

🎯 APPLICATION TARGET: SENIOR APPLIED AI ENGINEER – DEVELOPER TOOLS & WEBCONTAINERS

TARGETED EXECUTIVE SUMMARY

Senior Applied AI Engineer specializing in **autonomous developer agents, in-browser execution runtimes (WebContainers / Bolt.new), and deterministic context reduction**. Deep hands-on experience building agents that reason over large codebases, generate surgical code diffs, and execute verified git checkpoints. Author of ``git_bolt_jules_gate.py``, an automated token-cleaning gate saving **46,000+ tokens per session in WebContainer environments**. Builder of ``judgeguard-gemma4-agent``, an autonomous software engineer calibrated for high-speed code repair on held-out repositories.

SPECIALIZED TECHNICAL STACK

- **Autonomous Coder Agents:** Call-Paced Agent Topology, Repository Traversal, AST Parsing, Multi-File Patching, Context Window Pruning, In-Browser Tool Execution.
- **Developer Runtime & Tooling:** WebContainers, Node.js / Browser Runtimes, Git Checkpoint Internals, Surgical Diff Generation, Docker, Linux Shell.
- **LLM Orchestration & Governance:** Frontier LLMs (Gemini, Gemma 4, Claude, Llama 3), Zero-Trust Execution Gates, Hallucination Prevention, Semantic Drift Auditing.
- **Languages & Core Frameworks:** Python, TypeScript, JavaScript, Go, Bash, LangChain, ChromaDB, FastAPI.

KEY FEATURED ENGINEERING & PROJECTS

Bolt.new & WebContainer Optimization Engine (``git_bolt_jules_gate.py``)

Lead Architect | GitHub: kizabgd123

- Engineered automated token audit and context reduction gate specifically tuned for **WebContainers** and **Bolt.new** workflows.
- Automated isolation of heavy dependencies (``node_modules``, ``.cache``, ``dist``, build logs), compressing LLM prompt context by **46,000+ tokens per run** and eliminating in-browser memory exhaustion.
- Implemented surgical diff algorithms that replace multi-thousand-line full rewrites with targeted AST-level replacement blocks, preventing code truncation.

JudgeGuard Autonomous Developer Agent v2.0 (Google Gemma 4 Benchmark)

Architect & Competitor | Kaggle Kernel

- Built a single-agent autonomous software engineer running Gemma 4 31B, evaluated on held-out, unseen code repositories under a strict 12-hour evaluation budget.
- Formulated a high-throughput **call-paced doctrine** (Calls 1-8: locate/read, Calls 9-12: mandatory edit, Calls 13-20: pytest verification, Call 20+: diff commit), achieving **43% faster turnaround** than multi-agent baseline architectures.
- Implemented resilient inline patch fallbacks on edit collisions and scratch isolation in temporary runtimes.

JudgeGuard Reality-Grounded Verification Engine Core Creator | GitHub: kizabgd123/judge-guard-core-master

- Designed a 4-layer safety gate framework that verifies AI agent actions before and after execution:
- Layer 00:** Hard regex block on destructive commands (``rm -rf /``, ``sudo``, permission leaks).
- Layer 000:** Pre-flight LLM gatekeeper validating task admissibility and completion tickets.
- Layer 3:** Semantic drift evaluation ensuring the agent remains strictly within scope.
- Layer 4:** Physical reality verification checking disk state, SHA256 file hashes, non-empty outputs, and exit codes.

Diskort — Asynchronous Go Developer Toolkit

Open Source Creator | GitHub: kizabgd123/Diskort

- High-concurrency automation and developer tooling framework written in Go (**123+ GitHub Stars**).

PROFESSIONAL EXPERIENCE

Senior Applied AI Engineer (Developer Tools & Agentic Systems) | Independent / Remote

2023 – Present

- Designed and deployed agentic code-generation and debugging pipelines that operate autonomously in restricted execution environments.
- Built telemetry layers tracking prompt tokens, completion tokens, latency, and cost estimates per agent interaction.
- Partnered with developer tool teams to optimize prompt caching, context reuse, and token-saving strategies.

Lead AI Automation Engineer | Systems & Tooling

2020 – 2023

- Automated developer workflows, synthetic testing (Datadog Synthetics), and CI/CD pipelines.
- Built CLI tooling and internal copilots reducing engineering debugging cycles by 60%.

HONORS, AWARDS & IMPACT

- Kaggle Notebooks Expert:** Ranked **#734 / 61,348+ worldwide (Top 1.2%)**.
- GitHub Developer Program Member:** 60+ repositories maintainer.