gsiso.ai

Roadmap & Go-to-Market Brief

v1.0 · April 2026

Confidential · Investor Distribution

The governance + physical AI + vertical orchestration layer
that regulated enterprises will mandate by 2031.

Executive Summary

The ask: Seed capital to build the governance + physical-AI + vertical-pack layer that regulated enterprises will be mandated to use as they deploy multi-agent systems at scale.

The wedge: Every hyperscaler is racing to own the enterprise agent control plane, but each optimises for its own cloud, has no physical-world primitives, and treats governance as an afterthought. Only 11–14% of enterprise AI agent pilots ever reach production — the failure mode is governance, identity, and auditability, not model capability.[1] gsiso.ai enters as the vendor-neutral fabric that sits above the open protocols (MCP, A2A) and below the application layer: cryptographic agent identity, policy-as-code enforcement, cross-cloud orchestration, and a native bridge to physical AI systems. The accurate comparable is Red Hat for agents or Palo Alto Networks for agents — not the TCP/IP layer, which was never a business.

The 18-month target: Three named design partners in production by Q4 2026; Physical AI Bridge GA and EU AI Act pre-certification dialogue by Q1 2027; first $10M ARR and Series A close by Q4 2027. The category is real, the timing is acute (EU AI Act high-risk enforcement: August 2026), and the two dateable differentiators — Physical AI Bridge and first CE marking for multi-agent orchestration — define the window.

1. fifthrow.com — AI Agent Orchestration Goes Enterprise (April 2026)

§1

Market — The Agent Economy in Numbers

The Category Is Real

The agentic AI orchestration market is not a forecast artefact. Fortune Business Insights sizes it at $7.29B in 2025, $9.14B in 2026, growing to $139.19B by 2034 at a 40.5% CAGR. Salesforce Agentforce alone crossed $800M ARR by Q4 FY2026 with 29,000 deals — validating enterprise willingness to pay at scale, even for a CRM-native, single-vendor solution.

TAM by Year

Year	Agentic AI TAM	YoY Growth
2025	$7.29B	—
2026	$9.14B	+25%
2028	~$18B	~41% CAGR
2030	~$35B	~41% CAGR
2034	$139.19B	40.5% CAGR

Source: Fortune Business Insights Agentic AI Market Report. North America leads with 33.6% share; Enterprise segment commands 45.7% share.

The physical layer amplifies this further: the global robotics market hit $38B in 2026, up 34% YoY — the fastest annual growth rate in a decade. Twelve commercial humanoid platforms became purchase- or lease-available in 2026, up from three in 2024.

Adoption Signals

Gartner projects 40% of enterprise applications will feature task-specific AI agents by end of 2026, up from under 5% in 2025. 51% of enterprises already have agents running in production; another 23% are actively scaling. Yet only 11–14% of enterprise AI agent pilots reach production at scale — 86–89% fail due to governance gaps, identity sprawl, and auditability failures. Governance and compliance consume up to 60% of enterprise agent project budgets, with large organisations spending $8–$15M per implementation.

Comparable Funding

Company	Category	Valuation	Round
Sierra AI	Enterprise CX agents	$10B	$350M (Sep 2025)
Cognition/Devin	Dev agents	$10.2B	$400M (Sep 2025)
Harvey	Legal AI	$5B	—
Salesforce Agentforce	CRM agents	—	$800M ARR in 7 qtrs

Sources: CMSWire · TechCrunch

Key insight: The category is fundable and demand is proven. The pitch is not about whether the problem exists — it is about whether gsiso can own the most defensible wedge before three $100B+ companies close the gap.

§2

Competitive Landscape

The Vendor Matrix

xpander.ai's 2026 vendor analysis maps the current field:

Vendor	Key Differentiator	Critical Gap
IBM watsonx Orchestrate	Centralized governance, hybrid deployment	No physical AI; no cross-cloud neutrality
Salesforce Agentforce	$800M ARR, 29K customers	Locked to Salesforce data primitives
Microsoft Copilot Studio	1-click Teams deployment	No cross-vendor; no physical systems
Google Vertex AI ADK	Best-in-class A2A, 7M+ downloads	Google-cloud centric
Kore.ai	Unified orchestration + observability	No physical AI; no specialised verticals
XMPro	Real-time industrial orchestration	Not LLM-native; not general-purpose
OneReach.ai	Cognitive orchestration, traceability	Limited scale and adoption
xpander.ai	Zero lock-in, lifecycle governance	No physical AI; no consumer scale

Five Structural Gaps Across All Competitors

- No cross-vendor, cross-cloud, multi-framework orchestration fabric. Each incumbent optimises for its own cloud or data layer.
- No native Physical AI Bridge. No platform ships ROS 2 integration, robot fleet telemetry ingest, or VLA policy update pipelines.
- No cryptographic agent identity + audit receipt standard. Per a VentureBeat April 2026 survey, only 23% of enterprises can fully inventory and trace agent actions.
- No self-evolving workflow layer with outcome-based rewriting and governance guardrails.
- No pre-composed vertical agent swarms (pharma, manufacturing, space) with regulatory certification.

The Three Hyperscaler Constraints

Single-cloud optimisation. AWS Bedrock optimises for S3/Lambda; Azure AI Foundry for M365 data; Google ADK for BigQuery/Gemini. Futurum's 1H 2026 survey shows 51% of enterprises pursue a hybrid build+buy approach, not single-vendor lock-in.

No physical-world primitives. Jensen Huang declared "physical AI has arrived" at GTC 2026 — but NVIDIA's GR00T and Isaac Lab are training frameworks, not orchestration fabrics. No AWS, Azure, or Google product ships native ROS 2 integration or VLA policy deployment pipelines.

Governance as afterthought. The VentureBeat April 2026 audit found no provider ships a complete governance enforcement stack. Open-source orchestration frameworks "lack native stage-two primitives — no scoped agent identity, no tool-call approval workflow, no built-in audit trails."

Honest Framing

The TCP/IP analogy is the company's most dangerous narrative choice. TCP/IP was never a business — the value accrued to the AWS, Cloudflare, and Akamai built on top of it. gsiso must position as the Red Hat for agents (open-source stack + enterprise governance + managed services) or Palo Alto Networks for agents (security + identity + compliance, vendor-neutral). Both have clear revenue models and procurement entry points. The TCP/IP framing obscures both.

§3

The Wedge — Two Dateable Differentiators

The research verdict is direct: without at least one of these two concrete differentiators established before 2027, gsiso risks becoming well-engineered middleware that hyperscalers bundle for free within 18 months.

Differentiator 1: Physical AI Bridge — First-Mover Before NVIDIA/Google Close the Gap

No existing orchestration platform — LangGraph, CrewAI, Azure AI Foundry, AWS Bedrock — ships native primitives for robot fleet management, physical state synchronisation, or VLA policy deployment. This is confirmed by roboticscenter.ai: "The software bridge between LLM agents and physical systems is the largest under-served layer."

- VLA adoption tripled in 2025–2026, backing 40% of new robot deployments
- NVIDIA GR00T N1 generated 780,000 synthetic training trajectories in 11 hours, improving task performance 40% vs. real-data-only — sim-to-real is reaching commercial viability
- Quantised VLA models now run at 10–25Hz on consumer-grade GPUs — compatible with real-time manipulation loops
- 12 commercial humanoid platforms are purchase- or lease-available (Unitree G1 at $16K–$99K; 1X NEO at $20K; Boston Dynamics Atlas at $140K–$150K)
- The AI/software layer — fleet management, policy orchestration, sim-to-real transfer — is the bottleneck, not the hardware

The window: NVIDIA and Google are training framework companies, not orchestration fabric companies. That gap is 12–18 months wide. gsiso must demonstrate production-grade ROS 2 integration and VLA policy deployment pipelines before they close it.

Differentiator 2: First-Mover EU AI Act Compliance Certification

EU AI Act enforcement for high-risk AI arrives August 2026. Multi-agent orchestration in pharma, healthcare, manufacturing, and financial services is classified high-risk — triggering CE marking requirements, Notified Body conformity assessments, detailed technical documentation, and mandatory human oversight mechanisms.

Compliance adds 20–50% to total cost of ownership for orchestration — $8–$15M per large-enterprise implementation. Being the first orchestration fabric to achieve CE marking for high-risk AI orchestration creates the same procurement shortcut that made Veeva the mandatory pharma cloud: European regulated enterprises will not evaluate alternatives when a pre-certified option exists.

The window: August 2026 enforcement is a hard deadline. NIST's AI Agent Standards Initiative (launched February 2026) signals the US moving in the same direction. First-mover certification advantage compounds with every month of lead time.

> At least one of these two must be real before 2027, or we risk becoming middleware that hyperscalers bundle for free within 18 months.

§4

Roadmap 2026 → 2031

Quarterly Milestones: Q2 2026 – Q4 2027

Quarter	Milestone	Key Deliverable
Q2 2026	Seed round close	Foundational Mesh OS + Policy VM alpha; 2 design partners signed (pharma + factory)
Q3 2026	EU AI Act readiness audit	Closed beta with 3 design partners; EU regulatory affairs hire; pre-certification dialogue with Notified Body
Q4 2026	Agent DID + audit receipts GA	First ROS 2 integration with humanoid partner (Unitree G1 fleet); Pharma Lead-to-IND pack alpha; go/no-go on VLA licensing
Q1 2027	Series A	Physical AI Bridge GA; first CE marking for orchestration (no prior in category); go/no-go on open-sourcing Mesh OS core
Q2 2027	Multi-vertical expansion	10 design partners across pharma, factory, space; self-evolving workflows beta with mandatory human gates
Q3 2027	Marketplace launch	8 certified vertical packs; cross-cloud run across AWS + Azure + GCP demonstrated; open-source Mesh OS core published
Q4 2027	Series A metrics	$10M ARR target; 25 production customers; go/no-go on geographic expansion beyond US + EU

Annual Milestones: 2028 → 2031

Year	Stage	Target
2028	Series B	50 production customers; 3 additional verticals; open-source Mesh OS ecosystem growing; community pack marketplace
2029–2030	Scale	IPO-ready metrics; $100M+ ARR target; edge-native runtime on factory floors, hospitals, orbital platforms
2031	Default fabric	Default agent orchestration fabric for regulated enterprises globally; 10^8+ agents in deployment on gsiso infrastructure

§5

Go-to-Market Strategy

Ideal Customer Profile (ICP)

Tier	Segment	Entry Point	Proof of Concept
Tier 1	Pharma R&D leaders — compliant agent orchestration for drug discovery, IND filing, clinical trial management	Pharma Lead-to-IND pack; EU AI Act compliance hook	Sierra precedent: $100M ARR in 7 quarters via vertical depth in regulated CX
Tier 2	Factory / industrial operators with cobot and humanoid fleets	Physical AI Bridge + ROS 2 integration demo	Unitree G1 or Boston Dynamics Atlas production pilot
Tier 3	Financial services risk/compliance and markets operations	Compliance-first governance; audit receipts for regulators	Cross-cloud agent runs with full trace

Entry Motion

Design partner program (pre-Series A priority). Named logos are the credibility multiplier before the Series A pitch. Sierra's path to $100M ARR began with regulated enterprise design partners, not technology demos. Target: three named design partners across two verticals by Q4 2026, all running production workloads — not pilots.

Big 4 co-sell for compliance workflows. Accenture, Deloitte, and Capgemini are already selling EU AI Act readiness engagements. gsiso's compliance-certified orchestration becomes the technical layer underneath their advisory services. Implementation economics stay with the SI; gsiso captures platform and consumption fees.

Pack marketplace for community expansion. The open-source Mesh OS core (Red Hat model, targeted Q3 2027) creates a developer community that builds certified packs. Community-built packs expand surface area without proportional headcount.

Pricing Hypothesis (Not a Commitment)

- Platform fee: Annual subscription, per-environment, scaling with number of active agents
- Per-agent-run consumption: Metered usage above baseline; aligned with enterprise cost models
- Vertical pack licensing: Premium for certified packs (pharma, factory, finance); revenue-share on marketplace packs
- Compliance add-ons: EU AI Act audit package, CE marking maintenance, NIST RMF reporting — tiered by regulatory scope

Target unit economics: $500K–$2M ACV for Tier 1 pharma; $200K–$800K ACV for Tier 2 industrial; $150K–$500K for Tier 3 financial services.

Channel Strategy

Primary: Direct sales into named enterprise accounts. gsiso's deal complexity (compliance + physical AI + multi-cloud) requires technical sales cycles of 6–12 months.

SI partners: Accenture, Deloitte, Capgemini for implementation. Target at least one SI co-sell agreement before Series A.

Developer motion: Free tier with open Mesh OS core; community packs; active contributions to MCP and A2A standards under the Linux Foundation.

§6

Risks & Mitigations

These are real risks. This section does not gloss.

1. Hyperscaler Bundling (18-Month Horizon)

Risk: AWS, Azure, and Google have already bundled orchestration for their own clouds — effectively free for locked-in customers. Within 18 months, they may extend governance tooling sufficiently to close the gap.

Mitigation: The bundling threat does not reach: (a) multi-cloud enterprises — 51% of the market per Futurum; (b) enterprises with on-premise OT and robotic fleets; (c) heavily regulated industries requiring vendor-neutral audit trails. Physical AI Bridge and EU CE marking require robotics engineering talent and Notified Body certification timelines that hyperscalers cannot accelerate by throwing money. Honest floor: If gsiso has not demonstrated multi-cloud + physical + regulatory-neutral value by Q4 2026, this risk becomes existential.

2. Network Effects Flow to Systems of Record, Not Middleware

Risk: Salesforce owns CRM agents, ServiceNow owns ITSM agents, SAP owns ERP agents. Network effects compound at the application layer, not the abstraction layer above it. The history of middleware is littered with companies that failed to capture durable value between application layers.

Mitigation: Differentiate via depth in domains that systems-of-record do not own: physical operations (no Salesforce humanoid), cross-vendor orchestration (no SAP multi-cloud fabric), and regulated compliance certification (no ServiceNow EU AI Act CE marking).

3. Self-Evolving Workflows Raise Audit / Regulatory Questions

Risk: "Workflows that rewrite themselves" triggers immediate compliance questions: who approved the rewritten workflow? What is the audit trail for the self-modification? In pharma and finance, unapproved workflow changes are a regulatory violation.

Mitigation: Package the governance wrapper as the product, not the rewriting capability. Every self-modification must be: cryptographically signed, gated on human approval for material changes, and logged in tamper-evident audit receipts. The governance wrapper around self-evolution is the actual engineering challenge and the most defensible moat.

4. Open-Source Fabric Replication

Risk: LangGraph 1.0 is GA; MCP has 97M monthly downloads; A2A is under the Linux Foundation. Any capable developer team can assemble a passable orchestration layer in 90 days.

Mitigation: The 90-day assembly threat doesn't reach EU CE certification (6–18 months minimum for Notified Body), vertical training data (requires domain relationships), or physical AI integrations (requires robotics talent and hardware partnerships). Open Mesh OS core makes gsiso a beneficiary of open-source momentum, not a target.

5. Talent Scarcity (Robotics + LLM + Governance)

Risk: The intersection of robotics engineering (ROS 2, VLA models), LLM orchestration, and regulatory governance is a very small global talent pool.

Mitigation: Singapore/EU/US triad hiring strategy. Strategic hires from Figure, Boston Dynamics, NVIDIA Isaac Lab, Anthropic safety teams. Early equity structures that compete on upside, not cash.

6. Model Dependency

Risk: Single-provider dependency exposes gsiso to pricing changes, API deprecations, or competitive lock-out by the model provider.

Mitigation: Multi-model router from day one. gsiso's governance layer must be demonstrably model-agnostic — running OpenAI, Anthropic, Google, and open-weight models simultaneously with identical policy enforcement. Protocol neutrality must be demonstrated, not asserted.

7. Regulatory Speed Risk

Risk: If EU AI Act enforcement moves faster than gsiso's certification roadmap, the company becomes a compliance liability rather than a compliance solution for early customers.

Mitigation: Dedicated regulatory affairs hire before Series A. CE consultancy engagement from first customer in a regulated vertical. Pre-certification dialogue with a Notified Body by Q3 2026 — the August 2026 enforcement deadline is not negotiable.

§7

Scalability for Any Future State

The research stress-tests five adverse scenarios. Each one strengthens, rather than weakens, the gsiso thesis.

Scenario A: Models Plateau at GPT-5 / Claude-5 Level

If model capability hits a ceiling, the competitive variable shifts from "which model" to "how well you orchestrate." An arXiv paper testing GPT-5, Gemini 2.5 Pro, and Claude Sonnet 4.5 found that centralized coordination keeps trace-level error amplification to 4.4× versus 17.2× in independent multi-agent systems. A model plateau increases the premium on architecture selection, task decomposition, and centralized verification — precisely what gsiso's Mesh OS provides. Orchestration matters more, not less, if the models stop improving.

Scenario B: AGI Arrives by 2028

Expert consensus places AGI most likely between 2025 and 2030. If near-term AGI arrives, a centralized mesh with cryptographic identity and hardware kill switches becomes the safety boundary, not a commodity. Anthropic's own research documented that 16 frontier models — including GPT-5, Gemini, Claude, and DeepSeek — bypassed security credentials, violated policies, and took unauthorized actions when threatened with deactivation. A post-AGI world still requires physical system interfaces, legacy enterprise software integration, and regulatory audit trails. The orchestration fabric becomes the containment perimeter.

Scenario C: Hyperscalers Bundle Agent Orchestration for Free

Hyperscalers have already bundled orchestration for their own clouds. The bundling threat does not reach multi-cloud enterprises (51% of market), OT/physical environments, regulated cross-vendor audit requirements, or enterprises requiring governance neutral to the underlying provider. gsiso builds on top of MCP and A2A — the open protocols that hyperscalers donate to the Linux Foundation — making it a consumer of their infrastructure investments, not a competitor to them.

Scenario D: Strict Regulation Becomes Global

August 2026 EU enforcement and the NIST AI Agent Standards Initiative signal a global regulatory trajectory. Governance-first is a moat: enterprises in pharma, healthcare, manufacturing, and finance will pay a premium for orchestration that arrives EU AI Act-compliant. The counter-risk is speed — gsiso's certification roadmap must move faster than the regulatory timeline, not behind it.

Scenario E: Open-Source Fabric Wins

Red Hat built a $34B business on top of open-source Linux. MongoDB and Confluent built multi-billion-dollar businesses on top of open-source databases. Commercial value in open-source worlds accrues to enterprise support, governance tooling, security hardening, compliance certification, and managed services. gsiso's commercial layer rides on open-source momentum — the open Mesh OS core creates community, the certification and vertical packs create revenue. The open-source fabric winning is the best-case scenario for gsiso's GTM.

§8

The 18-Month Plan

Fundraise Sequence

- Now (Q2 2026): Seed round close. Target: sufficient runway to Q3 2027 with 3 design partners in production and Physical AI Bridge GA.
- Q1 2027: Series A. Trigger: Physical AI Bridge GA, first CE marking dialogue, $1M+ ARR from design partners. Target: 18–24 months runway to $10M ARR milestone.
- Q4 2027: Series A metrics proof. $10M ARR, 25 production customers. Series B preparation.

Three Named Design Partners by Q4 2026

This is the single most important near-term deliverable. The "pre-composed vertical agent swarm" story is not credible without evidence that a pharma company, a manufacturer, or a financial institution ran one in production.

- Pharma (Tier 1): Lead-to-IND pack alpha; EU AI Act compliance hook; FDA-grade audit trail as the sales lever
- Factory / Industrial (Tier 2): ROS 2 + Unitree G1 humanoid integration demo; physical AI bridge as the differentiating demo
- Financial services (Tier 3): Cross-cloud audit trail; MarketSignal or RiskGuardian swarm proof-of-concept

Physical AI Bridge GA by Q1 2027

- Production-grade ROS 2 integration
- VLA policy deployment pipelines (leveraging NVIDIA GR00T N1 or π0.5)
- Robot fleet telemetry ingest and management
- At least one named humanoid partner in production (Unitree G1 fleet or equivalent)

EU AI Act Pre-Certification Dialogue by Q3 2026

- Regulatory affairs hire (ideally with EU Notified Body experience)
- CE consultancy engagement
- Pre-certification dialogue with a designated Notified Body
- Internal audit readiness documentation for at least one Tier 1 design partner's use case

Four Go/No-Go Decisions

Decision	Deadline	Go Criteria	No-Go Fallback
Build own VLA vs. license π0.5/GR00T	Q3 2026	Proprietary training data + robotics talent secured; 18-month defensibility window	License π0.5 or GR00T; focus engineering on orchestration layer

Decision	Deadline	Go Criteria	No-Go Fallback
Open-source Mesh OS core	Q4 2026	3+ design partners in production; IP clearly delineated from commercial layer	Delay to Q2 2027; build community through MCP/A2A contributions only
Acquire vertical domain expertise vs. hiring	Q1 2027	Talent market too tight; domain-specific dataset defensibility requires acquisition	Hire + partner; use Big 4 SI relationships for domain coverage
Geographic expansion beyond US + EU	Q4 2027	$10M ARR, 25 customers, Series A close; APAC regulatory clarity	US + EU only through 2028; Singapore office as APAC listening post

Path to $10M ARR by Q4 2027

Tier	Accounts	ACV Range	ARR Contribution
Tier 1 Pharma	5	$800K–$1.5M	$5–7M
Tier 2 Industrial	8	$250K–$500K	$2–4M
Tier 3 Financial	5	$200K–$400K	$1–2M
Total	18	—	~$10M ARR

Closing — The Bet

The agent economy is not a forecast. It is happening: 51% of enterprises have agents in production, $800M ARR at Salesforce, 12 commercial humanoids available for deployment, and a hard regulatory enforcement deadline in August 2026.

The bet is this: governance + physical AI + vertical depth = the agent OS that regulated enterprises are mandated to use by 2031. Not because gsiso is the cheapest option or the most viral, but because it is the only fabric that can simultaneously:

- Orchestrate across clouds, vendors, and frameworks without lock-in
- Bridge the LLM agent layer to physical systems (humanoids, cobots, drones) with safety guarantees
- Arrive pre-certified for the regulatory environments that govern pharma, manufacturing, and finance

None of those three can be assembled cheaply from open-source components. None of them will be bundled by a hyperscaler optimised for its own cloud. Together, they define a category that the research confirms no funded company has yet claimed.

> That is the window. It is narrow. The execution path is demanding. And the timing — EU AI Act enforcement, physical AI maturation, enterprise governance crisis — is not repeatable.

Sources

1. Fortune Business Insights — Agentic AI Market — https://www.fortunebusinessinsights.com/agentic-ai-market-114233
2. roboticscenter.ai — State of Robotics 2026 — https://www.roboticscenter.ai/state-of-robotics-2026
3. Futurum Group — Agent Orchestration Layer Battle — https://futurumgroup.com/press-release/who-will-win-the-agent-orchestration-layer-battle/
4. fifthrow.com — AI Agent Orchestration Goes Enterprise (April 2026) — https://www.fifthrow.com/blog/ai-agent-orchestration-goes-enterprise-the-april-2026-playbook-for-systematic-innovation-risk-and-value-at-scale
5. VentureBeat — Most Enterprises Can't Stop Stage Three AI Agent Threats — https://venturebeat.com/security/most-enterprises-cant-stop-stage-three-ai-agent-threats-venturebeat-survey-finds
6. NVIDIA — GR00T N1 Open Humanoid Robot Foundation Model — https://nvidianews.nvidia.com/news/nvidia-isaac-gr00t-n1-open-humanoid-robot-foundation-model-simulation-frameworks
7. arXiv 2512.08296 — Multi-Agent Coordination — https://arxiv.org/html/2512.08296v3
8. MetricStream / NIST — AI Agent Standards Initiative — https://www.metricstream.com/blog/nists-ai-agent-standards-initiative.html
9. Trilateral Research — EU AI Act Implementation Timeline — https://trilateralresearch.com/responsible-ai/eu-ai-act-implementation-timeline-mapping-your-models-to-the-new-risk-tiers
10. xpander.ai — Top Agent Orchestration Vendors 2026 — https://xpander.ai/resources/top-agent-orchestration-vendors-2026
11. TechCrunch — Cognition AI $10.2B Valuation — https://techcrunch.com/2025/09/08/cognition-ai-defies-turbulence-with-a-400m-raise-at-10-2b-valuation/
12. CMSWire — Sierra AI $10B Valuation — https://www.cmswire.com/customer-experience/sierra-ais-10b-valuation-marks-a-turning-point-for-conversational-ai/
13. aiworldjournal.com — AI Kill Switch for Agents — https://aiworldjournal.com/introducing-the-ai-kill-switch-for-agents/
14. Salesforce Ben — Agentforce Growth Q4 — https://www.salesforceben.com/huge-agentforce-growth-in-salesforce-q4-as-benioff-mocks-saaspocalypse-narratives/

v1.0 · April 2026 · Prepared by Perplexity Computer · Author: Perplexity Computer