

Sean Kim

+1 (201) 937-7441 | taemin.kim0327@gmail.com | github.com/skarokin | in/kimtaemin | taeminkim.com

EDUCATION

Rutgers University—New Brunswick

Sep 2022 — Jan 2026

Bachelor of Arts in Computer Science

GPA: 3.86

EXPERIENCE

Palo Alto Networks

Feb 2026 — Present

Site Reliability Engineer

Santa Clara, CA

- Engineered a FedRAMP High AI SRE agent on Lambda/API Gateway with 12 MCP integrations, orchestrating subagents for parallel root-cause and triage which cut investigation time from ~30m to ~5m and MTTR by 47%.
- Hardened agent safety with Bedrock Guardrails, HITL, and RBAC, which masked sensitive data, caught 93% of hallucination cases, blocked unauthorized tool use, and gated high-risk tools through human approval flows.
- Centralized LLM observability in Datadog, adding monitors that reduced time to detect regressions by 93%.
- Architected a self-serve autoremediation platform with SNS, Lambda, DynamoDB, and S3, turning ad-hoc Python scripts into safe production automations, reducing MTTR by 78% across 42 Datadog monitors.
- Built a DynamoDB policy engine and Slack control plane, combining runtime config, human escalation, idempotent retries, atomic locks, and recurrence windows to block 100% of unsafe or conflicting executions.
- Guaranteed safe, auditable executions via 97% test coverage and Pydantic validation, while leveraging S3 audit trails and Datadog traces, logging, and custom metrics to enable sub-5 minute RCA for failed remediations.

CyberArk

May 2025 — Jan 2026

Site Reliability Engineer Intern

Newton, MA

- Prototyped an SNS and Step Functions auto-remediation pipeline to validate zero-touch alert resolution.
- Automated 3 Datadog alerts by implementing runbooks in Python, cutting the MTTR for those alerts by 95%.
- Built an AI SRE agent with Confluence, GitHub, and Datadog integrations to accelerate incident response.

PROJECTS

copium.dev | [GitHub](#) | copium.dev

Go, TypeScript, SvelteKit, PostgreSQL, Algolia, Docker, GCP, Cloudflare

- Built an internship management platform with SvelteKit and PostgreSQL, processing 3000 weekly applications.
- Scaled Algolia search engine indexes to 500 concurrent ops with Cloud Pub/Sub & Go consumers on Cloud Run.
- Delivered an 81% reduction in query latency for BigQuery data analytics by leveraging CQRS architecture.
- Implemented compensating transactions for consistency across 3 data stores with a 99.99% successful retry rate.

ref:note | [GitHub](#) | refnote.app

Go, TypeScript, Node.js, Next.js, Docker, Redis, GCP, Vercel

- Launched a Next.js collaborative note editor on Vercel, supporting 1000 concurrent collaborative sessions.
- Reduced deployment costs by 37% by optimizing Node.js WebSocket synchronization with Redis Pub/Sub.
- Deployed a Go backend on Cloud Run via Docker with sub-50ms cold starts for low-latency Firestore queries.

TECHNICAL SKILLS

- **Programming Languages:** Go, Python, SQL, TypeScript, JavaScript, Java, C
- **Libraries & Frameworks:** Node.js, SvelteKit, React, Next.js, PyTorch, spaCy, NLTK, NumPy, Pandas
- **Cloud & Devops:** AWS, GCP, Cloudflare, Terraform, Jenkins, Docker, Datadog
- **Platforms & Data:** Git, PostgreSQL, Nginx, RabbitMQ, Redis, Kafka