, capacity, RACI, backup/restore.
- ✓ **Economics:** Implementation, platform, consumption, value ledger.
- ✓ **Exit:** Exports, configuration handover, known lock-in.



13. *How to Initiate a Qualified Sales Dialogue*

The most effective first step is not a general presentation, but a short qualification conversation focused on one specific process or an overview of AI costs. GEM works with the client to identify the current state, define a safe boundary for the pilot and prepare evidence for the decision.

- 1 | **We identify the decision:** One task or management problem.
- 2 | **We measure the baseline:** Volume, time, quality, rework.
- 3 | **We design the first pilot:** Limited scope, evaluation, data, checkpoints.
- 4 | **We substantiate go/adjust/stop:** Outcome against the baseline, security and unit cost.

Recommended call to action: Let us start with one process or a 90-day overview of AI access and costs. We enter discovery with a specific hypothesis and leave with the evidence needed for a decision.