



NVIDIA INCEPTION GRAND CHALLENGE 2026

# Ontology-based Agentic AI Platform for Sovereign Enterprises

The ontology-based platform for building, governing and scaling AI agents on the data your business runs on, without moving it.

**Deployment status:** in production pilots with Australian Government agencies, including a paid pilot and an agency POC progressing to pilot



**COMPANY**  
Synoptix AI (IHA Consulting Pty Ltd)

**COUNTRY**  
Canberra, ACT, Australia

**FOUNDER & CEO**  
Hanan Asghar (ex-Deloitte, ex-Australian Government)

**FOUNDED**  
2024

**STAGE / FUNDING**  
Bootstrapped

**TEAM**  
20 FTE

**TARGET SEGMENT**  
Government, national security & regulated industries (banking, telecommunications and mining)



# Critical knowledge work is still manual because enterprise AI cannot be trusted with high-stakes decisions

Government agencies, regulated enterprises and critical infrastructure organisations rely on experts to assess cases, interpret policy, investigate evidence and make operational decisions. These workflows remain slow, expensive and difficult to scale because existing AI cannot operate with enterprise-grade trust, governance or context.



## Why knowledge work remains manual

Mission-critical decisions still depend on specialists reviewing hundreds of documents, policies and systems manually because the work requires judgement, traceability and compliance.



## Why current AI solutions are insufficient

Most AI solutions retrieve information but do not understand business entities, relationships or policy. They generate plausible answers yet cannot consistently produce trusted, explainable and decision-grade outcomes.



## The cost of manual decision-making

Knowledge-intensive processes remain slow and expensive because every assessment depends on expert review. Organisations struggle to improve throughput, reduce costs or scale operations without increasing headcount.

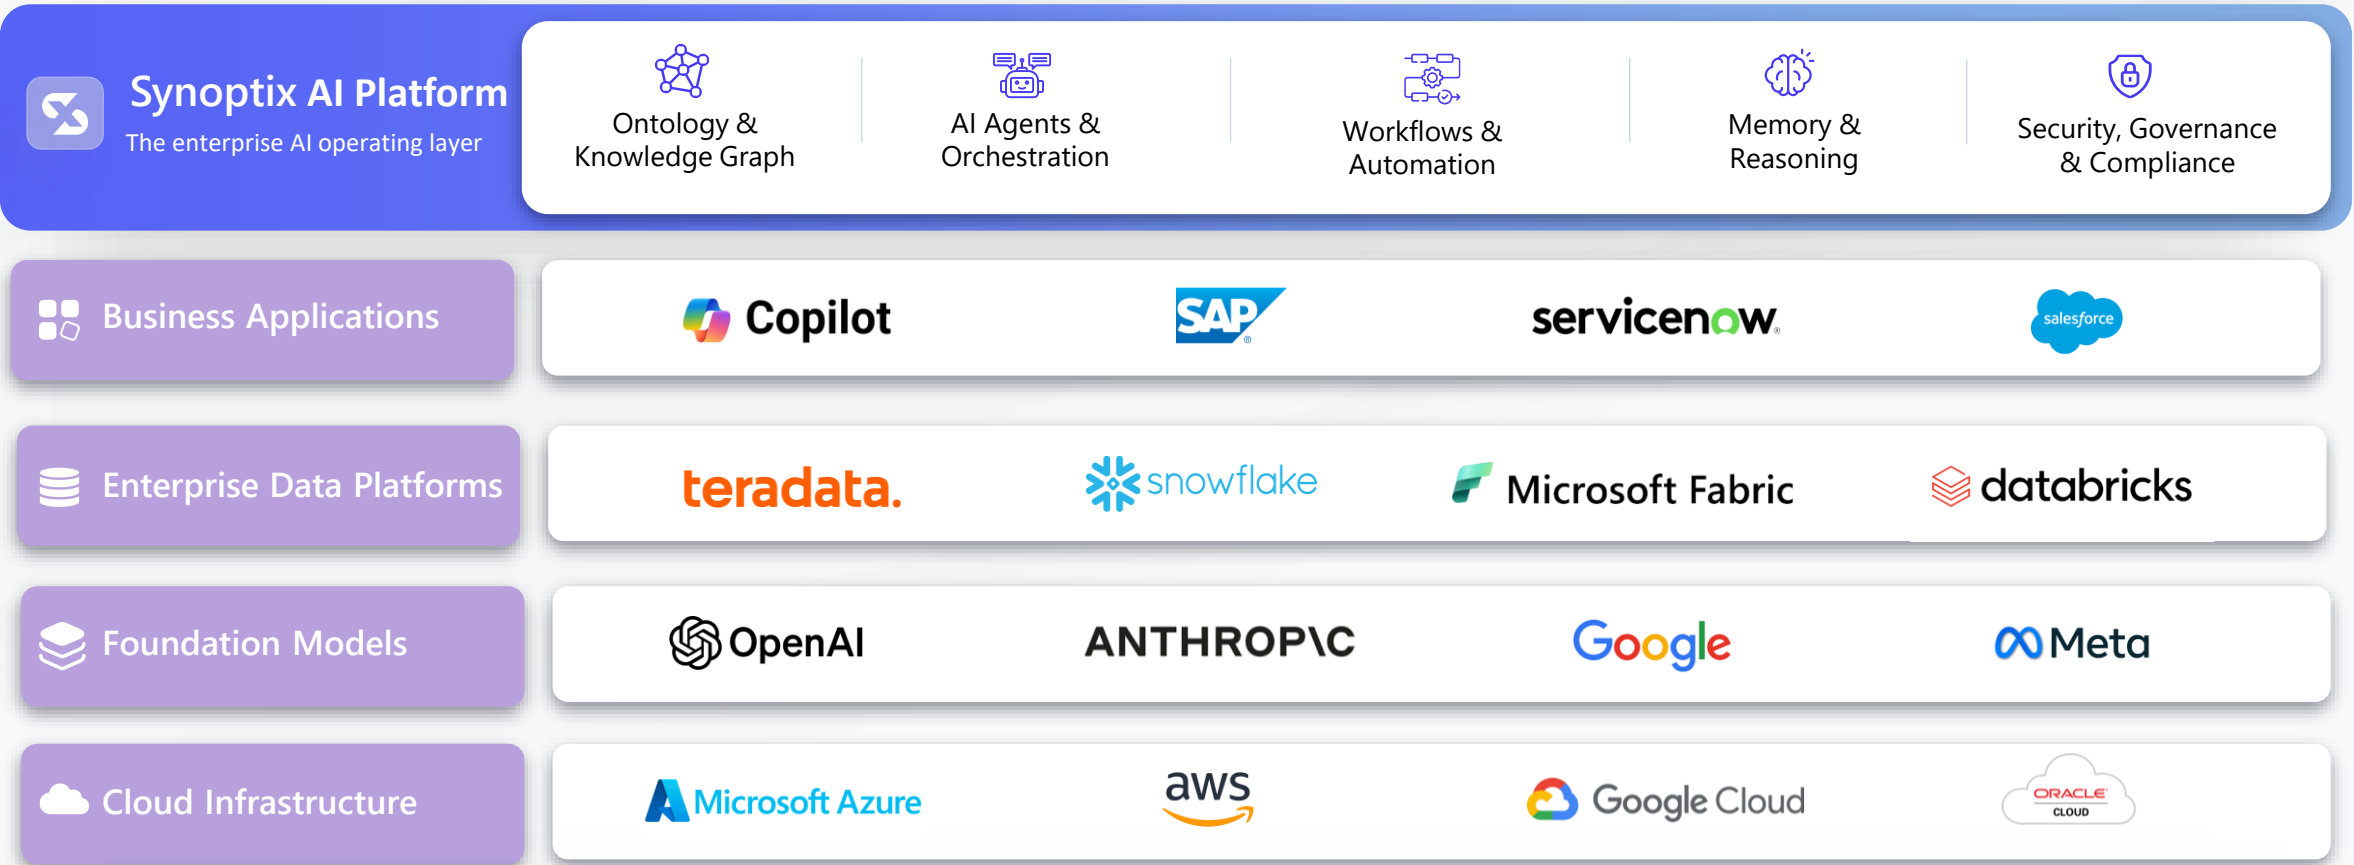


## Why governance, monitoring and security matter

High-stakes decisions require AI that enforces policy, protects sensitive information, provides complete auditability and explains every recommendation. Without governance, AI cannot be safely deployed across mission-critical operations.



# Where Synoptix AI Fits in the Enterprise Stack



Synoptix is the enterprise AI orchestration layer that sits above your existing investments, connecting infrastructure, models, data and applications into governed AI workflows that deliver real business impact.

**We Integrate. We Orchestrate. You Innovate.**



Integrate Everything



Orchestrate Intelligently



Innovate Confidently

Synoptix AI is compatible with all major cloud, on-premises and AI infrastructure

# A semantic layer over your data, a governance layer under your AI

## WHAT IS SYNOPTIX AI

One sovereign platform to Build, Deploy, Manage & Monitor AI agents. An ontology-based semantic layer models the business as entities and relationships, while a governance layer (system-level security, observability and monitoring) sits under every agent

## HOW IT WORKS

Deploys inside the customer's own tenancy with zero data movement. 100+ connectors reach structured DBs, documents, images, email and operational systems (ERP, CRM, ServiceNow, EDRMS) in place; agents, analytics and APIs consume one governed semantic layer

## WHAT IS AI-DRIVEN

Agent reasoning and multi-step workflows, ontology-grounded context assembly, sovereign ASR and vision embeddings, and in-path AI Firewall classification, all served on NVIDIA GPU-accelerated inference via Microsoft Azure

## WHAT IS BUILT AND USABLE TODAY

Ontology & Taxonomy Workbenches, no-code Agent Builder, a security, monitoring and observability console, domain taxonomies, live in Australian Government environments



# Two live engagements with Australian Government agencies

## CASE A: ANONYMISED (GOVERNMENT)



### National credentialing authority

*AI marking of spoken language exams*

#### INDUSTRY / USE CASE

Government credentialing: Automated exam assessment

#### DEPLOYMENT TYPE

Paid pilot progressing to production

#### TIMELINE / STATUS

Engaged May 2026. Architecture delivered; pilot marking in progress

#### SCOPE

33,000+ exams p.a. across 55+ language pairs. Sovereign ASR + multi-tier language models.

## CASE B: ANONYMISED (GOVERNMENT)



### Federal border security agency

*Image similarity search for illicit tobacco detection*

#### INDUSTRY / USE CASE

Border protection: Illicit goods intelligence

#### DEPLOYMENT TYPE

POC delivered; structured 8-week pilot plan agreed

#### TIMELINE / STATUS

Engaged February 2026. Agent operational in agency environment

#### SCOPE

GPU-accelerated vision embeddings over seizure images; brand/batch matching linking seizures to supply networks

# Proof of performance

## Case A: AI exam marking

### MARKING TURNAROUND

3 to 4 weeks → **<8 hours**

### COST PER EXAM

A\$200 → ↓ **~95%**

### MARKER CONSISTENCY (AGREEMENT)

80% → **>95%**

### ESTIMATED HOURS SAVED

**1,206,000 hrs p.a.**

*Measured against human double-marking baseline across pilot language pairs.*

## Case B: Illicit tobacco image search

### IMAGE TRIAGE THROUGHPUT

100/day manual → **10,000/day**

### TIME TO CANDIDATE MATCH

Hours → **seconds**

### MATCH PRECISION @ TOP-10

n/a (manual) → **>90%**

### ANALYST HOURS PER CASE

10 hours → ↓ **~60%**

*Measured on agency seizure-image corpus during POC; pilot will validate at scale.*

# Proof points judges can verify



## Customer Use Case: National credentialing authority (Australia)

Architecture completed and accepted. Synoptix AI selected to design the AI-powered interpretation assessment platform. Project progressing toward implementation under NDA.



## Pilot & Project Summary

Completed discovery, ontology design, AI evaluation framework and sovereign deployment architecture with customer stakeholders. Deliverables formally reviewed and accepted.



## Deployment-ready architecture available through Azure Marketplace

Synoptix AI is packaged for enterprise deployment via Azure Marketplace, enabling secure, scalable implementation within customer environments.



## Customer Validation

Customer has commissioned multiple work packages following initial engagement, validating technical capability and business value. Additional customer correspondence available under NDA.



## Partner Validation

Formal technology partnership established. Joint GTM activities, Microsoft Partner Center co-sell, customer introductions and Partner Forum presentation scheduled.



# NVIDIA stack in the Synoptix AI architecture



## Infrastructure

- NVIDIA GPU-accelerated inference (A100 / H100) via Azure AI Foundry, in Azure Australia sovereign regions
- All model serving (ASR, LLM and vision embedding workloads) runs on NVIDIA GPUs
- GPU-accelerated vector similarity search powering the illicit tobacco image-matching agent



## Frameworks & AI Models

- NVIDIA NIM microservices consumed through Azure AI Foundry model catalogue
- NVIDIA Nemotron model family under evaluation for agent reasoning tiers
- TensorRT-LLM-optimised serving for latency-sensitive marking pipelines; CUDA-accelerated embedding generation



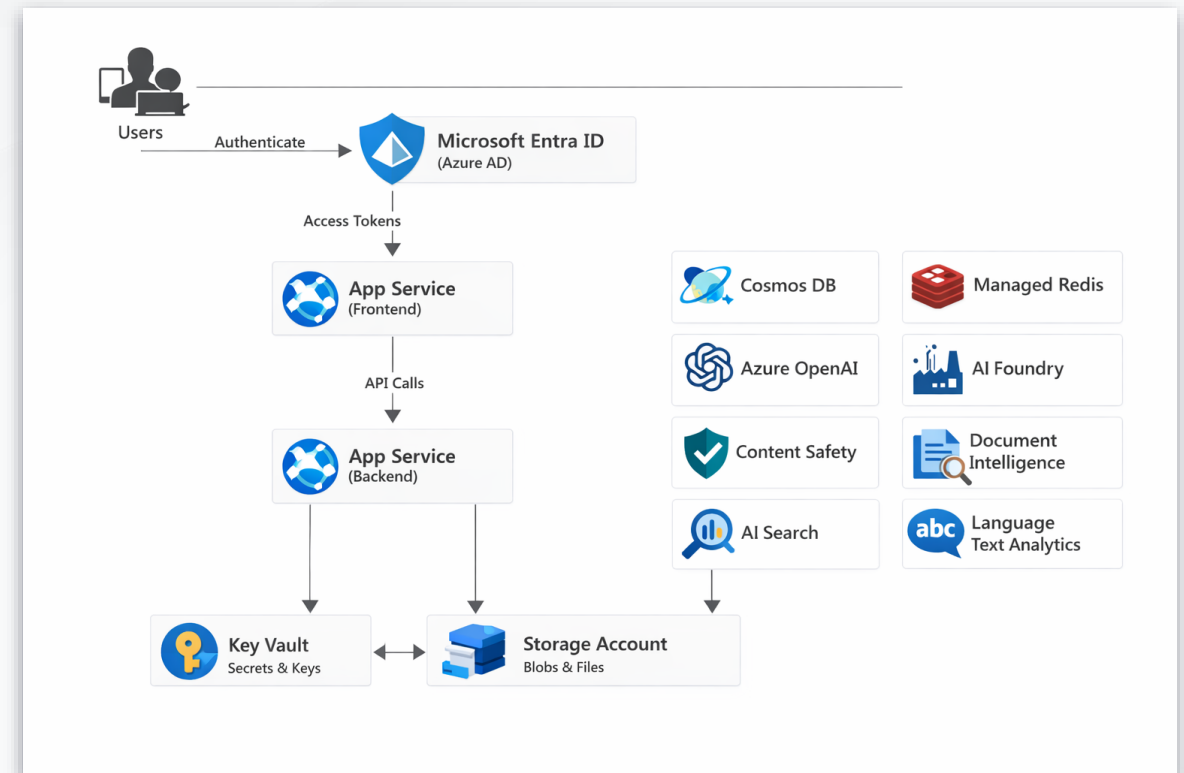
## Adoption Roadmap

- On-premises and air-gapped NVIDIA GPU deployment for PROTECTED-level government workloads
- NVIDIA Cloud Partner (NCP) infrastructure for sovereign hosting across Singapore and Southeast Asia
- Deepen NIM and Nemotron adoption as the default inference layer across the platform

**Why NVIDIA matters here:** sovereign agentic AI is GPU intensive; every exam marked and every image matched drives NVIDIA-accelerated inference. Scaling to national workloads scales GPU, NIM and NCP consumption directly.

# Deployment Architecture & Enterprise Integration

- ✓ **Zero Data Movement**  
Customer data never leaves the enterprise boundary. AI models, agents and orchestration are deployed directly into the customer's cloud, sovereign environment or on-premises infrastructure.
- ✓ **Flexible Deployment**  
Deploys into Azure customer tenancy, sovereign cloud, hybrid environments and private data centres, with edge deployment on the roadmap.
- ✓ **Enterprise Integration**  
Native REST APIs, event-driven workflows and connectors integrate with enterprise systems including Microsoft services, databases and existing business applications.
- ✓ **Security & Compliance**  
IRAP-aligned deployment with role-based access control, encryption, audit logging, AI Firewall and controls for PII, data residency and classified information.
- ✓ **Enterprise Onboarding**  
Typical deployment completed in **4 to 6 weeks**, including tenancy setup, identity integration, ontology import, connector configuration, agent setup and user acceptance testing.



# Commercial Readiness & Go-to-Market



## Target Customers

Australian Government agencies, regulated enterprises and critical infrastructure organisations requiring sovereign, explainable and secure AI. Primary buyers include CIOs, CTOs, Chief Data Officers and AI transformation leaders.



## Sales Pipeline

Active pipeline across Australian Government, national science agencies, critical infrastructure, financial services and regulated enterprises. Opportunities include paid pilots, enterprise platform deployments and multi-phase AI transformation engagements.



## Commercial Traction

Available on Azure Marketplace with paid enterprise pilots underway. Formal Teradata partnership established for joint go-to-market, supported by Microsoft Partner Center co-sell and multiple enterprise engagements progressing toward production deployment.



## Commercial Model

Enterprise SaaS licensing with implementation and ontology design services. Flexible deployment through Azure Marketplace or customer-managed sovereign, hybrid and on-premises environments.



## Partner-led Go-to-Market

Joint go-to-market strategy with Teradata, Microsoft and Unisys, supported by Microsoft Partner Center co-sell, AWS, CSIRO, specialist consulting partners, university research collaborations, and industry bodies including the Tech Council of Australia and the AIIA, to accelerate enterprise adoption.



## APAC Expansion Strategy

Platform designed for sovereign AI requirements across Singapore and the broader APAC region. Partner-led expansion through hyperscalers, global system integrators and regional enterprise channels.