

AI

AI is no longer a tool it is an operating advantage

Clayton Cross
Technologist

THE NEW CHALLENGE: CONTROLLING THE ECONOMICS

PERSONAL USE



ChatGPT (Codex)

Talking while driving • Image generation
Health • food • everyday questions



Claude

Prompts • review
Strategy • red teaming



Gemini

Very rarely
Only when other platforms are busy

MODELS & AGENTS




DeepSeek R1 0528 Qwen3 8B

Cost & capability testing
Open-weight • locally run •
lower-cost alternative



OpenClaw ×5

ACTIVE AGENTS



Hermes ×2

ACTIVE AGENTS

THREE COST DRIVERS

01



MODEL PRICE

Cost per token


02



REASONING DEPTH

How long it thinks

03



AGENT BEHAVIOUR

Calls, loops and tools

SAME TASK. DIFFERENT MODEL. DIFFERENT COST.

ONE 60-MINUTE TRANSCRIPT	12,000 INPUT TOKENS	1,000 OUTPUT TOKENS		13,000 TOTAL TOKENS PROCESSED	ONE UNCACHED API CALL	NO TOOLS OR ADDITIONAL REASONING
-----------------------------	------------------------	------------------------	---	-------------------------------------	--------------------------	-------------------------------------

MODEL		TOKENS PROCESSED	APPROX. COST (AUD)
	Claude Fable 5	13,000	A\$0.24
	OpenAI GPT-5.6 Sol	13,000	A\$0.13
	Gemini 3.1 Pro	13,000	A\$0.05
	DeepSeek Reasoner	13,000	A\$0.01


(GPT-5.6 SOL ULTRA: APPROXIMATELY A\$0.56 PER TASK)

LESSON: DO YOUR LEARNING AND HEAVY EXPERIMENTATION ON SUBSCRIPTIONS WHILE COSTS ARE CAPPED; MOVE TO APIs ONLY WHEN AUTOMATION JUSTIFIES THE METER.

ATLAS (AI AGENT / AGENTIC WORKFLOW)



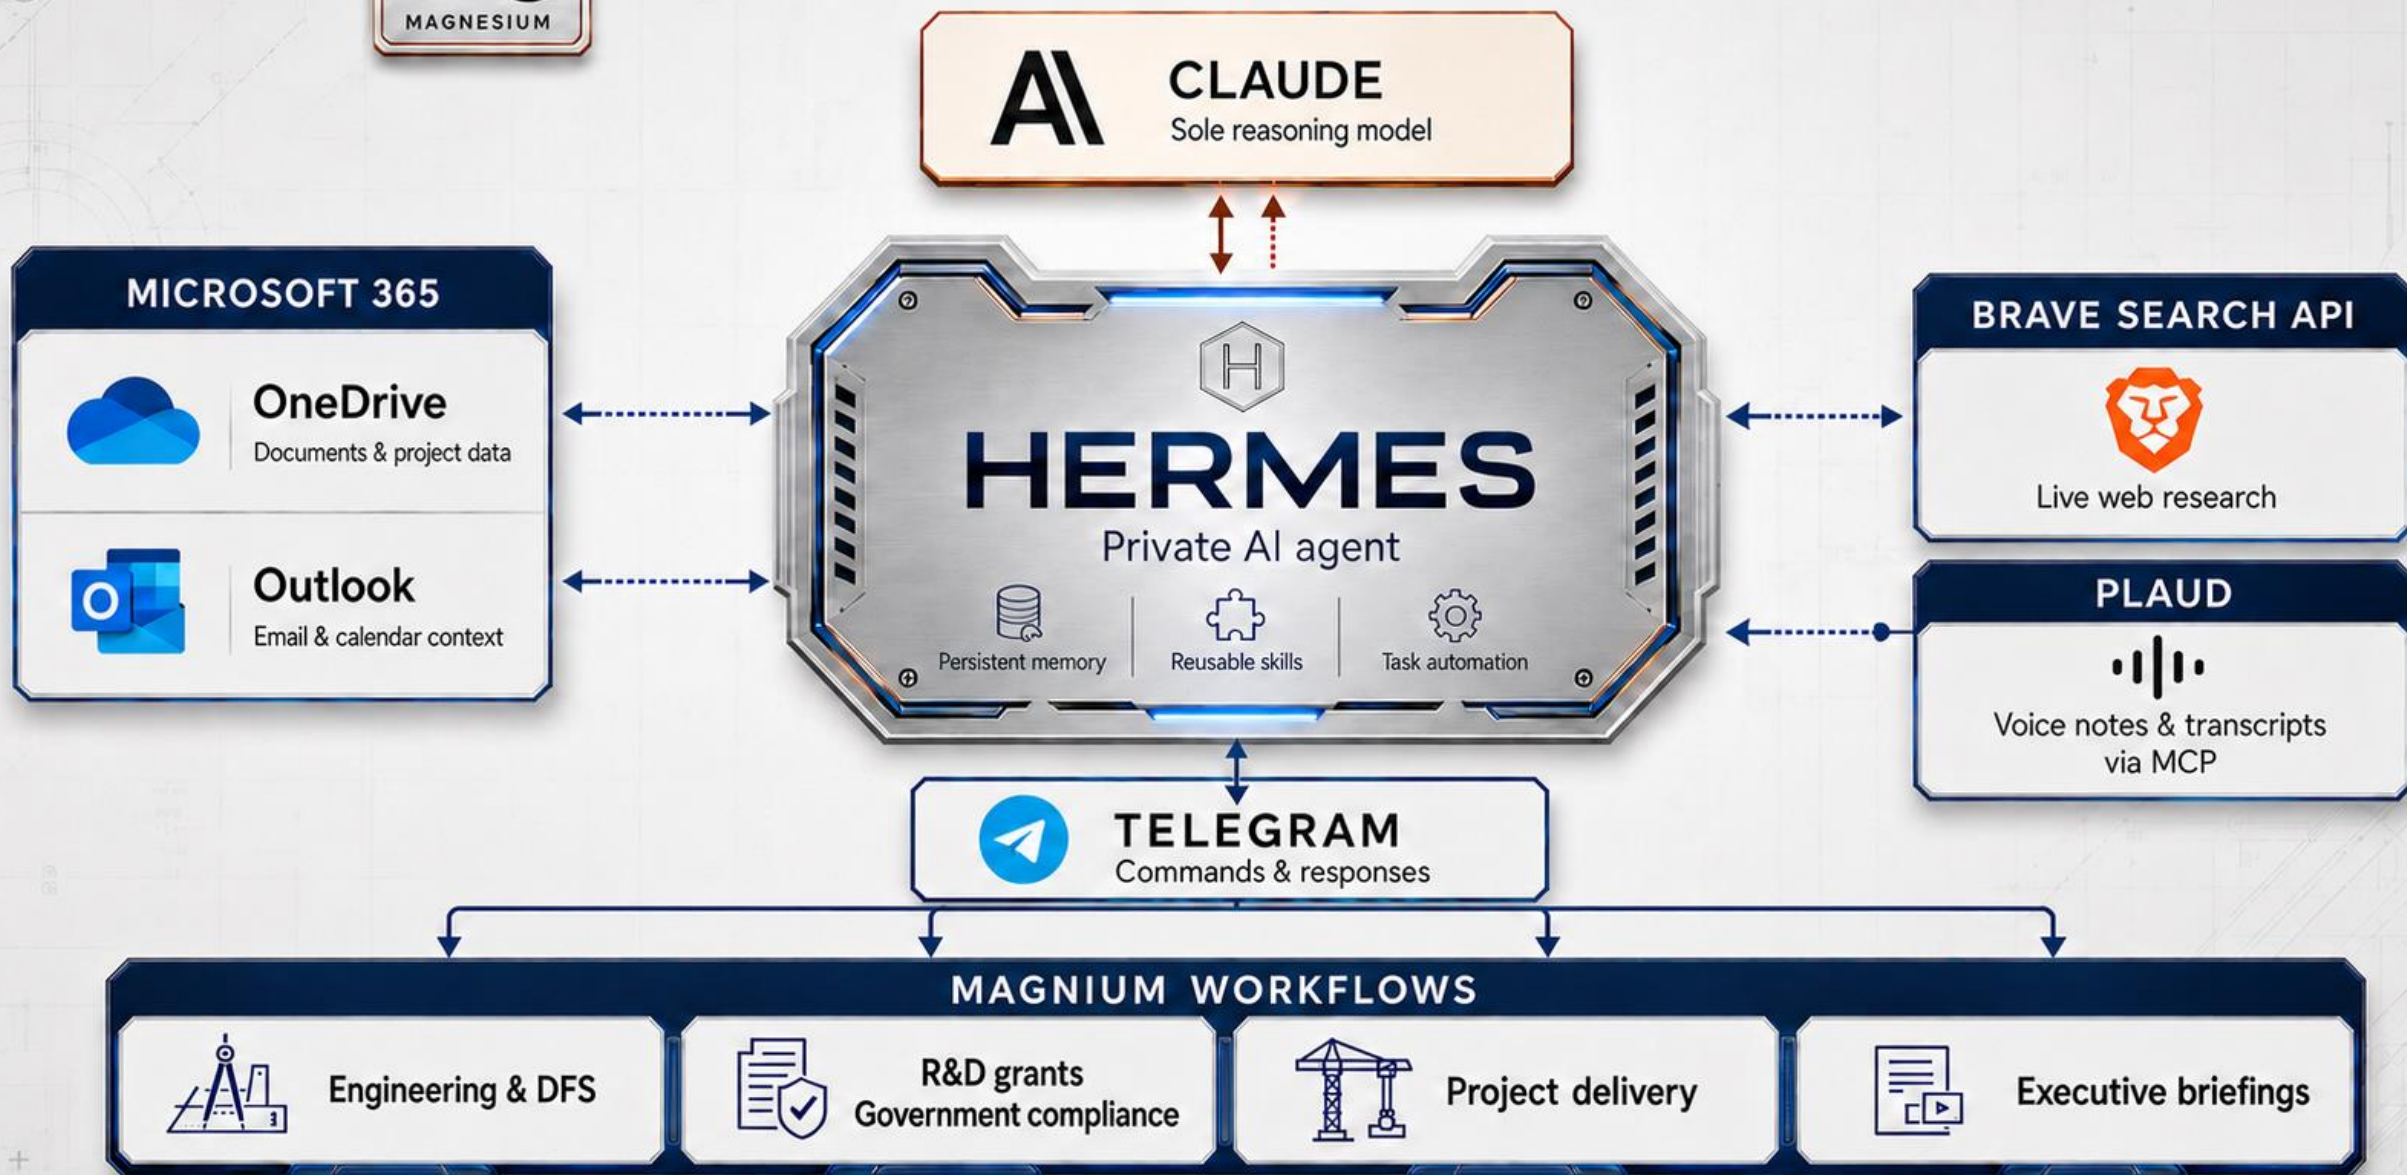
AI ECOSYSTEM (ATLAS)





MAGNIUM AUSTRALIA


HERMES AI OPERATING SYSTEM



Opportunity




Most businesses are using AI for simple tasks.
The leaders will use it to redesign how business works.



AI IS NOT ONE INDUSTRY. IT CHANGES SHAPE BY SECTOR




It is a sphere of many domains. The opportunity is to go deep in one domain you already understand.



FROM AI TO ONE DOMAIN: MINING



Even within one domain, expertise is still a subset within a subset.



The opportunity is not to know everything. It is to know your domain deeply.

STOP SEARCHING FOR THE AI UNICORN



Specialisation is coming. Maturity is not here yet.



Primitive generalist

AI is here

Medicine is here

Deep specialist ecosystem

Medicine then

Generalist phase



- One practitioner attempts everything
- Broad capability across many tasks
- Limited depth in any one domain
- Useful, but not yet specialized

Medicine now

Specialist ecosystem



Cardiology



Expert teams



Neurology



Expert teams



Pulmonology



Expert teams



Pathology



Expert team



Orthopedics



Expert teams



Many more




Expert teams

- Thousands of narrow expert specialties
- Deep capability in defined disciplines
- Clear roles, standards, and pathways
- This is where AI will mature over time


Broad claims of AI expertise should be treated with caution. Build the capability

EXPERTS KNOW THEIR DOMAIN. AI SEARCHES THE EDGES.

HUMAN VIEW



AI VIEW




AI SURFACES THE PATTERN. EXPERTS DECIDE WHETHER IT MATTERS.

AI HAS JAGGED INTELLIGENCE



It can be brilliant at one task and surprisingly weak at another that looks similar.



Often excellent

- Writes high quality code
- Summarises complex material
- Detects patterns in large datasets
- Produces strong first pass analysis



Can still fail badly


- Misreads context
- Struggles with some simple tasks
- Can give a confident but wrong answer
- Misses what a human finds obvious

That is why judgement still matters.

The skill is knowing where to trust AI, where to test it, and where a human must step in.

JAGGED INTELLIGENCE EXISTS BETWEEN MODELS TOO

Frontier LLMs are uneven in different ways. No single model is best at every task.




OpenAI GPT-5.5 family

- ✓ Best fit: general knowledge work, coding, tool use
- ⚠ Watch-out: not always the clearest leader in native multimodality



Anthropic Claude 4 family

- ✓ Best fit: careful analysis, writing quality, agentic coding
- ⚠ Watch-out: narrower multimodal positioning than Gemini



Google Gemini 3 family

- ✓ Best fit: multimodal, search-rich tasks, long context
- ⚠ Watch-out: less differentiated on polished enterprise writing

Jagged intelligence exists between models. Model choice should follow the task.

Illustrative executive view based on current official vendor materials. Real performance remains task-dependent.

CHOOSE THE TASK FIRST. THEN CHOOSE THE MODEL

The mistake I made here is treating AI as a vendor decision. The right approach is to route each task to the best model suited to it.

1. Department of War



Planning, red teaming, complex coding, analytical drafting

AI Anthropic Claude 4

Strong for careful reasoning, argentic coding, red teaming and long form analysis



Mission workflows, operational automation, research and data analysis

OpenAI GPT 5.5

Strong for tool use, workflow execution, research synthesis and structured business work



ISR, imagery, video, maps and document heavy intelligence review

Google Gemini 3

Strong for multimodal analysis, long context and evidence gathering

Defence should not standardise around one model. Planning, red teaming, coding, intelligence review and operational workflows each require different strengths.

2. Banking



Credit memos, policy drafting, risk write ups and complex judgement

AI Anthropic Claude 4

Strong for nuanced reasoning, structured drafting and careful communication



Analyst work, spreadsheets, operations and workflow automation

OpenAI GPT 5.5

Strong for data analysis, research, tool use and execution heavy workflows



Fraud, KYC, document review and evidence gathering

Google Gemini 3

Strong for multimodal files, long context and document rich review

Banking or any business should not pick one AI vendor and call it a strategy. Credit, compliance, fraud, customer service and automation all need different model capabilities.


The model is not the strategy. The operating model is.

Model strengths change quickly. Task routing, governance and data control matter more than model loyalty.

THE NEXT PHASE IS CONTEXT


Context is what turns AI from a clever tool into operating capability..

Standard chatbot



- Fast answers
- Useful general knowledge
- Limited business context
- Stops at the answer

Context + Harness (openclaw) + agent




- Uses your context
- Works across systems
- Completes tasks.
- Keeps humans in the loop where judgement matters

AI moves from answering questions to executing work when it understands the context.
The prompt gets an answer. The context gets the work done.

Clean Data Becomes Operating Context




A well organised Google Drive turns documents, data and domain knowledge into usable context for Atlas.



The prompt gets an answer. Clean context gets the work done.

The Mac Mini stores the permanent core layer

Only the most important business context is stored locally. Broader files, emails, and live sources remain connected and available on demand.




The model may change. The lasting advantage is the curated core business context stored on the Mac Mini, with broader business data still accessible on demand.

DATA IS THE FOUNDATION OF USEFUL AI



The model can only reason well from the context and data it is given.



Practical test

Before connecting AI to the business, ask:


- 01 Is the data clean enough to trust?
- 02 Can the system access the right sources?
- 03 Are permissions and governance clear?
- 04 Is the information current and complete?

Messy data turns a powerful model into a confident guesser – which is exactly what you do NOT want.

Core question: Is your data organised well enough for AI to use it?

Build the operating layer, not just AI literacy

The practical lesson from Atlas: advantage comes when trusted business context is connected to models, tools, workflows and human judgement.



AI rewards organisations that convert knowledge into execution.



Artificial intelligence driven gold discovery and targeting world class systems

FOR FURTHER INFORMATION

australia@ellipsis.com.au

+61 2 8024 5913