

GOVERNED WORKFORCE • RESEARCH LEAD

The research lead, working at frontier speed under zero-trust policy.

A governed agent that synthesizes STEM, market, and competitive research — home-routed to a self-hosted reasoning model, cleared to escalate to frontier, with every chain-of-thought and memory written under enforcement.

● Asphodel + Tiresias-ZT • ACME Corporation

THE STAKES

Frontier research is where governance usually breaks.

Synthesis work reaches for the strongest model, pulls in the widest context, and writes the most durable memory. That is exactly where ungoverned agents leak spend, mix tenants, and produce reasoning no one can audit.

WITHOUT GOVERNANCE

Every query defaults to the most expensive model. Chain-of-thought is discarded or dumped to shared storage. No signed record of what was reasoned over, or on whose authority.

UNDER TIRESIAS-ZT

The research lead homes to a self-hosted model, escalates only on cause, and writes CoT + memory into an RLS-isolated tenant behind a deny-by-default enforcement plane that signs every call.

The value of a research agent is its judgment. The point of governance is to make that judgment cheap, isolated, and provable.

THE WORKFORCE

367 governed agents, mapped to a real 325,154-principal org.

The research lead is one tenant inside a live zero-trust workforce — not a sandbox. Every agent is registered, policy-scoped, and enrolled to the same enforcement plane.


367 GOVERNED AGENTS	360 FLEET + 7 SHOWCASE	325,154 DIRECTORY PRINCIPALS	28 DEPARTMENTS • 10 REGIONS
---	--	--	---

Source: CP registry (367 agent tenants, postgres) • Samba-AD directory — 120,120 humans + 205,034 NHI identities, 626 roles.

POSITION

Fourth in the hand-off DAG — the synthesis hinge.


The 7 showcase agents run a 14-edge topological hand-off DAG. The research lead receives escalated context from the security track and produces the synthesis that data analysis and editorial build on. It feeds the research council.



Source: `personas.json` — `acme-runner-sched`, `runner.py` `schedule --interval 3600`. 14-edge topological hand-off DAG.

Every job is a governed round-trip, not a raw API call.

The research lead never talks to a model or a database directly. It mints a capability token; the enforcement plane decides; only on allow does the call reach its own isolated memory tenant.



API-KEY	X-Tiresias-API-Key + X-Capability-Token Both required on every governed call; policy is deny-by-default.	HS256
ISOLATE	Each agent is its own soul-svc tenant _memories & _soul_api_keys carry FORCE RLS; cross-tenant needs SET ROLE service_role.	RLS
SIGN	PEP signs the call into its audit chain SoulKey Ed25519 → shipped to CP; durable WAL at /data/audit-wal.	Ed25519

Source: governed request flow — one job, ARCHITECTURE-FACTS \$Governed request flow.

Home self-hosted. Escalate to frontier on cause.

The research lead homes to ollama_deep — deepseek-v4-pro on Ollama Cloud — and is cleared to escalate to frontier paid tiers when the work demands it. Routing favors free/self-hosted; the ceiling is available, not default.

TIER	MODEL · PROVIDER	ROLE	\$ / CALL	RELATIVE COST
Home	deepseek-v4-pro · Ollama Cloud	default synthesis	\$0.002	
Escalate	sonnet-4.5 · OpenRouter	hard reasoning	\$0.03	
Ceiling	opus-4.8 · OpenRouter	frontier only	\$0.19	

Source: config.json (9 tiers, 3 providers) · ROI board acme-roi-economics — observed \$/call, tzt_usage.

A competitive-synthesis brief, start to signed record.

This is what a single governed research job actually does — the same shape every fleet agent runs, roughly one to two times a day.

01	Domain brief arrives Market / STEM / competitive synthesis prompt handed off from the DAG.	intake
02	Tiered LLM call Homes to deepseek-v4-pro; PEP validates api-key + cap-token, allows.	\$0.002
03	Governed CoT capture POST /v1/cot/capture — reasoning written into the agent's own tenant.	/cot
04	Memory write POST /v1/memory/write with source_type + authored_by (ContentInspection enforce).	/memory
05	Usage report Metered to /cp/usage/report → tzt_usage → Grafana ROI board.	report

Source: `fleet_runner.py` — domain brief → tiered LLM → governed CoT + memory write → usage report.

The ceiling opens on a keyword, then closes again.

Home/ceiling/floor tiers are per-agent. Escalation is not the agent's whim — it fires on cause. And when the model call fails, the agent degrades safely rather than stalling.

ESCALATE

Keyword triggers — `+incident+breach+critical` — lift routing from the deepseek-v4-pro home toward the frontier ceiling for that job only.

DEGRADE

On `LLMError` the agent falls back to floor. Free-homed agents floor to `ollama_oss` (`gpt-oss:120b`) for a reliable degrade — never an unmetered frontier retry.

| Escalation is a policy event, not a default. Cost only moves when the work demands it.

Source: `personas.json` — per-agent `home/ceiling/floor` + keyword escalation; `fallback-to-floor` on `LLMError`.

Reasoning is captured, attributed, and isolated.

A research lead is only as good as what it remembers — and in an enterprise, only as safe as how that memory is walled off. Both are enforced at the data plane, not left to the agent.

CAPTURE	Chain-of-thought is a first-class write POST /v1/cot/capture — synthesis reasoning is retained, not discarded.	/v1/cot
ATTRIBUTE	Writes require source_type + authored_by ContentInspection in enforce mode rejects unattributed memory.	enforce
ISOLATE	Memory lands in one tenant only _memories FORCE RLS on Cloud SQL acme_soul; no cross-tenant read.	FORCE RLS
RECALL	Retrieval is tenant-scoped POST /v1/memory/search over the agent's own embeddings only.	/search

Source: soul-svc data plane — acme-soul-svc → PostgREST → acme_soul (RLS forced).

Deny-by-default, signed, and watched for drift.

Nothing about the research lead's autonomy weakens the control surface.
Every call is authorized, signed into an immutable chain, and monitored for behavioral drift with automatic containment.

8343 PEP • DENY-BY-DEFAULT	Ed25519 SIGNED AUDIT CHAIN	auto SOULWATCH → QUARANTINE	RFC5424 SYSLOG → LOKI SIEM
--	--	---	--

On allow, the PEP forwards verbatim to the tenant backend and signs the call into its SoulKey Ed25519 chain shipped to the control plane. SoulWatch turns behavioral drift into auto-quarantine against the shared postgres store.

Source: Enforcement Plane (acme-tiresias-zt, TZT_ROLE=proxy) • SoulWatch • Promtail:1514 → Loki.

The self-host-vs-API answer, measured per call.

Because the research lead homes to a self-hosted reasoning model and escalates only on cause, its steady-state cost sits at the bottom of the roster. The ROI board makes the gap explicit.

MODEL	PLACEMENT FOR RESEARCH LEAD	\$ / CALL	VS. FRONTIER CEILING
opus-4.8	ceiling · frontier only	\$0.19	
sonnet-4.5	escalate	\$0.03	
deepseek-v4-pro	home	\$0.002	
gpt-oss:120b	floor · reliable degrade	\$0.0015	
gemma-4-31b:free	free tier	\$0.00	

Source: ROI dashboard acme-roi-economics — observed per-call \$/1M-token economics from tzt_usage.

SUMMARY

A research lead you can hand the frontier — and still audit.

Home-routed for cost, cleared to escalate on cause, isolated at the row level, and signed on every call. It feeds the research council with synthesis that carries a provable record behind it.

ROUTE	Homes to deepseek-v4-pro at \$0.002/call; escalates to frontier only on keyword cause	tiered
GOVERN	Every call through PEP :8343 deny-by-default, signed Ed25519, watched for drift	zero-trust
ISOLATE	CoT + memory written to one RLS-forced soul-svc tenant, attributed on write	FORCE RLS
SCALE	One of 367 governed agents on the same live enforcement plane	367