



BARTHOLOMEW AI • BTP v5.4.16

The In-Process Runtime Execution Gateway for Autonomous AI Agent Swarms

Seed Round: \$2,000,000 USD on a \$15,000,000 Valuation Cap Post-Money SAFE (\$0 Prior Dilution)

4,500+	< 35 us	2,500+	\$360,000
Active Developers (npm, PyPI, Open VSX)	In-Process Latency (50,000x faster than LLMs)	Metered Actions in Cryptographic Ledger	Active Pipeline ARR (12 Enterprise Pilots)

THE AUTONOMOUS EXECUTION GAP & THE BARTHOLOMEW GATEWAY

The Problem: LLMs have evolved from passive text generation into autonomous agents with terminal, database, and API execution authority. Giving models tool access creates the **Autonomous Execution Gap**: prompt-level guardrails fail determinism, while secondary LLM judges (e.g. LlamaGuard) introduce 1,500–2,500 ms of latency and double token bills. Traditional network firewalls cannot see in-process function calls.

The Solution: Bartholomew is an in-process execution gateway. Running directly inside the agent's memory space and edge proxy, it evaluates multi-language AST syntax invariants in **under 35 microseconds** (< 0.035 ms), halts destructive commands, scrubs leaked API keys in real time, and issues RFC 8785 Ed25519 cryptographic receipts for automated compliance.

FOUR CORE ARCHITECTURAL INVENTIONS

1. Sub-35us Polyglot AST Gating: Evaluates Bash, Python, SQL, JS/TS, Go, and Rust syntax trees in sub-35us before execution. Halts system wipes (rm -rf), SQL drops, and spend loops with zero network latency.	2. In-Flight Secret Vault: Scans tool calls and prompts in real-time. Detects and redacts OpenAI keys (sk-*), AWS tokens, GitHub PATs, and private keys before leakage to model providers or logs.
3. Ed25519 Merkle Attestation: Issues mathematically signed RFC 8785 receipts for every action. Bundled into local Merkle trees for instant SOC 2 Type II and HIPAA compliance audits.	4. Zero-Install Edge Gateway & M2M: Drop-in reverse proxy for Grok, Meta AI, Claude, and Muse bots. Features autonomous L402 Lightning micropayments (\$0.01/call) for machine-to-machine swarms.

COMPETITIVE ADVANTAGE MATRIX

Dimension	LLM Judges (LlamaGuard)	Perimeter Cloud WAFs	Docker / E2B Sandboxes	Bartholomew (BTP v5.4)
Latency	1,500 - 2,500 ms	30 - 100 ms	200 - 500 ms	< 35 us (Sub-0.035 ms)
Execution Seam	External Cloud API	Perimeter Network Edge	Container Hypervisor	In-Process Runtime + Edge
Determinism	Stochastic (Bypassed)	HTTP Codes Only	OS Container Boundary	100% AST Invariant Proof
Token Cost	Doubles Token Spend	Data Egress Fees	VM / Compute Billing	\$0 Token Spend (Local CPU)
Audit Evidence	Plaintext String Logs	Centralized Syslog	Container Exit Codes	Ed25519 Canonical Receipts

CAPITAL STRUCTURE, USE OF FUNDS & CONTACT

• **Capital Structure:** 100% founder-bootstrapped by Chief Architect Itsub Alemayehu with \$0 prior institutional equity raised. Accelerated by non-dilutive grant and compute jumpstart backing from **Google Cloud, Microsoft for Startups, and Amazon Web Services (AWS)**.

• **The Financing Ask:** Raising **\$2,000,000 USD** on a **\$15,000,000 Valuation Cap Post-Money SAFE**.

• **Capital Allocation:** **60% Low-Latency Systems Engineering (\$1.2M)** to scale compiler and eBPF kernel engineers; **25% Enterprise Compliance & Field Pilots (\$500k)** for formal SOC 2 Type II and closing 20+ enterprise pilots; **15% Ecosystem Expansion (\$300k)** for native integrations across Claude, Gemini, OpenAI Agents SDK, and Cursor.

• **Contact:** Itsub Alemayehu • Founder & Chief Architect • itsub@bartholomew.info • <https://bartholomew.info>