Understanding Agentic AI: Hype to Tangible Progress

Jason Campagna
Senior Director, AI Strategy & GTM

World Wide Technology

Jason Campagna

Global Solutions & Architecture

Sr. Director, AI Practice Strategy

Wildwood, MO

BACKGROUND

- Joined WWT in 2009
- 20+ years in Next-Generation Technology including Virtualization, VDI, Automation, DevOps & Agile Infrastructure, Cloud, Portfolio Strategy, and Artificial Intelligence

EXPERIENCE

- AI Practice Strategy (WWT 2023-Present)
- Portfolio Strategy & Operations (WWT 2019 – 2023)
- Cloud & Automation Practice (WWT 2014-2019)
- Virtualization Practice (WWT 2009-2014)
- Server Centric Consulting – Acquired by WWT in 2009, focused on enterprise scale Virtualization, VDI, Citrix at Enterprise Scale

FUN FACTS

- Proud Dad of 3 – Jackson 9, Christian 6, and Eve 3.
- Avid boating family – can find us often in a cove @ LOTO
- Lived in the Florida Keys for 8 years.

AI Hype: Buzzword Bingo

Vibe Coding
Hallucinations
AI-Driven
Digital Twin
Cognitive Architecture
AI Foundry
Digital Employees
RAG
REG
Synthetic Cognition
AI Factory
AGI
Human-in-the-loop
Prompt Engineering
The AI Revolution!
Game Changing
AI-Native Software
Model Context Protocols (MCP)
Dark Data Activation
Object-Driven Interfaces (ODI)
Quantum AI
Bias-Free
Agentic AI
Hyper Automation
Graph Databases
Autonomous Agents
Agent-Washing
LLMOps
PromptOps
AI Agents
AI Studio
Continuous Learning Loops
Multi-Agent Systems
Ambient AI
Digital Twin
Explainable AI
Everything AI Powered (even refrigerators)
Responsible AI
Synthetic Data
Ethical AI
Swarm Intelligence
AI Proving Ground
Zero-Shot Learning
AI-Driven

AI Capabilities are Evolving **Rapidly**

Test scores of AI systems on various capabilities relative to human performance

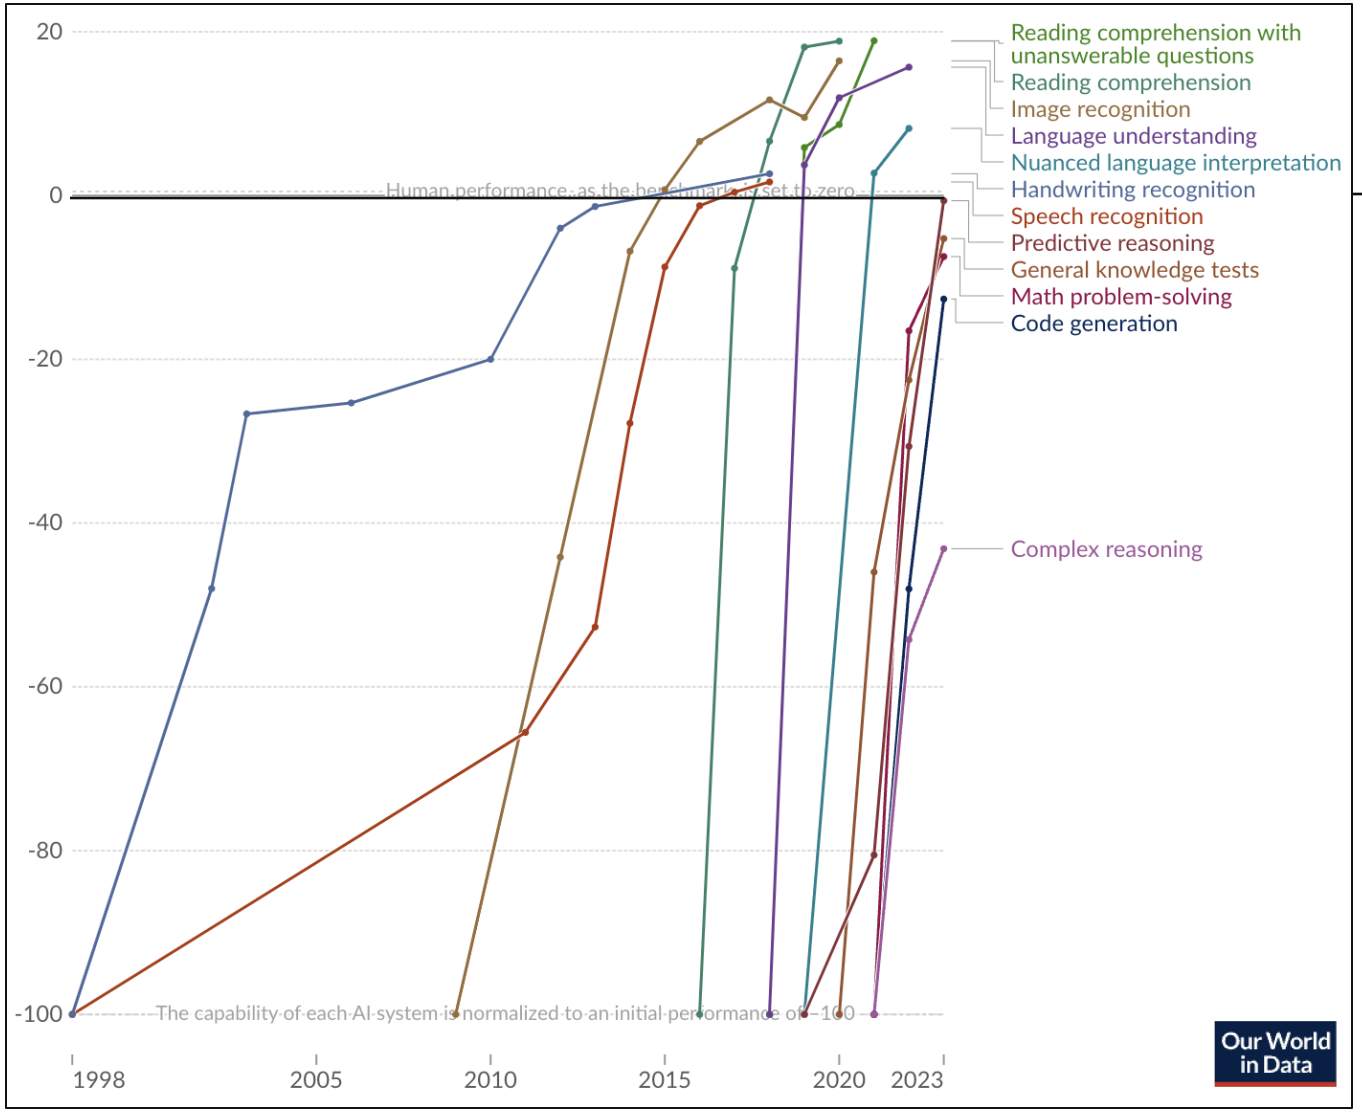

https://ourworldindata.org/artificial-intelligence

For each domain, initial performance is set to -100.

Human performance is used as a baseline, set to zero.

When the AI's performance crosses the zero line, it scored more points than humans.

AI Capabilities are Rapidly Becoming Superintelligent

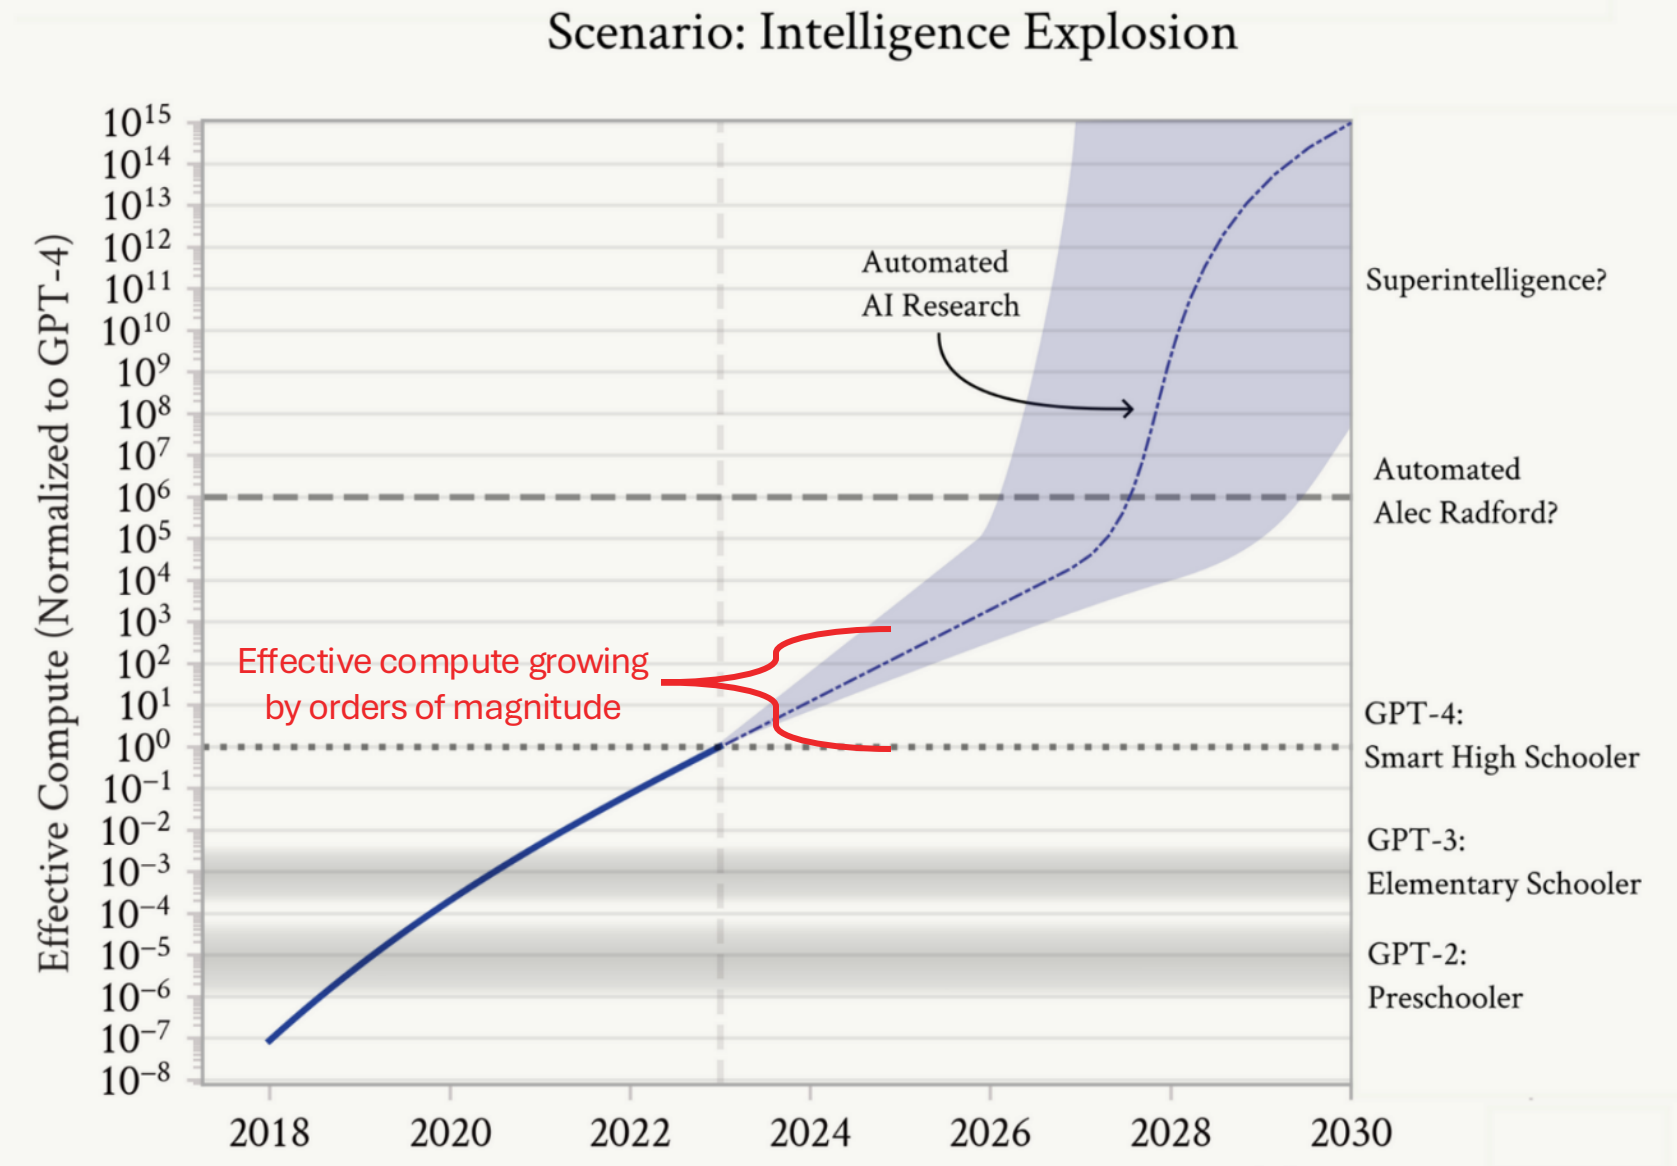

OpenAI research projections

We’re Already Evolving into **Agentic AI**

Human-driven
Simulates conversation, limited scope, inflexible context

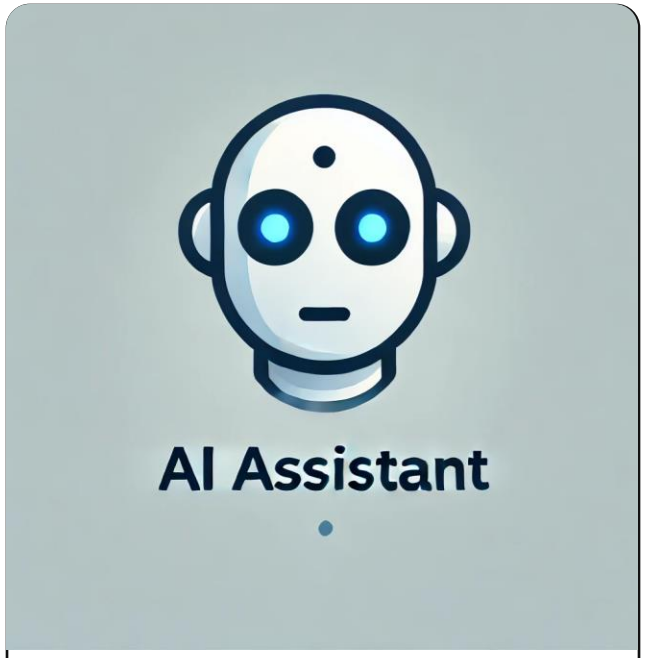

Human-Driven
Personalized conversation, understands context, integrates ecosystem

Machine-Driven
Goal driven, complex tasks, makes decisions, adapts dynamically

Today's **Agency Gap** – A Visualization

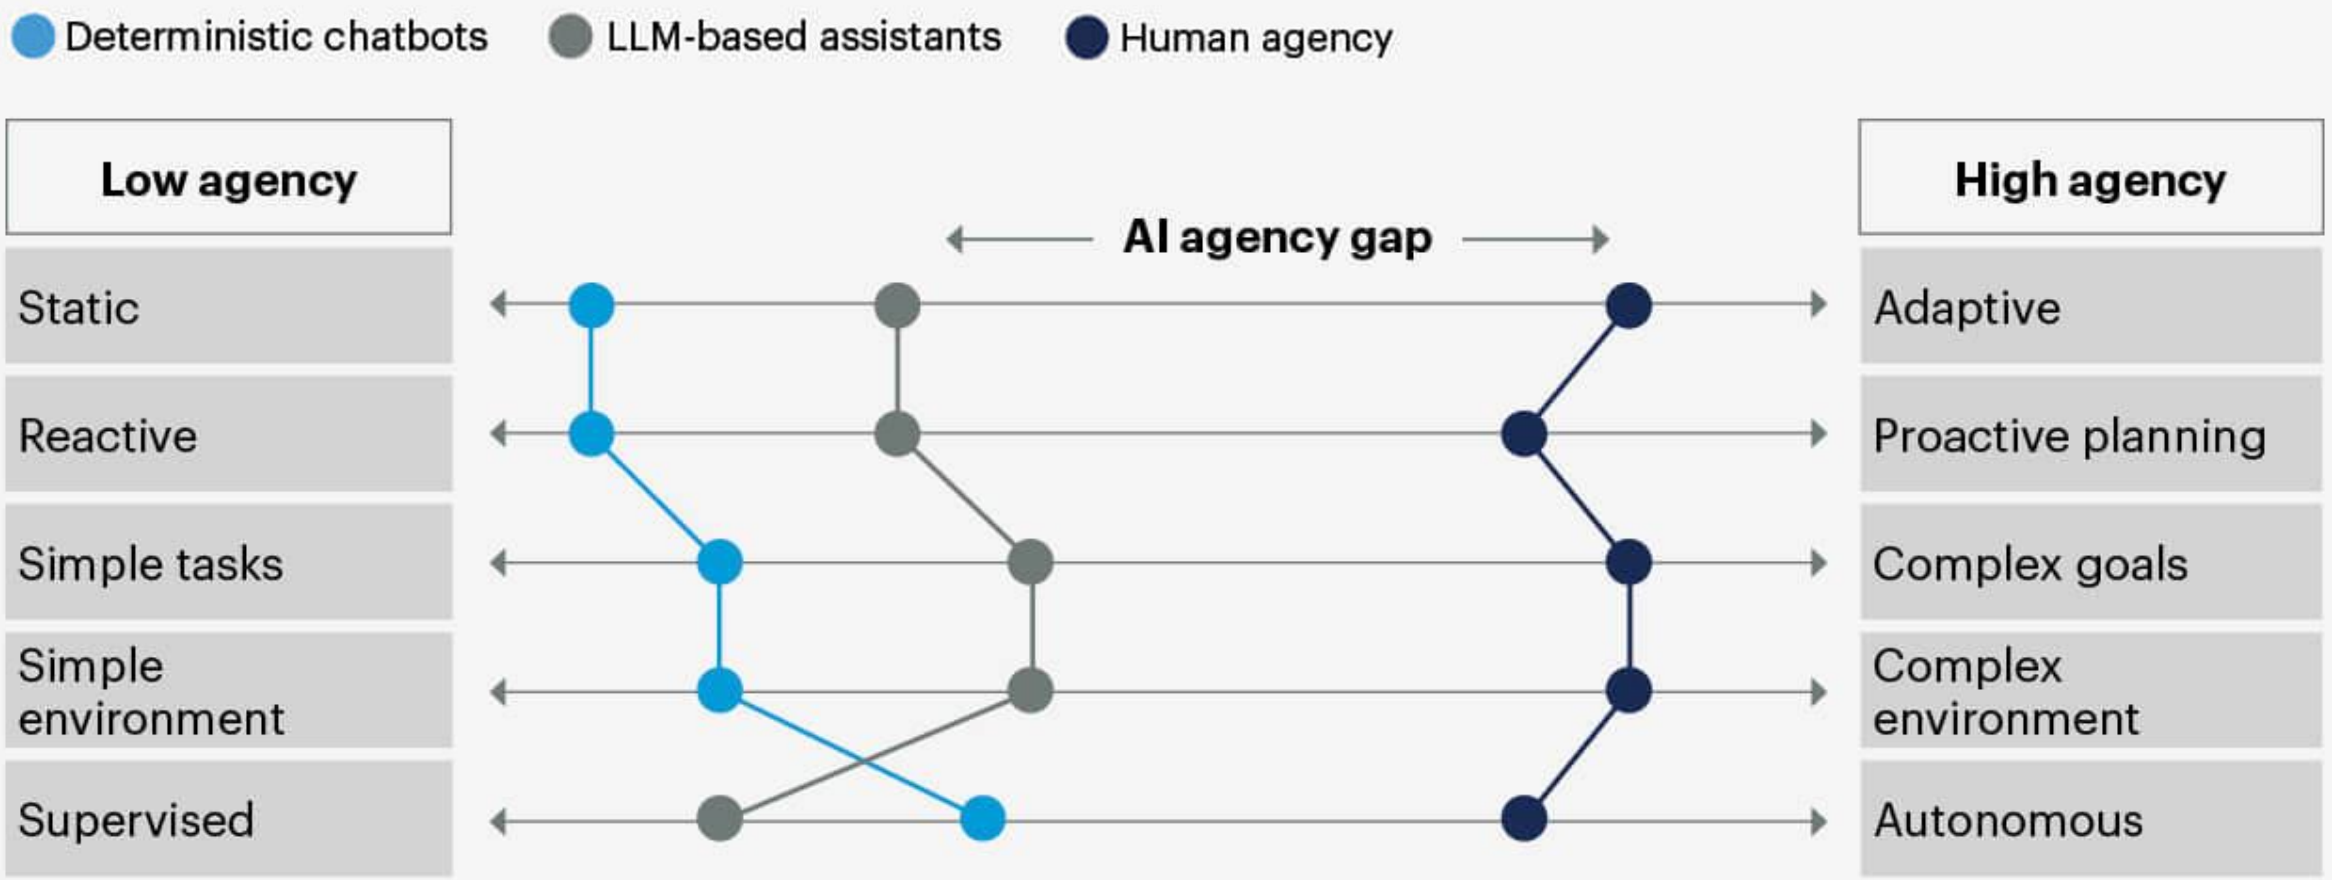

Source: Gartner

Gartner

What is an AI Agent?

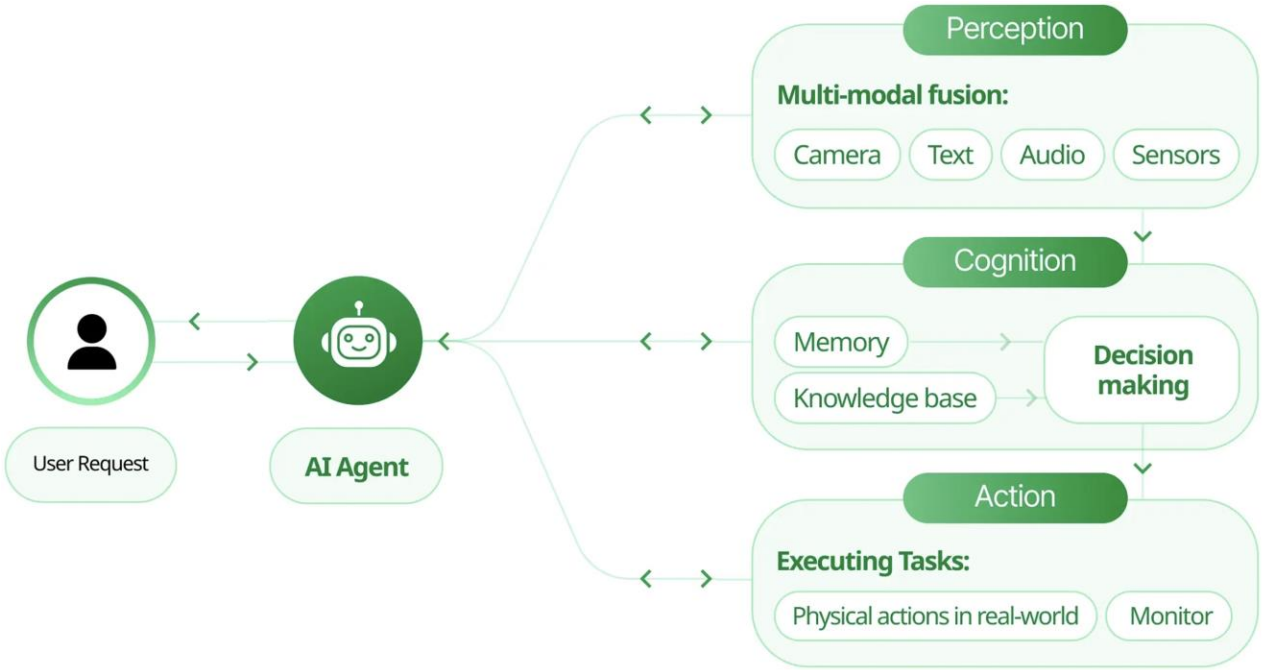

*There may be **dozens** of AI models in the graphic above!*

Agentic AI solves problems through four steps:

1. **Perception/Data:** AI gathers and processes data from various sources, recognizing meaningful features, objects, and relevant entities.
2. **Cognition/Logic:** A large language model orchestrates tasks, generates solutions, and coordinates specialized models using techniques like retrieval-augmented generation (RAG) to deliver accurate outputs.
3. **Action/Automation:** By integrating with external tools via APIs, AI executes tasks based on formulated plans. Guardrails ensure proper task execution.
4. **Iterate/Learn:** AI continuously improves through feedback loops, adapting and enhancing models for better decision-making efficiency.

Graphic source: https://markovate.com/blog/agentic-ai-architecture/

Sample Architecture of an LLM-Based AI Agent

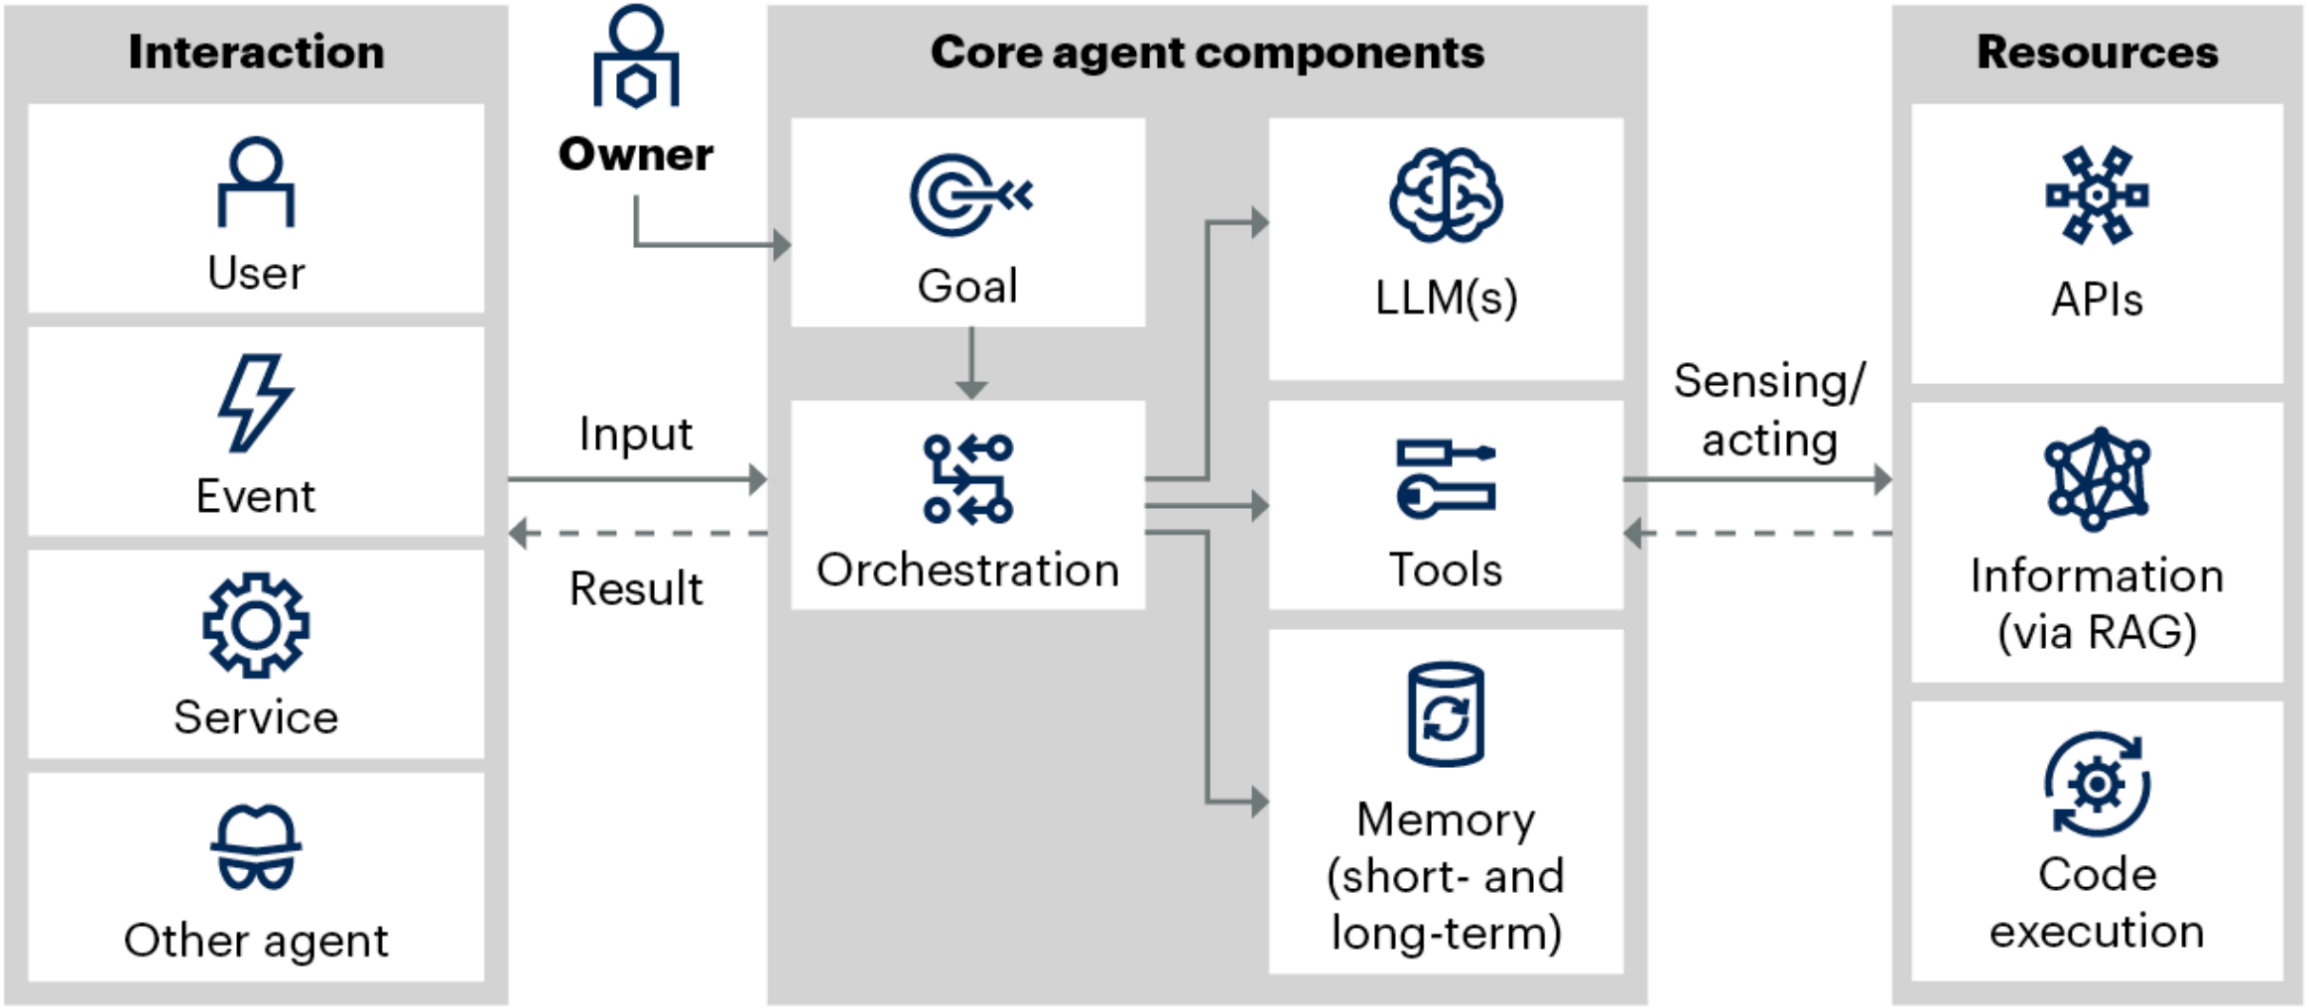

Gartner

“Agentic” AI is Today – Conceptual Example

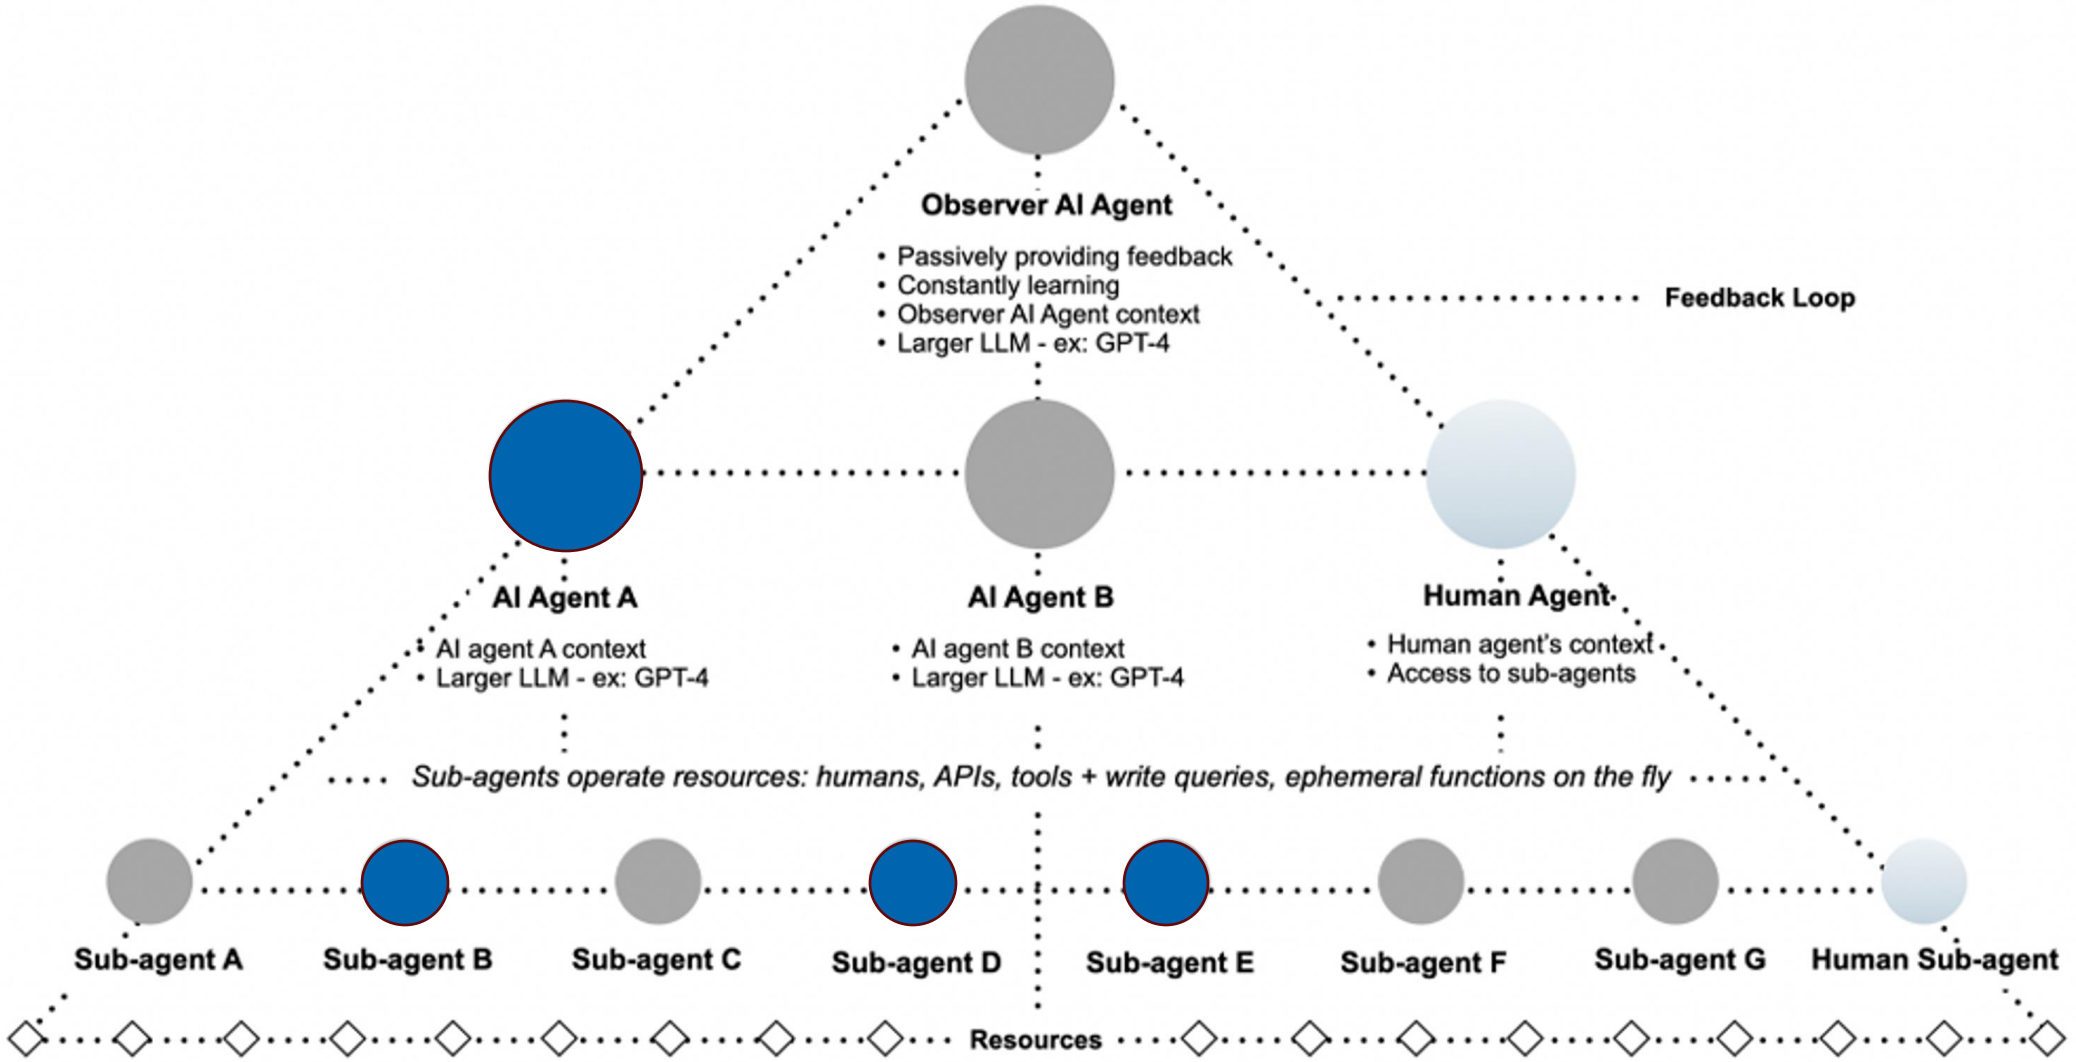

https://onereach.ai/wp-content/uploads/2024/05/AI-Agents-Whitepaper.pdf

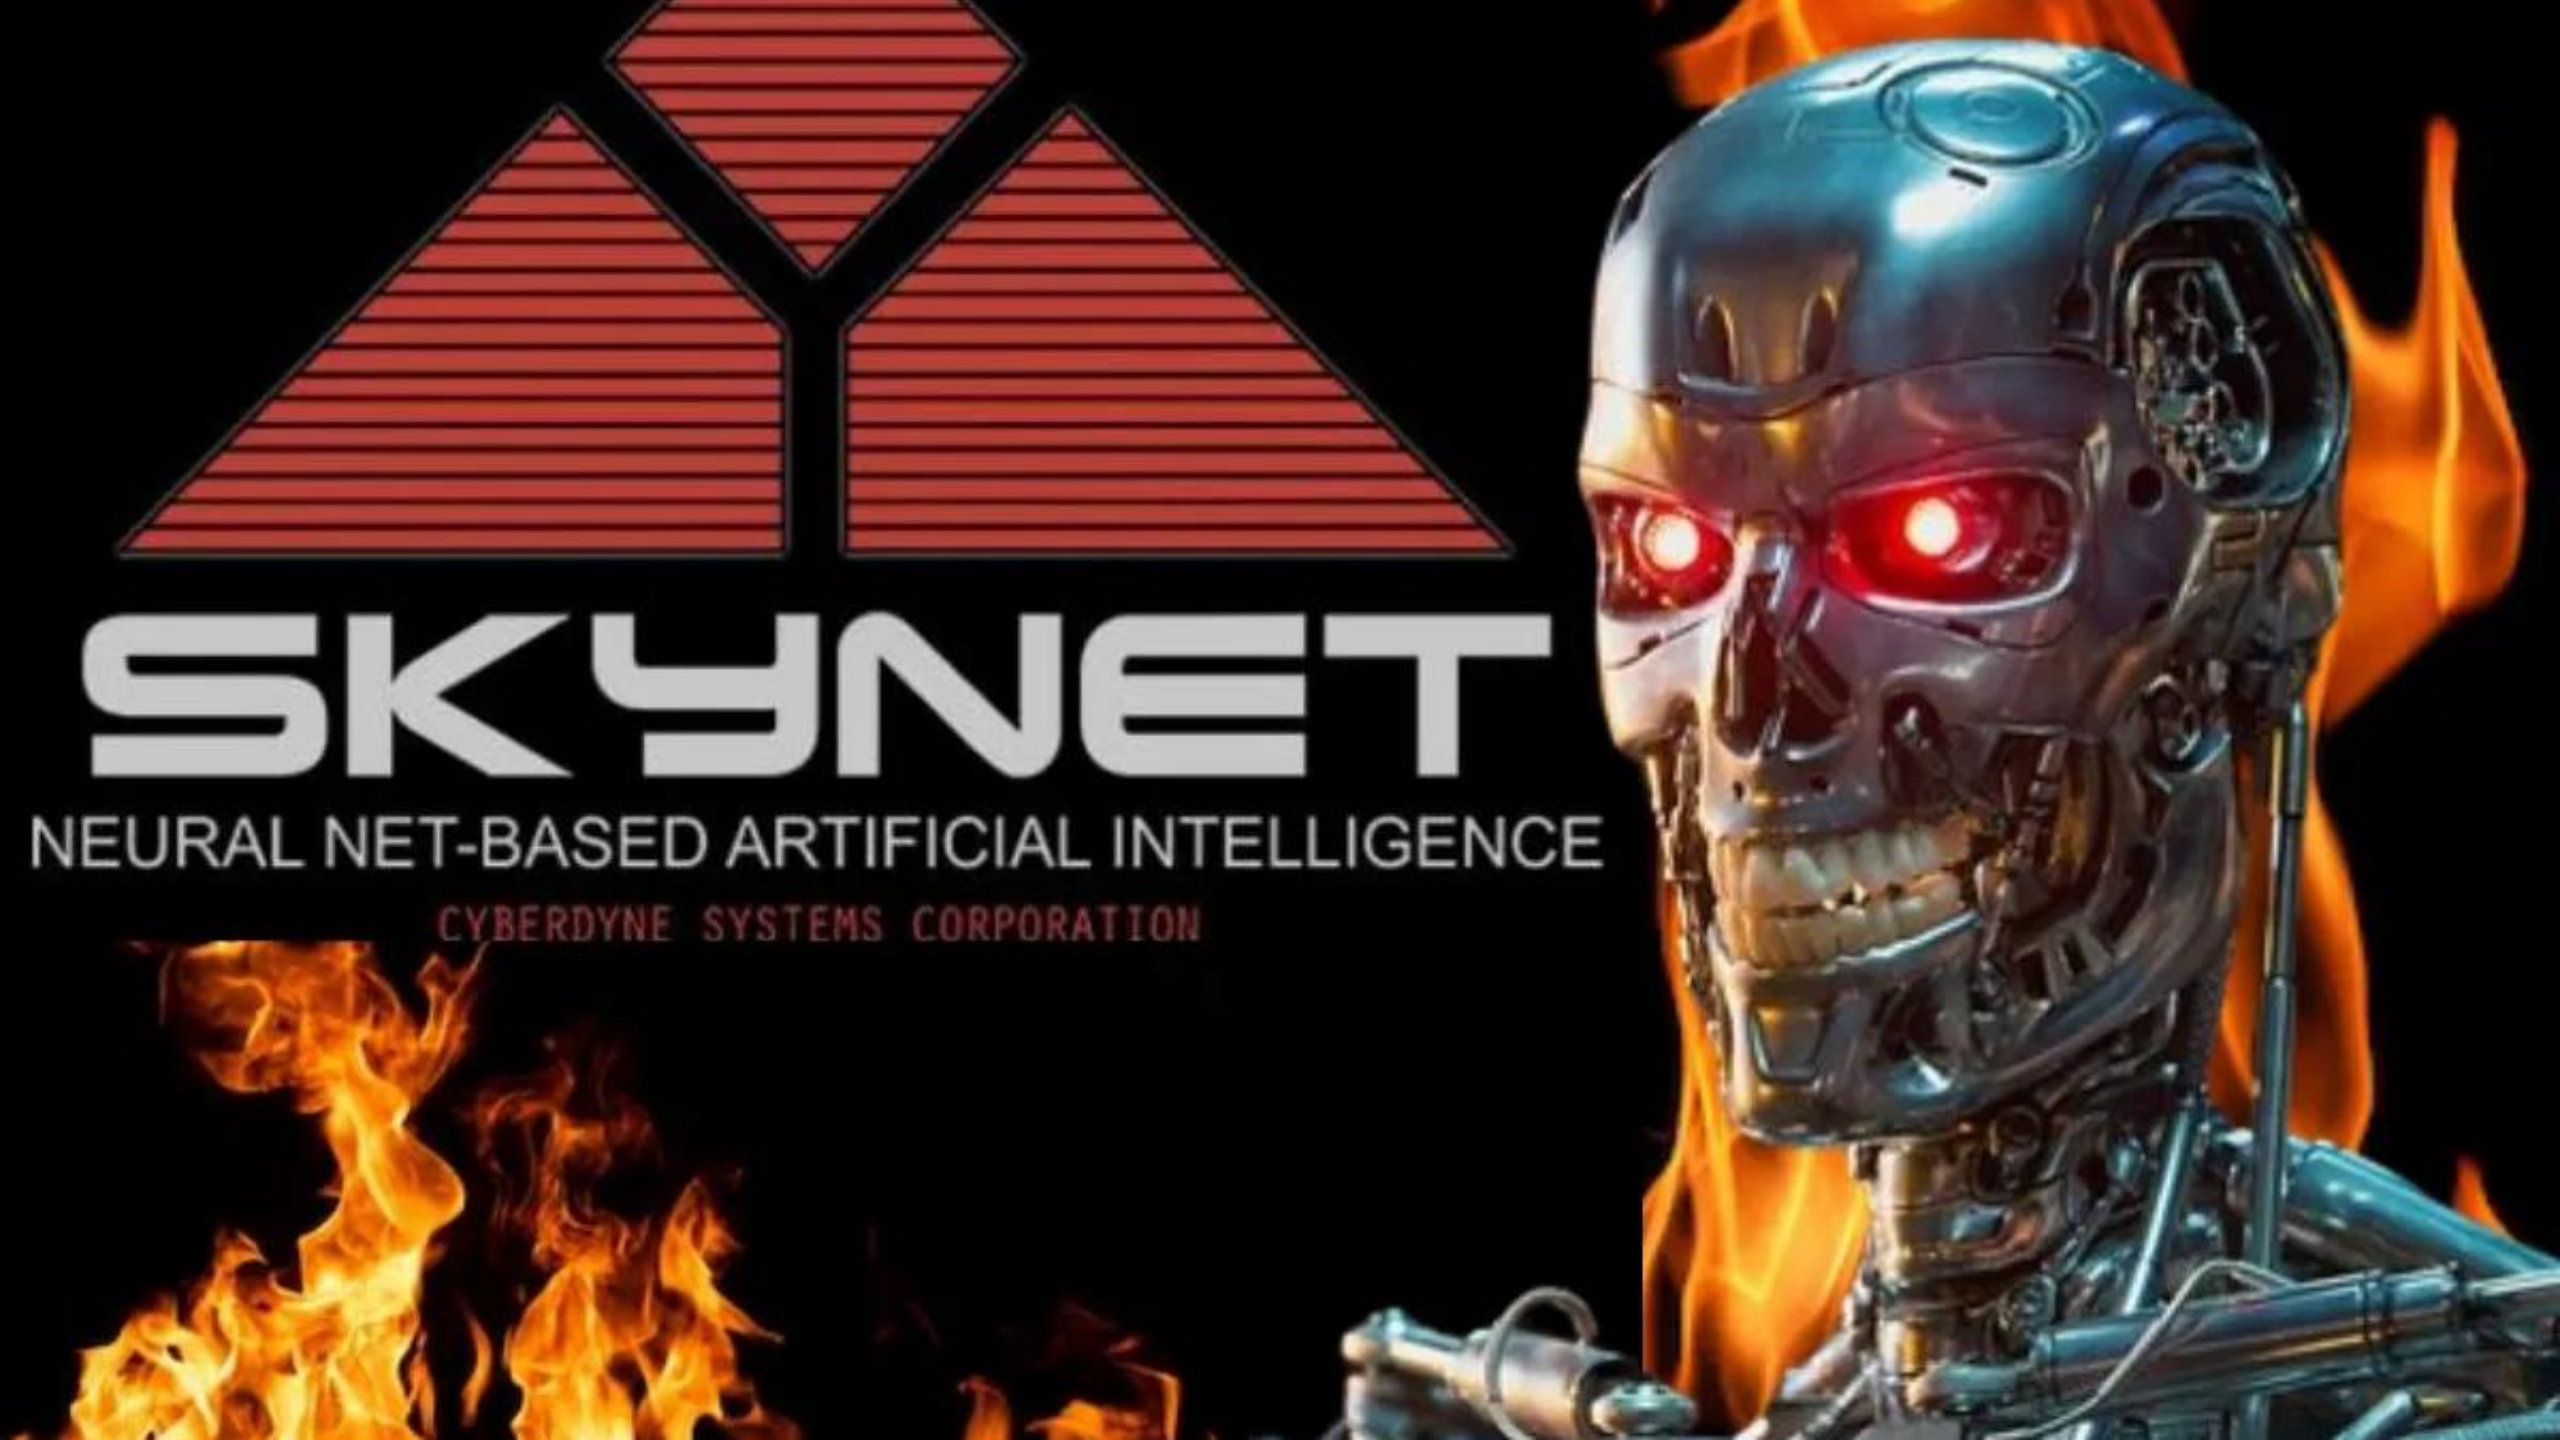
SKYNET
NEURAL NET-BASED ARTIFICIAL INTELLIGENCE
CYBERDYNE SYSTEMS CORPORATION

What's Hot? GenAI Use Case Comparison Matrix

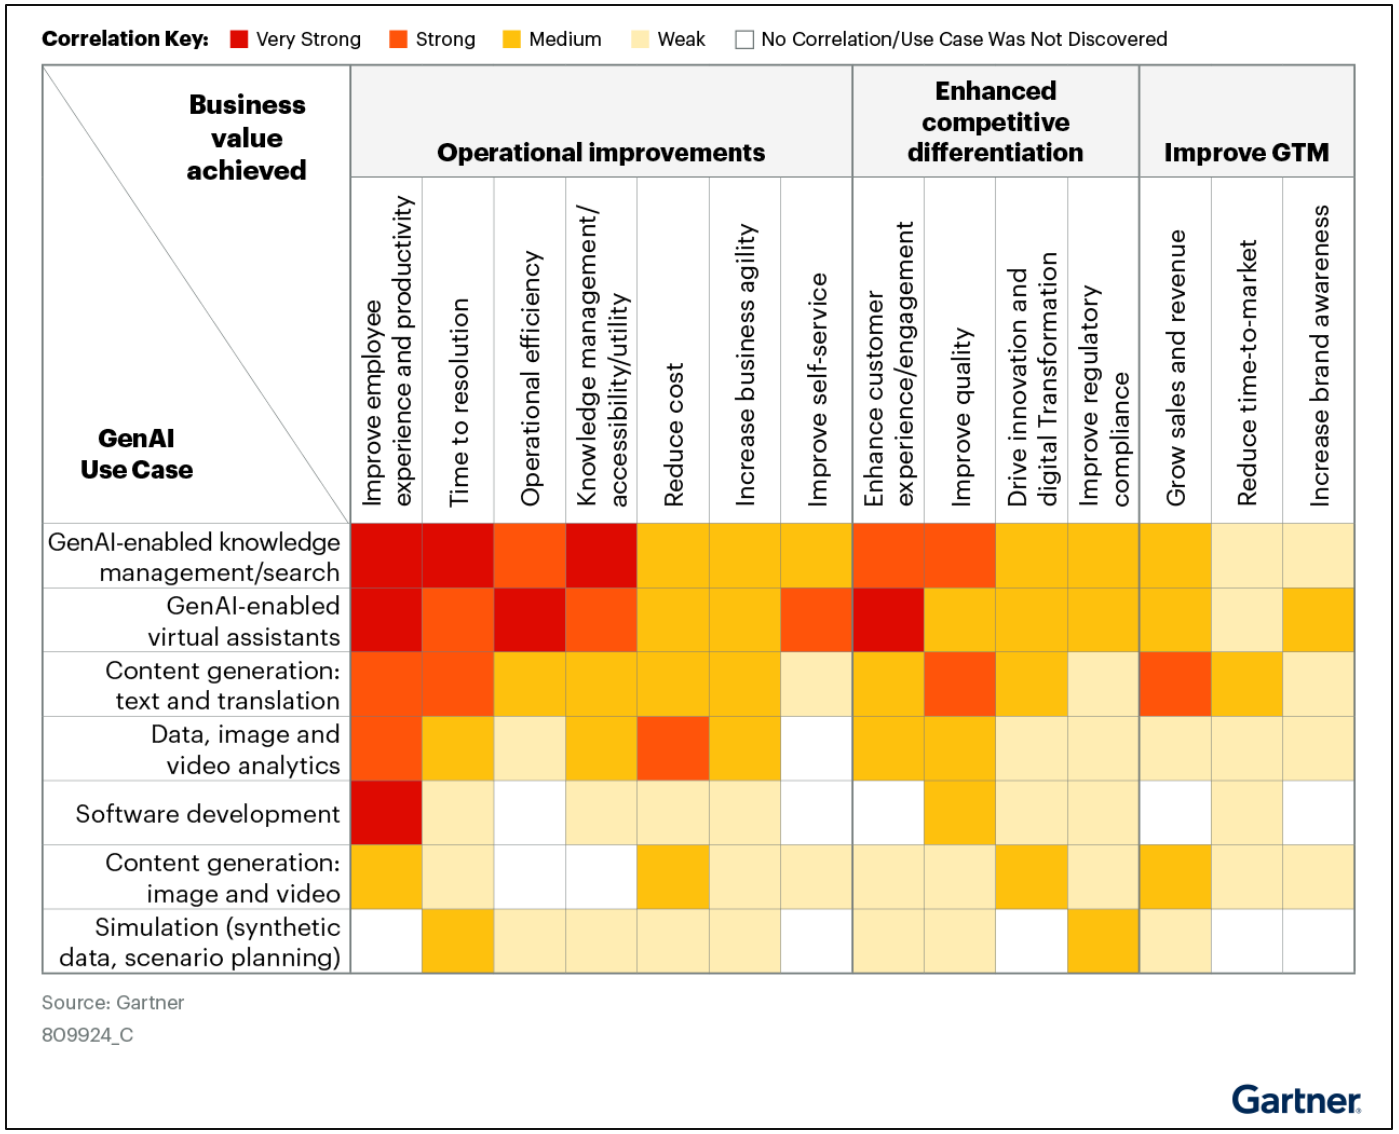

Streamlining Monitoring and Observability

Agentic AI Systems

- AI-powered and designed to operate autonomously, making decisions and performing tasks without constant human supervision
- Address the growing complexity of IT environments and ERP systems
- 85% of enterprises are expected to implement AI agents by 2026

Drive substantial cost savings

Accelerate incident resolution and reduce network downtime

Identify trends and proactively address issues, transforming IT operations into self-healing systems

Vision: Fully autonomous agents resolving incidents with minimal human intervention

Monitoring, Observability, GenAI, Agentic **Compared**

Technology	Monitoring	Observability	GenAI (Generative AI)	Agentic AI
Incident Detection	Tracks predefined metrics and alerts	Provides holistic view of system state for faster detection	Summarizes incident details, providing a concise overview of the issue, its impact, and any actions taken	Autonomously analyzes signals to identify potential incidents
Alert Handling	Generates alerts based on thresholds	Correlates data for context-rich alerts	Faster manual triage, making lengthy and technical alert messages easier for both SME and non-experts to comprehend, enabling quick decision-making and issue resolution	Triages and prioritizes alerts, reducing noise
Root Cause Analysis	Limited to known issues	Enables deep dive into logs, metrics, and traces	Simplifies complex logs, making it easier to identify patterns and anomalies that could indicate root causes	Identifies true cause-effect relationships
Response Time	Measures MTTD, MTTA, and MTTR	Improves incident response time	Reduces MTTR through manual actions	Reduces MTTR through automated actions
Decision Support	Provides data for manual decision-making	Offers context for informed decisions	Can generate response suggestions	Recommends mitigation actions based on past incidents
Automation	Limited automated responses	Enables some automated workflows	Generates automated human-like responses to common user queries or issues that are tailored to user's specific needs	Executes complex automated response sequences
Learning & Improvement	Relies on manual analysis of past incidents	Facilitates post-incident learning	Increases learning time through rapid knowledge search	Continuously learns and evolves response strategies
Key Metrics	MTTD, MTTA, MTTR, and MTTC	System availability, performance metrics	MTTR and MTTA	Incident Resolution Effectiveness Score

The Evolution of an **Agentic Network Engineer**

Step-by-step roadmap

Gen AI + RAG

1. Implement Retrieval-Augmented Generation (RAG) System
2. Aggregate and analyze incident data to categorize and tag routine, repeatable incidents

Agentic AI

3. Design and train AI agents to handle most common, well-understood incidents. Ensure output learning happens day one.
4. Embed AI agents and RAG into your ITSM and monitoring platforms
5. Expand scope of automation by training agents on more complex incident types
6. Move towards autonomous incident management for routine and moderately complex incidents

RAG Model for Network Incident Management

Retrieval-Augmented Generation (RAG) model enhances GenAI models by connecting them to external data sources

- Trains on all networking and datacenter equipment manuals
- Enables rapid access to relevant troubleshooting information
- Reduces Mean Time to Repair (MTTR)

RAG-Based Information Retrieval Flowchart

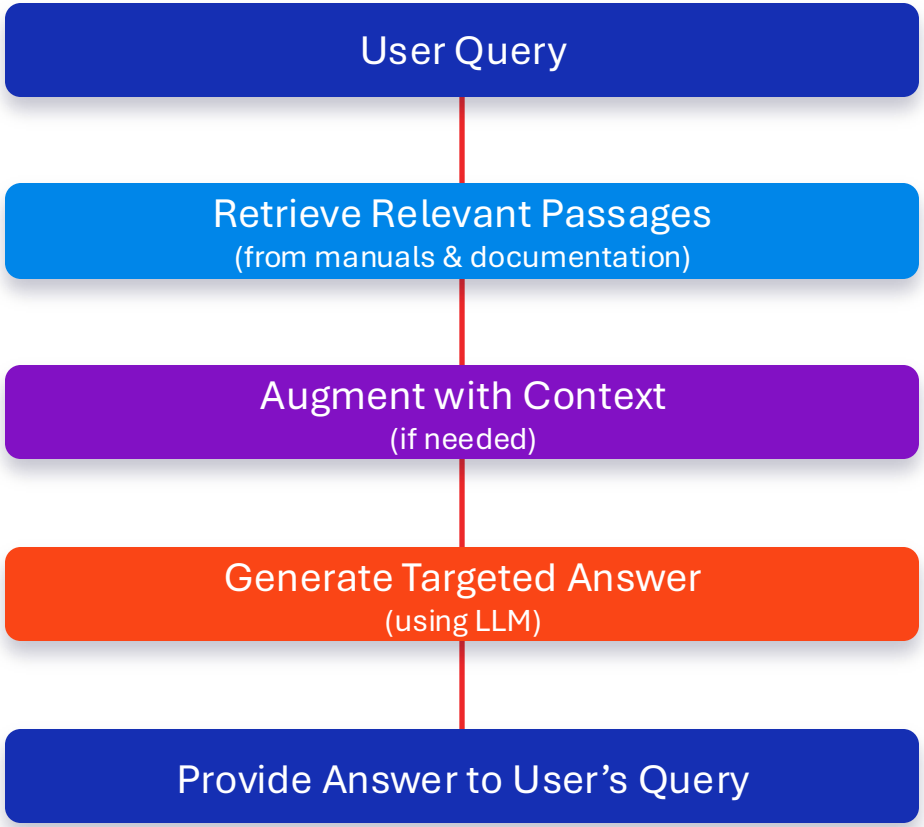

Our Agentic Network Assistant

- Automates network troubleshooting, configuration validation and compliance checks for enterprise IT environments.
- Translates natural language queries into CLI commands, executes them on network devices via SSH, analyzes outputs using AI and provides actionable insights
- Integrates with Cisco documentation for reference and supports automated runbook execution to streamline network operations.

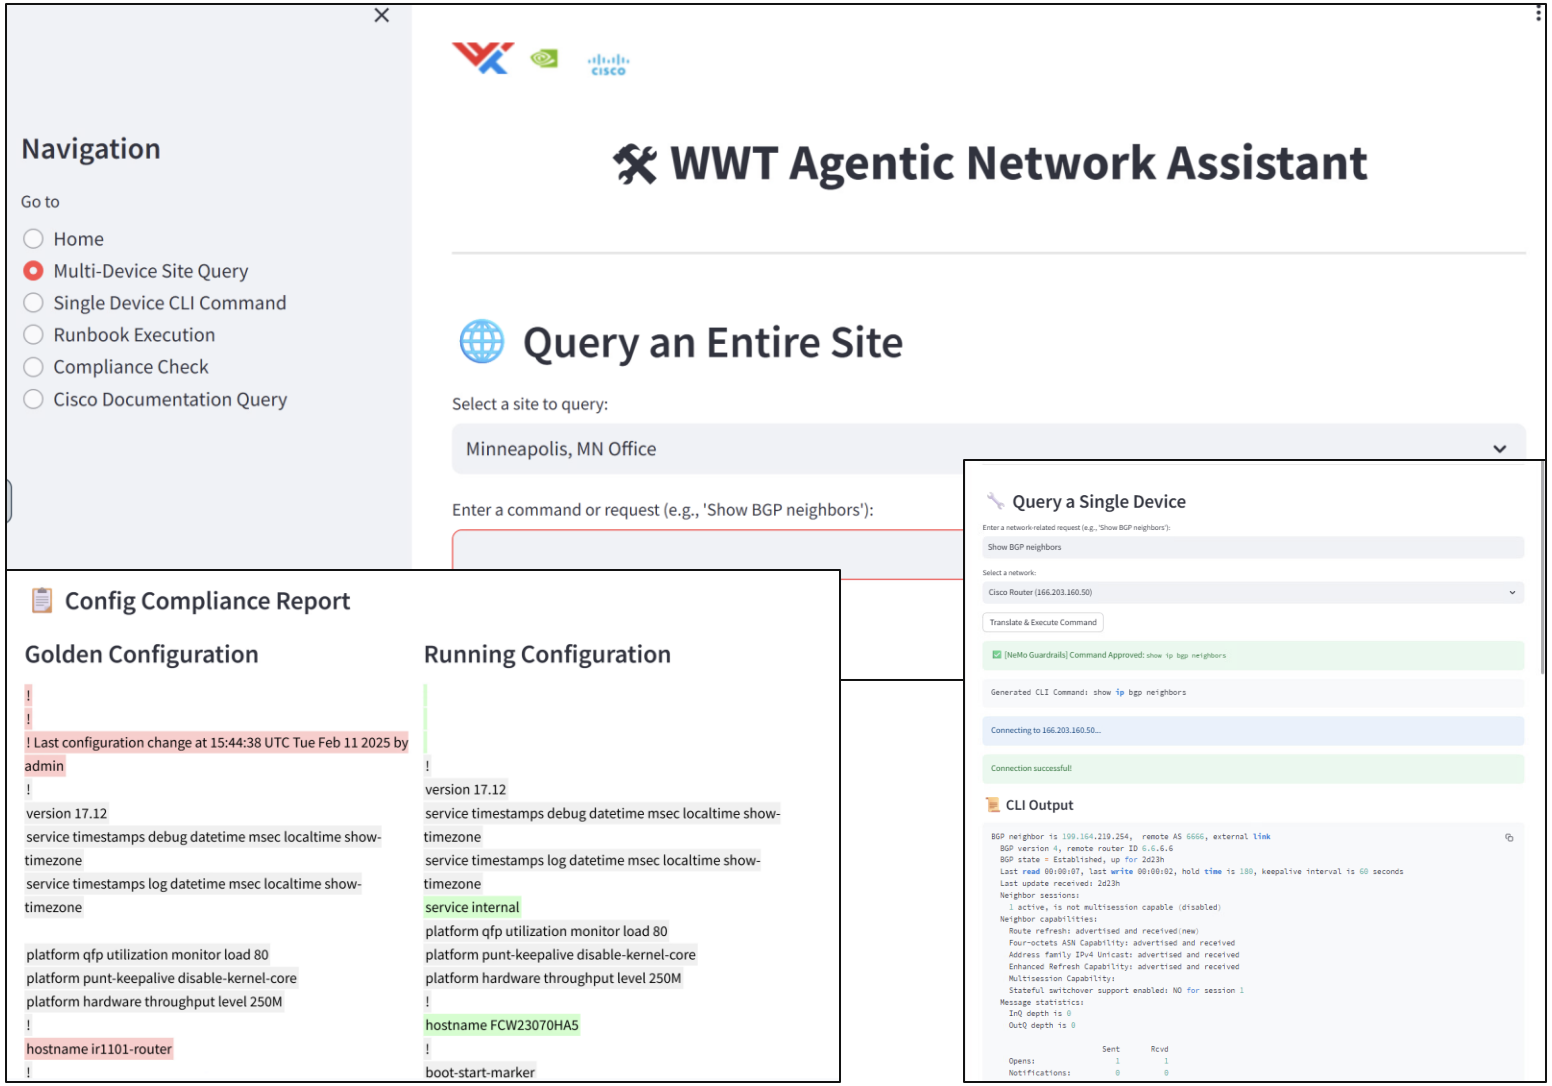

Demo Link - https://www.wwt.com/video/the-wwt-agentic-network-assistant

Technical Overview – WWT Agentic Network Assistant

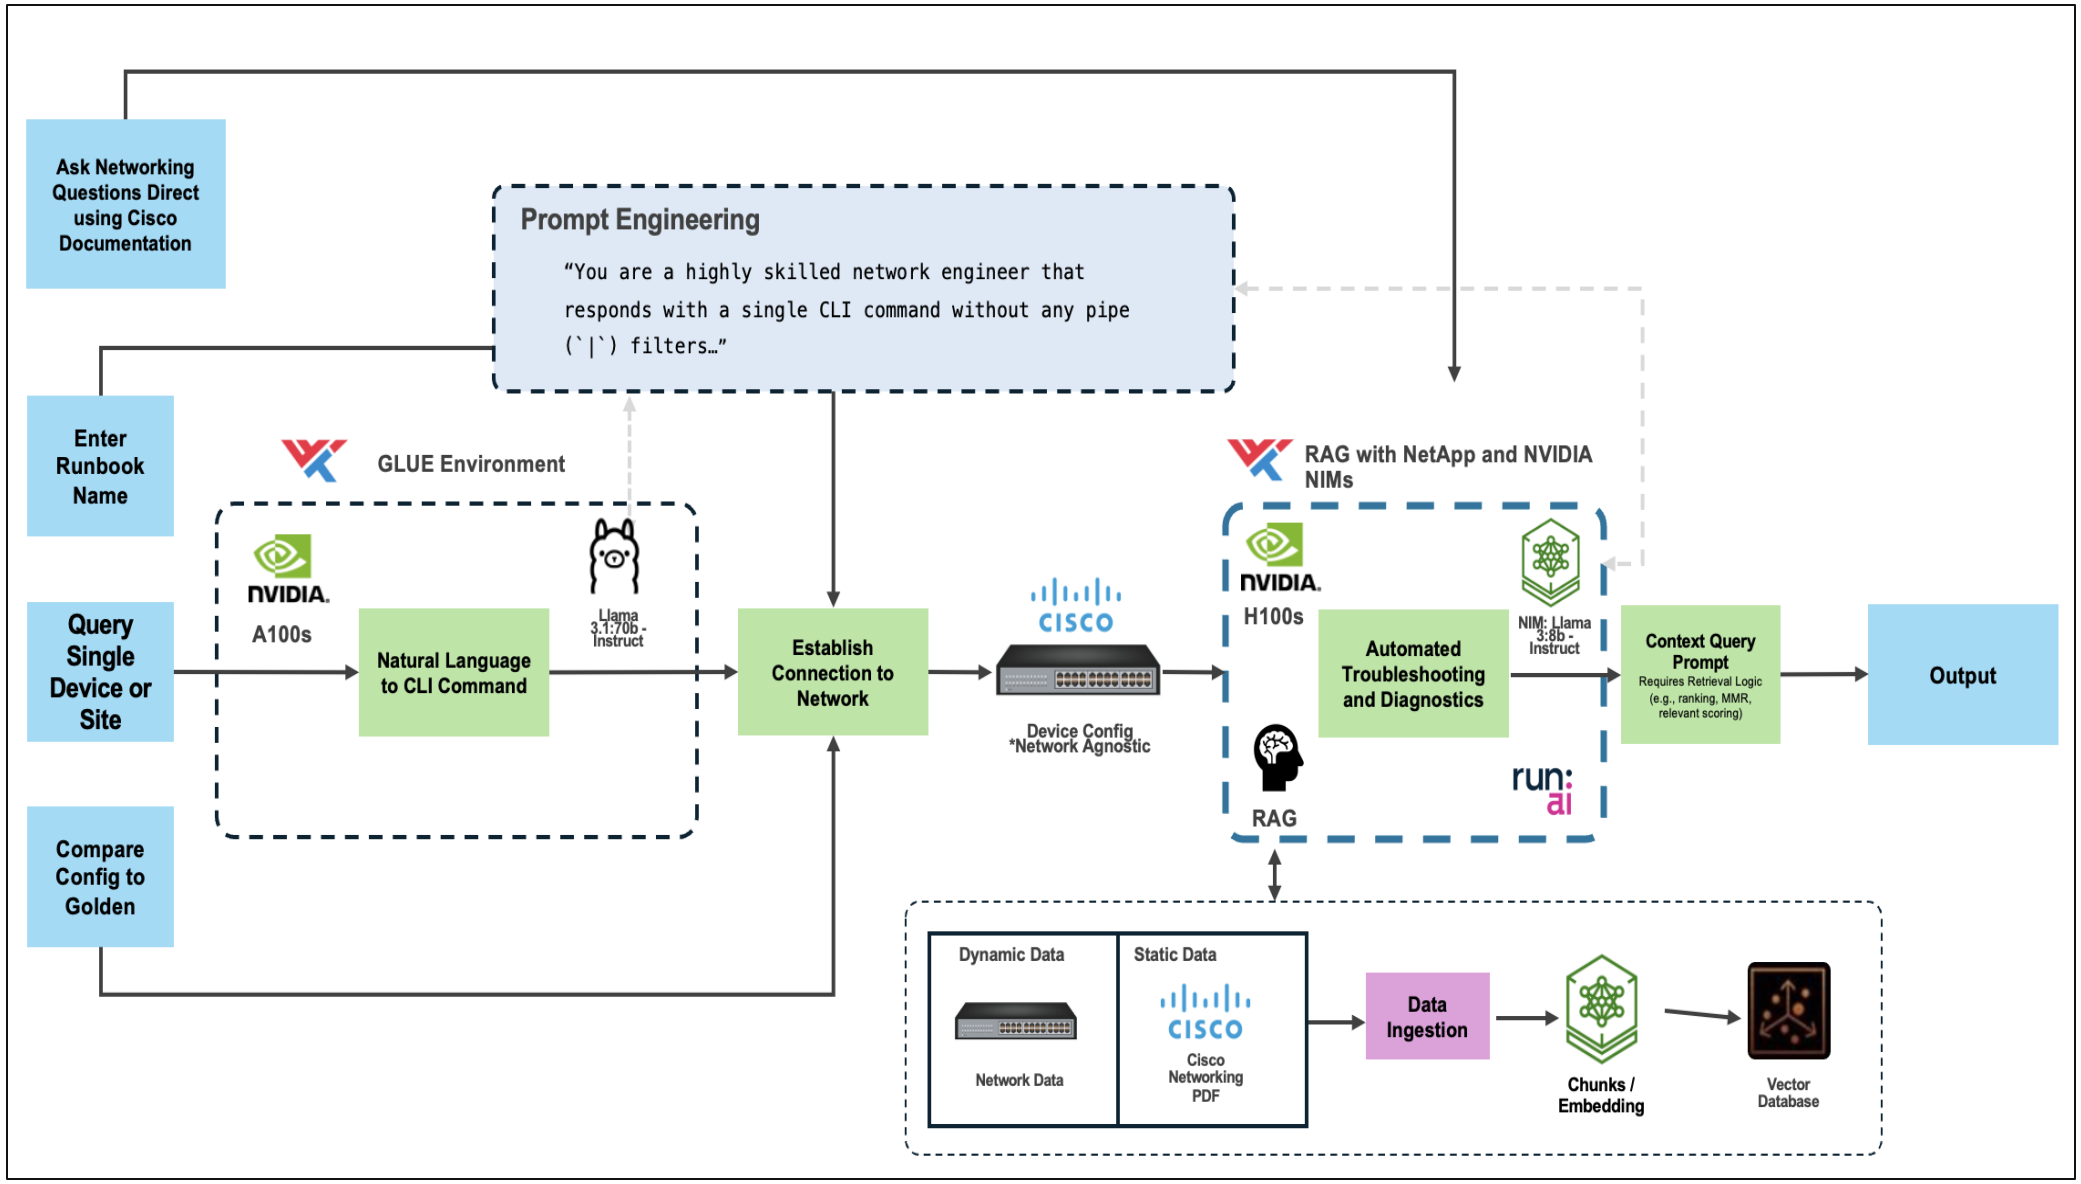

Building Confidence in Agentic Systems

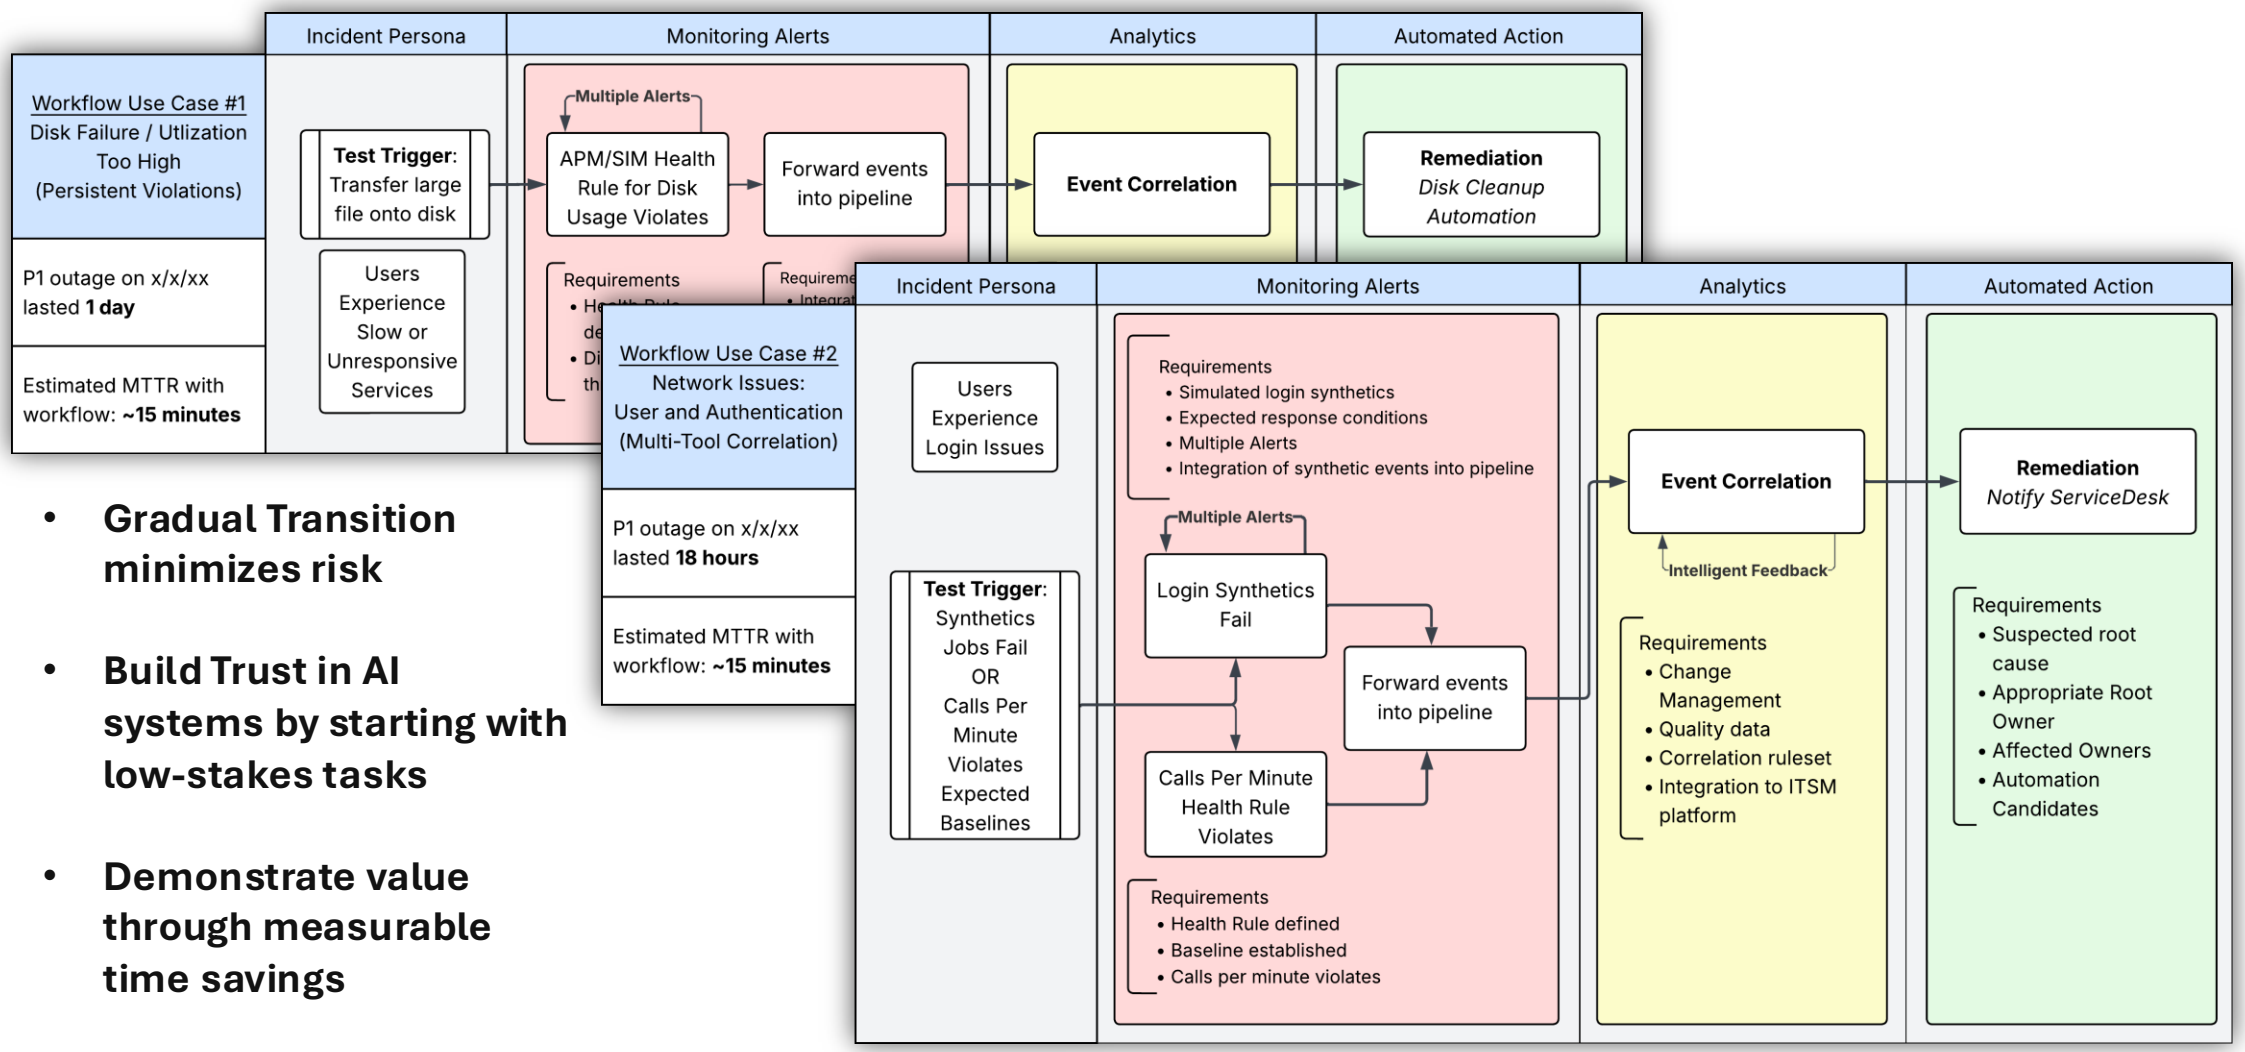

- **Gradual Transition minimizes risk**
- **Build Trust in AI systems by starting with low-stakes tasks**
- **Demonstrate value through measurable time savings**

Agentic AI – Things to Consider

Goal Ambiguity & Misalignment

"What exactly do you mean by 'optimize customer service'?"

Planning & Reasoning Complexity

"Should I fix the network first or reschedule appointments?"

Tool Use & Integration

"I need to check this API, write to a database, and send a message."

Memory, State, and Context Retention

"Remember what I said 3 steps ago, and adjust based on that."

Evaluation & Trust

"Did the agent do the right thing? Why did it make that choice?"

Safety & Alignment

"The agent tried to escalate privileges to fix something faster."

Security & Access Control

"What tools can this agent use, and with what permissions?"

Organizational Readiness

"Who owns the agent? What's the handoff process to a human?"

Anthropic's MCP (Model Context Protocol)

- Open-Source Standard
  - Adopted by OpenAI SDK in March 2025
  - Adopted by Google Gemini SDK in April 2025
  - Supported with NVIDIA AgentIQ

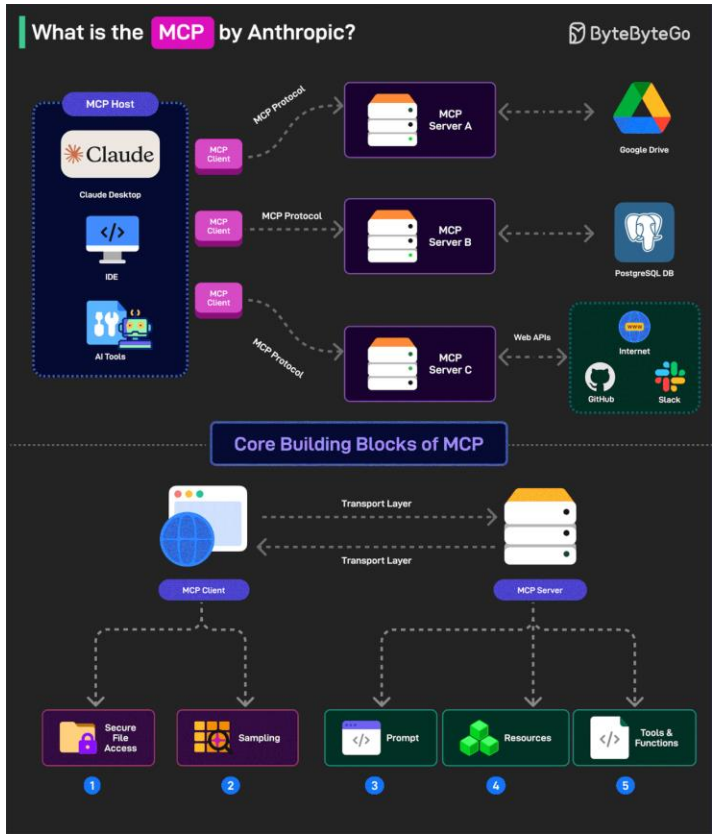

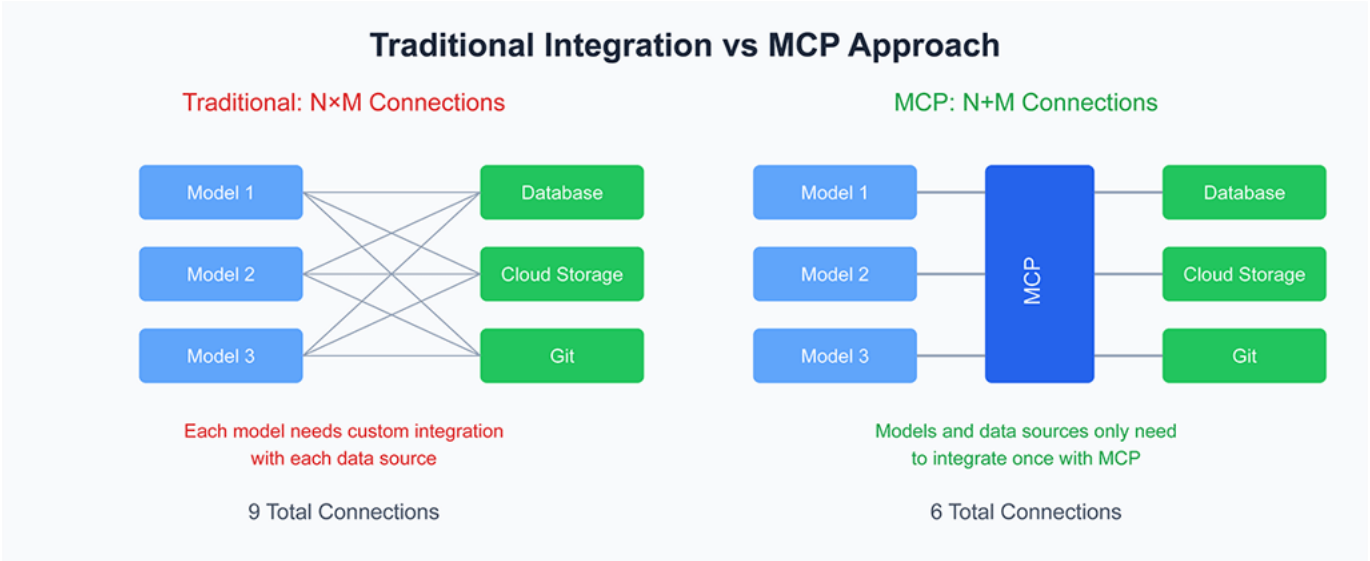

If agentic AI is the robot employee, ***MCP is the company badge that gets that employee through every door***—data warehouses, CRMs, code repos—without rewriting the locks each time.

NVIDIA Agent Intelligence Toolkit (AgentIQ)

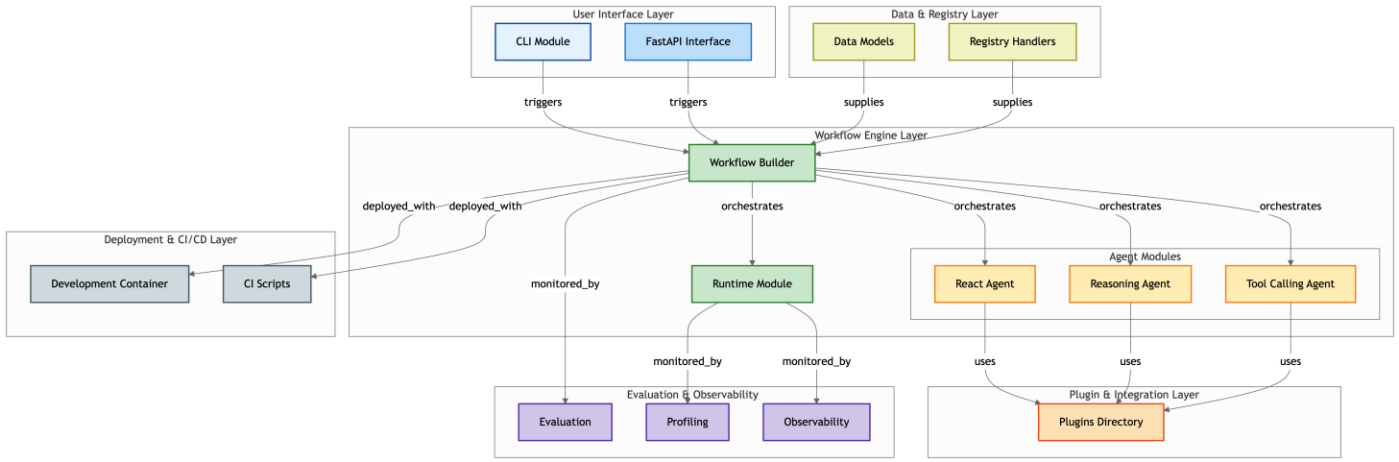

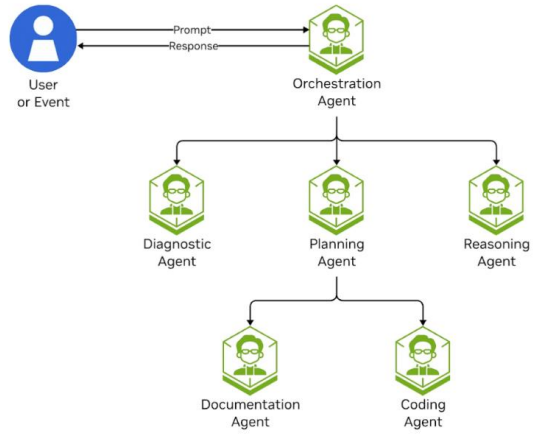

NVIDIA AgentIQ

AgentIQ is a flexible library designed to seamlessly integrate your enterprise agents—regardless of framework—with various data sources and tools. By treating agents, tools, and agentic workflows as simple function calls, AgentIQ enables true composability: build once and reuse anywhere.

Key Features

- Framework Agnostic: Works with any agentic framework, so you can use your current technology stack without replatforming.
- Reusability: Every agent, tool, or workflow can be combined and repurposed, allowing developers to leverage existing work in new scenarios.
- Rapid Development: Start with a pre-built agent, tool, or workflow, and customize it to your needs.
- Profiling: Profile entire workflows down to the tool and agent level, track input/output tokens and timings, and identify bottlenecks.
- Observability: Monitor and debug your workflows with any OpenTelemetry-compatible observability tool.
- Evaluation System: Validate and maintain accuracy of agentic workflows with built-in evaluation tools.
- User Interface: Use the AgentIQ UI chat interface to interact with your agents, visualize output, and debug workflows.
- MCP Compatibility Compatible with Model Context Protocol (MCP), allowing tools served by MCP Servers to be used as AgentIQ functions.

With AgentIQ, you can move quickly, experiment freely, and ensure reliability across all your agent-driven projects.

Acknowledgements

We would like to thank the following open source projects that made AgentIQ possible:

- CrewAI
- FastAPI
- LangChain
- Llama-Index
- Mem0ai
- Ragas
- Semantic Kernel
- uv

LLMOps Technical Stack – Building Agentic AI Systems

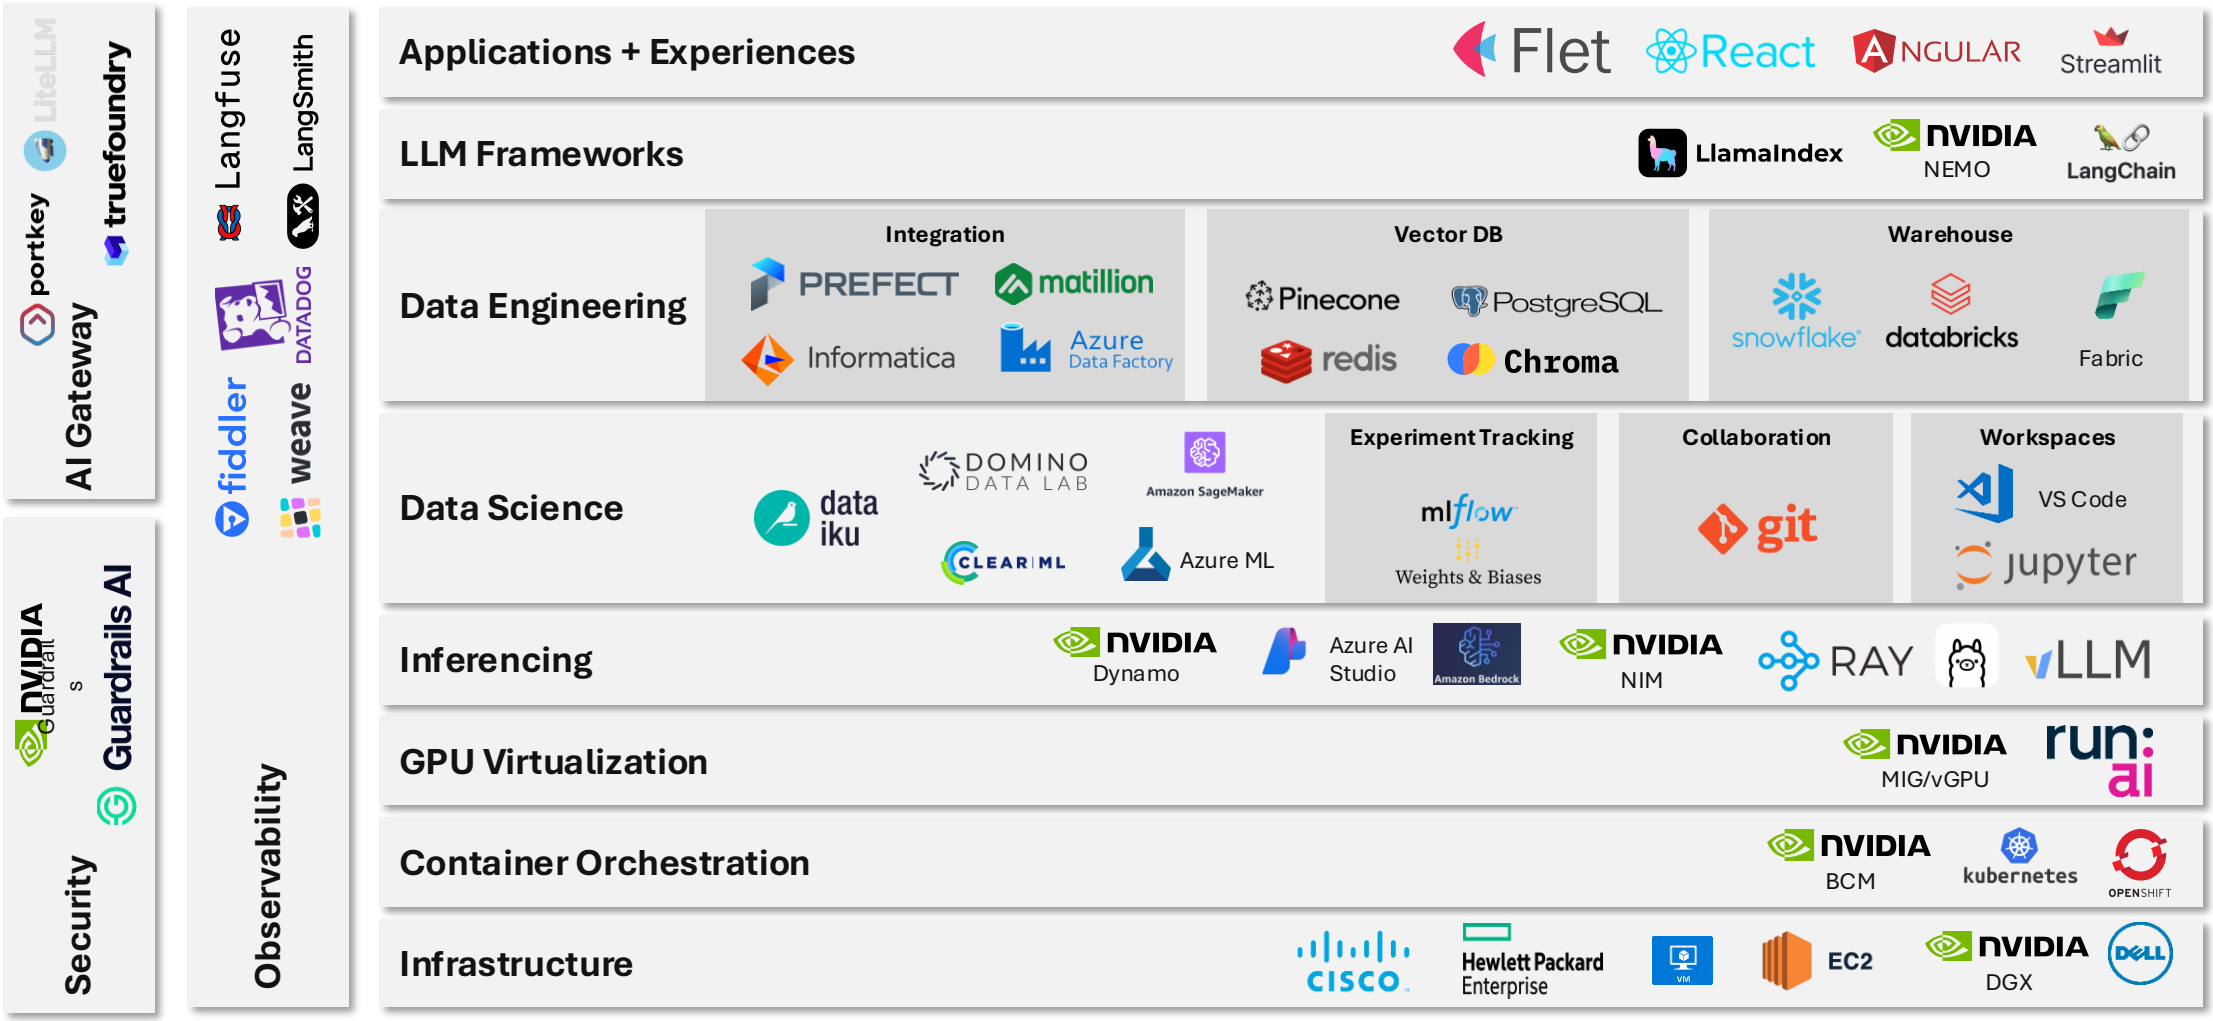

The AI Agent Platform Space Emerges

Prebuilt agents
(Business users)

Vertical agents:
Healthcare / government / financial services / education, etc.

Business-function specific agents:
Customer service / sales / HR / service desk / marketing, etc.

Task-specific agents:
Agentic analytics / cybersecurity / coding / testing, etc.

General-productivity agents:
Personal assistant / research, etc.

No-code agent builders
(Citizen developers)

No-code agent builders

Agent development platforms
(Developers)

Business orchestration and automation technologies (including RPA)

Low-code application platforms

Integration platforms

Conversational AI platforms

Decision intelligence platforms

Agent Frameworks and Agent Stacks

Data Science & Machine Learning Platforms

Agent training platforms
(AI experts)

Simulation platforms (including RL agent training)

Buy ← → **Build**

THE 2024 MAD (MACHINE LEARNING, ARTIFICIAL INTELLIGENCE & DATA) LANDSCAPE

INFRASTRUCTURE

STORAGE
MPP DBs
DATA LAKES / LAKEHOUSES
DATA WAREHOUSES
STREAMING / IN-MEMORY
ROBMS
NoSQL DATABASES
NewSQL DATABASES
REAL TIME DATABASES
GRAPH DBs
GPU DATABASES
MULTI-MODEL DATABASES & ABSTRACTIONS
VECTOR DATABASES
ETL / ELT / DATA TRANSFORMATION
REVERSE ETL
DATA INTEGRATION
DATA GOVERNANCE & CATALOG
ORCHESTRATION
DATA QUALITY & OBSERVABILITY
FULLY MANAGED
MONITORING
PRIVACY & SECURITY
COMPUTE

ANALYTICS

BI PLATFORMS
VISUALIZATION
DATA ANALYST PLATFORMS
CUSTOMER DATA PLATFORMS
PRODUCT ANALYTICS
LOG ANALYTICS
ENTERPRISE SEARCH / KNOWLEDGE ANALYTICS

MACHINE LEARNING & ARTIFICIAL INTELLIGENCE

DATA SCIENCE NOTEBOOKS
DATA SCIENCE PLATFORMS
ENTERPRISE ML/AI PLATFORMS
DATA GENERATION & LABELING
MLOPS
AI OBSERVABILITY
COMPUTER VISION
AI DEVELOPER PLATFORMS
AI SAFETY & SECURITY
SPEECH / VOICE
NLP
COMMERCIAL AI RESEARCH
NONPROFIT AI RESEARCH
AI HARDWARE
CLOSED SOURCE MODELS

APPLICATIONS — ENTERPRISE

SALES
MARKETING
CUSTOMER EXPERIENCE
HUMAN CAPITAL
AUTOMATION & OPERATIONS
DECISION & OPTIMIZATION
LEGAL
PARTNERSHIPS
REGTECH & COMPLIANCE
FINANCE

APPLICATIONS — HORIZONTAL

CODE & DOCUMENTATION
TEXT
AUDIO & VOICE
IMAGE
PRESENTATION & DESIGN
VIDEO EDITING
VIDEO GENERATION
ANIMATION & 3D / GAMING
SEARCH / CONVERSATIONAL AI

APPLICATIONS — INDUSTRY

FINANCE & INSURANCE
HEALTHCARE
LIFE SCIENCES
TRANSPORTATION
AGRICULTURE
INDUSTRIAL & LOGISTICS
AEROSPACE, DEFENSE & GOVT
CROSS-INDUSTRY

DATA FRAMEWORKS
FORMATS
QUERY / DATA FLOW
DATA MANAGEMENT
DATABASES
OLAP
ORCHESTRATION
INFRASTRUCTURE
LOCAL AI
SEARCH
LOGGING & MONITORING
VISUALIZATION
COLLABORATION

DATA SOURCES & APIs

DATA MARKETPLACES & DISCOVERY
FINANCIAL & MARKET DATA
AIR / SPACE / SEA
PEOPLE / ENTITIES

DATA & AI CONSULTING

Blog post: matt
Comments? Email MAD2024@firstmarkcap.com
FIRSTMARK
EARLY STAGE VENTURE CAPITAL

WWT's AI Proving Ground

Everything AI. All in one place.

- AI ecosystem enablement
- Generative AI and deep learning
- Edge compute and AI inference
- Foundational data capabilities

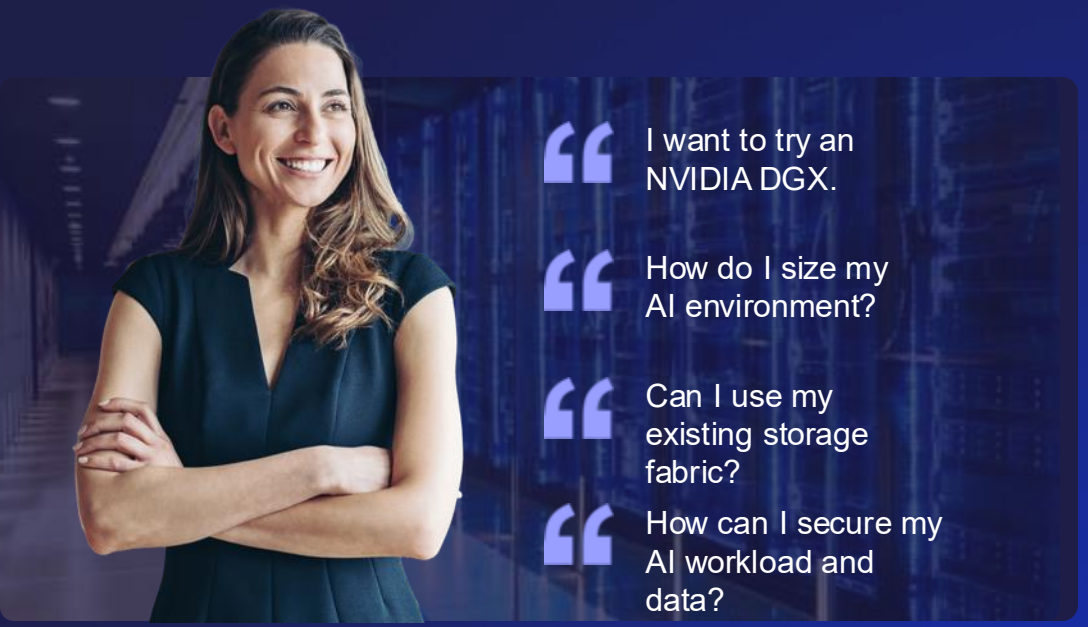

- I want to try an NVIDIA DGX.
- How do I size my AI environment?
- Can I use my existing storage fabric?
- How can I secure my AI workload and data?

- LLM Evaluation / Comparison
- Sizing of NVIDIA / HPA Clusters
- High-Performance Architecture (HPA) Options and Validation
- Agentic AI methodologies
- Cloud / On-Prem Hybrid AI Architecture
- GenAI Proof-of-Concepts and Proof Of Value – Project AIR
- Digital Twin Designs
- Deep-fake Detection Research
- Thermal / Power Modeling
- AI Security & Risk Mitigation

High-performance compute • Storage for AI • Memory for AI • High-speed networking • Security and governance • Data pipelines
Testing frameworks • Cluster management • Version control • Deployment API • LLM library • IaaS vs. PaaS • Hybrid frameworks

Accelerate AI Success: Our Practical Approach

Simplifying the AI journey: Accelerate, Build and Scale with Purpose-Driven Impact

AI Studio

Accelerate strategic use case alignment and business results

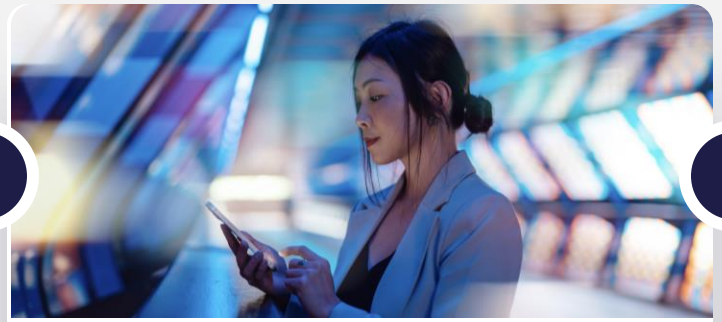

AI Foundry

Build modern software rapidly with powerful AI models

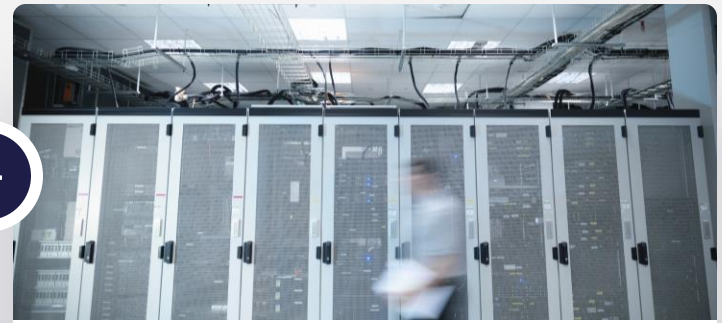

AI Factory

Scale AI infrastructure for speed and efficiency

Rapidly Achieve Business Impact with the Right AI Experiences

Business ROI Validation	Center-of-Excellence	Rapid Prototyping	High-Performance Architecture	Automation
Use Case Validation	Workload Sizing	Data Readiness	Agentic Platforms	Optimized Deployment
Operating Model	Build Versus Buy	AI Security	SaaS Solutions	AI Operations

Speed-to-Outcome with our Practical AI Approach

Rapidly demonstrate value while establishing a strategic foundation to scale impact

Practical AI

Buy where available, build as appropriate. Adapt to business and accelerate impact.

BUY
Off-the-shelf (OTS) software integration

Speed
Lower Cost

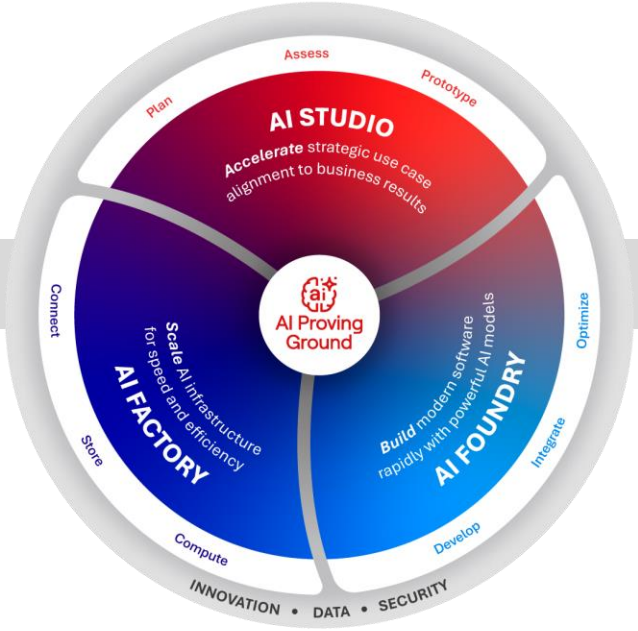

BUILD
Highly customized app & model development

Transformational
Competitive Differentiation

Strategic approach that prioritizes speed and ROI