AITHENTIC
Authentic AI. Augmented Outcomes.
A venture of HOFT Group

HOFT&
HUB OF FINANCE TRANSFORMATION

ENTERPRISE-GRADE AGENTIC AI · TRUST BY DESIGN

The Authentic Intelligence Layer

AI and Data Security · Governance

How AITHENTIC delivers Agentic AI that finance, risk and IT leaders can actually trust — by separating data from reasoning, hardening the runtime, and validating every output.

01	Data Governance	02	AI Governance	03	Data Security	04	Hallucination Control

WHY ENTERPRISES PAUSE ON AGENTIC AI

The trust problem facing CFOs, CIOs and Risk leaders

Most agentic AI projects stall not on technology — but on questions of trust. When AI touches financial data, regulated workflows or executive decisions, four risks block deployment.

Hallucinations on regulated data

LLMs invent numbers, sources or citations — unacceptable in finance, healthcare or compliance contexts.

No audit trail

Outputs cannot be traced back to source systems, KPIs or measures — making them indefensible to auditors and regulators.

Secrets sprawl & weak identity

API keys on the client, service accounts in code, no managed identity — security teams will not approve production.

LLM as calculator

Asking an LLM to fetch and compute KPIs guarantees inconsistency. Finance numbers must be deterministic — not generated.

AITHENTIC's answer: an Authentic Intelligence Layer that solves all four — by design, not as an afterthought.

THE AITHENTIC FRAMEWORK

Four pillars of Authentic Intelligence

Every AITHENTIC engagement is built on four interlocking pillars. Together they make Agentic AI auditable, trusted, hardened and explainable — from day one.

01

Data Governance

- Trusted sources only
- Fabric Gold tables · Power BI semantic models
- Deterministic data fetching
- KPI ownership & lineage

02

AI Governance

- LLM role explicitly bounded
- Structured Pydantic payloads in
- Validation layer on outputs
- Full traceability to source

03

Data Security

- Managed Identity, no stored secrets
- Auth at the Static Web App layer
- RBAC mapped to KPI scope
- Audit logging · CORS · egress lockdown

04

Hallucination Control

- Validated input payloads
- LLM guardrails in prompt & runtime
- Output validation vs source data
- Regenerate / flag on mismatch

The result: *AI insights that finance, audit, risk and IT can sign off on — together.*

FOUNDATIONAL GUARANTEE · DATA RESIDENCY

Trusted models. In your tenant. In your region.

AITHENTIC connects only to enterprise-hosted LLMs inside Azure AI Foundry. No public endpoints, no consumer LLM products, no data crossing borders.

What never happens: consumer ChatGPT calls · data sent to public LLM APIs · customer queries used as model training data

INSIDE YOUR AZURE TENANT · IN YOUR REGION

Microsoft Azure AI Foundry

Enterprise-hosted models, governed inside your subscription · runs in your Azure region

AZURE OPENAI	ANTHROPIC	GOOGLE	xAI	+ MORE
GPT-4o · GPT-5	Claude	Gemini	Grok	Llama · Mistral · Cohere

NO PUBLIC ENDPOINTS
No consumer LLM APIs in the data path

STAYS IN REGION
UAE / KSA / your chosen Azure region

TENANT-ISOLATED
Runs inside your Azure subscription

FOUNDATIONAL PRINCIPLE

The LLM is not your database — and never your calculator

AITHENTIC enforces a strict separation between trusted data retrieval and language-based reasoning. The LLM never decides what data to fetch, and never computes a KPI.

DETERMINISTIC

Data Layer

- Azure Functions · predefined endpoints
- Microsoft Fabric Gold tables
- Power BI semantic measures
- Azure AI Search (RAG · documents)

CONTRACT

Structured Payload

```
{
  "kpi": "EBITDA",
  "period": "Q1 2026",
  "value": 142.3,
  "source": "PBI:fct_pnl",
  "validated": true
}
```

Pydantic-validated · schema-bound · no free-text data

REASONING ONLY

LLM Layer

- Executive summaries
- Narrative generation
- Insight explanation
- Natural-language Q&A
- *Calculation / fetching*

WHY THIS MATTERS

The LLM is the language layer — never the source of truth. Financial numbers stay reliable. Audit trails stay intact. Hallucinations have nowhere to enter the calculation path.

PILLAR 01 · DATA GOVERNANCE

01

Trusted sources, controlled access, full lineage

AI insights are only as trustworthy as the data behind them. AITHENTIC binds every output to validated, lineage-tracked sources.

STRUCTURED DATA

Numbers, KPIs & metrics

SOURCES

- Microsoft Fabric — Gold tables (curated, conformed)
- Power BI semantic models — measures, KPIs
- ERP / EPM systems via certified connectors

USED FOR

Revenue · Cash flow · EBITDA · Working capital · any board-grade financial metric

UNSTRUCTURED DATA

Documents & context

SOURCES

- PDFs, PPTs, Word documents, policies, contracts
- Chunked & embedded into Azure AI Search
- RAG retrieval bound to user-permission scope

USED FOR

Narrative context · policy explanations · qualitative commentary supporting the numbers

Lineage everywhere: every value the AI cites can be traced to its exact source — table, measure, document and timestamp.

PILLAR 02 · AI GOVERNANCE

02

What the LLM may do — and what it never does

AITHENTIC governs the LLM the way a regulated firm governs any high-risk system: explicit role boundaries, controlled inputs, and validated outputs.

THE LLM IS USED FOR

- Executive summaries grounded in fetched payloads
- Narrative generation around verified KPIs
- Financial insight explanation (why not what)
- Natural-language Q&A over governed data

THE LLM IS NEVER USED AS

- A calculator for KPIs or financial measures
- A query engine deciding which data to fetch
- A source of facts (numbers, dates, citations)
- A direct write-path into systems of record

Schema-bound input

Pydantic-validated payloads — the LLM only sees a fully-typed contract

Prompt-level guardrails

Role, scope and refusal policies enforced in every system prompt

Output validation

Generated narrative checked against source data before reaching the user

PILLAR 03 · DATA SECURITY

03

Eight-point hardening — security teams sign off

Identity, secrets, network and audit are non-negotiable. AITHENTIC follows a single, opinionated security baseline on every engagement.

01

Managed Identity

Identity flows via Managed Identity — never stored secrets

02

API key off the client

Master keys stay server-side; clients never hold privileged keys

03

Auth at the edge

Users authenticated at the Static Web App layer before any call

04

RBAC = KPI scope

Role-based access maps to KPI scope — not just URL access

05

Key Vault only

Key Vault is the single, exclusive home for every secret

06

Audit logging

Every Function App call is logged for compliance and forensics

07

Network lockdown

Egress controls and CORS lockdown enforced at the platform layer

08

LLM guardrails

Runtime guardrails on prompts, completions and tool invocations

Aligned with Azure Well-Architected Framework · ISO 27001-ready · KSA & UAE data-residency by default

PILLAR 04 · HALLUCINATION CONTROL

04

Three layers of defense before the user sees anything

AITHENTIC treats hallucinations as a system-level risk — not a model-level quirk. The defense is engineered before, during and after generation.

L 1

Pre-LLM

Validated input

- Deterministic data fetch
- Pydantic schema enforcement
- Missing-field detection
- Type and range checks

L 2

In-LLM

Bounded generation

- Role-locked system prompt
- Refusal policies enforced
- Grounded-only response mode
- Tool-use guardrails

L 3

Post-LLM

Output validation

- Cross-check vs source payload
- Numeric consistency tests
- Citation presence check
- Regenerate / flag on mismatch

OUTCOME

What reaches the executive is grounded, validated and source-traceable — or it does not reach the executive at all.

TRACEABILITY · ENHANCED FLOW

From narrative to source — every insight, defensible

Beyond the baseline, AITHENTIC adds a validation step and a traceability surface so users can see exactly where each AI insight comes from.

BASELINE FLOW

Current

1. Azure Functions fetch from Fabric / Power BI
2. Build structured JSON payload
3. Send payload to LLM
4. LLM generates narration
5. Show response to user

AITHENTIC ENHANCED FLOW

Enhanced

1. Azure Functions fetch from Fabric / Power BI
2. Build structured JSON payload
3. Send payload to LLM
4. LLM generates narration
5. **VALIDATION LAYER · consistency vs source**
6. **If valid → show response | If not → regenerate / flag**

EXAMPLE · *"EBITDA decreased due to increased costs"* → **EBITDA: PBI measure · Costs: Gold table · Period: Q1 2026**

PLATFORM ALIGNMENT

Built end-to-end on the Microsoft trust boundary

Staying inside one cloud provider's trust boundary simplifies compliance, identity and data residency. AITHENTIC's reference architecture is Azure-native by default.

Microsoft Fabric
Lakehouse / warehouse · Gold tables

Power BI
Semantic models · KPIs · measures

Azure OpenAI
Hosted LLMs with enterprise SLA

Azure AI Search
RAG vector store for documents

Azure Functions
Deterministic data-fetch endpoints

Azure Key Vault
Single home for every secret

Azure Static Web Apps
Edge auth · clean front-end delivery

Managed Identity + RBAC
Identity flows · scoped access

COMPLIANCE
Single trust boundary · audit-ready

RESIDENCY
KSA / UAE region pinning supported

TOTAL COST
Existing E5 / EA estate is reused

WHAT ‘AUTHENTIC’ MEANS

Not just intelligent.

Authentically trustworthy.

Auditable

Every output traces to source data, owner and timestamp.

Trusted

Numbers come from systems of record — never from the model.

Hardened

Identity, secrets, network and audit pass enterprise security review.

Explainable

Users see why an insight exists — not just what it says.

Authentic AI. Augmented Outcomes.

AITHENTIC

LET'S TALK

Build an AI your auditors can sign.

AITHENTIC delivers Agentic AI as outcomes — packaged, governed and production-grade. From discovery to live in weeks, with the trust posture finance, risk and IT can defend.

HOFT&
HUB OF FINANCE TRANSFORMATION

WEB aithentic.co · hoft-global.co

EMAIL zeeshan@hoft-global.co

PHONE +966 540 52 3612, +971 56 264 6830

OFFICES Saudi Arabia · United Arab Emirates · United Kingdom · Pakistan · India · Egypt

Authentic AI. Augmented Outcomes.