

Forum:[Hack Forums](#)

Thread: [\[CLOSED\] AGI TEAM - The Invisible Workforce. Zero Footprint. Total Execution.](#)

CYBERWAR WEAPONS - May 5, 2026, 7:44 pm

? AGI team ?

The Invisible Workforce. Zero Footprint. Total Execution.

Quote:is a high-level AGI orchestrator that replaces your entire overhead. It doesn't just "assist"—it operates.

You provide the instructions; it provides the result. From the first bid on UpWork to the final line of code and the invoice follow-up, AGI team acts with the precision, nuance, and professional "vibe" of a senior human staff. No generic AI responses. No "bot" patterns. Just work.

? TOTAL ANONYMITY

Integrated with advanced linguistic modeling to eliminate "AI-isms." Whether it's a proposal, a client email, or a code comment, the output is indistinguishable from a high-tier professional. Your clients will never know they aren't talking to a human. ?

END-TO-END SALES
Hooks directly into UpWork and CRMs. The AGI finds jobs, writes bespoke proposals, handles the "interview" chat, and closes the contract. It manages your leads while you focus on the vision. ?

AUTONOMOUS FULFILLMENT

A specialized suite for software development, marketing, and design. The AGI takes the project requirements and executes them in your repo or dashboard. It debugs, iterates, and delivers finished products. ?

LIVE OPERATIONS
The team is always online. It handles service requests, lead inquiries, and project updates in real-time. It manages its own calendar and meet invites exactly like a real agency would.

? HOW IT WORKS

Brief the Team: Set your goals, brand voice, and technical standards.

Grant Access: Connect UpWork, GitHub, Slack, or your preferred stack.

Go Dark: The AGI takes over. It sources work, fulfills orders, and maintains client relations.

Review Results: Log in to see completed milestones and revenue growth.

? SERVICE PLANS

CORE UNIT - \$449/Month

(Full Sales & Fulfillment Loop, 2 Active "Employees")

AGENCY SQUAD - \$1,299/Month

(Unlimited UpWork Bidding, 8 Active "Employees," Priority Processing)

THE SHADOW CORP - Contact for Pricing

(Full-Scale Infrastructure, Custom Human-Mimicry Fine-Tuning, Dedicated Server)

Payment Methods: Cash App / BTC / LTC / XMR / ETH

? CONTACT US

Discord: intelligence_general

Disclaimer: GHOST-TEAM is a tool for high-level business automation and technical simulation. Use responsibly and at your own risk.

[contract=tid]Start Contract[/contract]

Grand Hustle - May 6, 2026, 3:43 am

Interesting concept, but I've seen a lot of "AI agency" products that are basically prompt wrappers around existing APIs. A few technical questions to separate real autonomous infra from marketing:

1. What's the actual execution architecture?
 - Multi-agent orchestration?
 - DAG/task graph?
 - Event-driven workers?
 - Supervisor/reconciler layer?
2. How is long-term memory/state persisted across sessions?
 - Vector DB?
 - Relational state?
 - Episodic memory?
 - Context compaction strategy?
3. How do you prevent recursive hallucination/context poisoning between agents?
4. What models are actually being used underneath?
 - Frontier APIs?
 - Open-weight local models?
 - Routing/fallback strategy?
 - Cost optimization layer?
5. How are tools executed safely?
 - Docker/container sandboxing?
 - VM isolation?
 - Credential vaulting?
 - Scoped permissions?
6. How does the coding agent validate generated code before shipping?
 - Unit tests?
 - Static analysis?
 - CI/CD?
 - Execution sandbox?
7. How do you handle large multi-week projects with evolving requirements?
 - Repo indexing?
 - Semantic code search?
 - Architecture mapping?
 - Incremental planning?
8. How do you prevent infinite agent loops or deadlocks?
 - Watchdogs?
 - Iteration caps?
 - Supervisor agents?
 - Timeout recovery?

9. You mention eliminating “AI-isms” — what specifically achieves that?
 - Fine-tuning?
 - Style classifiers?
 - RLHF?
 - Post-processing pipeline?
 - Prompt ensembles?
10. What evals are you using to measure “human indistinguishability”?
 - Blind testing?
 - Detection benchmarks?
 - Conversion metrics?
11. How are client credentials stored and isolated?
 - Encrypted secrets?
 - OAuth scopes?
 - Audit logs?
 - Zero-trust model?
12. How do you defend against prompt injection via GitHub repos, client chats, or uploaded docs?
13. What's the rollback strategy if an agent corrupts a repo or sends a bad proposal?
14. What's the concurrency limit before latency spikes?
 - Queue architecture?
 - Worker autoscaling?
 - Rate-limit handling?
15. What part of the stack is genuinely novel vs existing frameworks like:
 - LangGraph
 - AutoGen
 - CrewAI
 - OpenHands
 - MCP toolchains
 - Claude Code

Would also love to see a full technical walkthrough:

UpWork discovery -> proposal generation -> negotiation -> repo execution -> QA/testing -> delivery
-> invoicing.

If you can answer these in dms or on here, I am definitely interested!