

Eric Nguyen

✉ contact@eric-n.com | 🌐 www.eric-n.com | 📄 [eric-nguyenit](#) | 🗣️ [c0debyeric](#)

Summary

Senior AI Cloud Platform Engineer building agentic platforms and production RAG systems on AWS Bedrock. 9+ years across IT, software, and cloud, including 4+ years architecting internal developer platforms at Fortune 500 (Royal Caribbean) and FedRAMP (Coalfire) scale, spanning 200+ AWS accounts and 50+ Azure subscriptions. Ship LangGraph multi-agent workflows, Bedrock Knowledge Bases (Aurora pgvector, Cohere Rerank), and Terraform-managed IDPs governed by HCP Terraform and Sentinel policy-as-code.

Certifications

- AWS Certified Generative AI Developer – Professional (in progress)
- AWS Certified AI Practitioner
- AWS Certified Solutions Architect – Associate
- HashiCorp Terraform Associate 004
- Microsoft Certified: Azure Administrator Associate (AZ-104)

Technical Skills

Platform Engineering & IDP: Backstage (Software Catalog, Software Templates, TechDocs), Self-Service Infrastructure, Developer Enablement

AI-Native Platforms: LangGraph, Strands SDK, CrewAI, AWS Bedrock (Claude Sonnet/Haiku), RAG (Bedrock Knowledge Bases, Aurora pgvector, Titan v2, Cohere Rerank v3.5, hybrid search, LLM-as-judge eval), MCP Protocol

Cloud Platforms: AWS (EKS, EC2, Lambda, S3, RDS, DynamoDB, IAM, Bedrock, Control Tower), Azure (VMs, VNets, Entra ID, Key Vault)

Databases: PostgreSQL (RDS), DynamoDB, MongoDB

Container Orchestration: Kubernetes (EKS/on-prem), Docker, Helm, ArgoCD (GitOps)

IaC & Backend Tooling: Terraform (Enterprise-Scale, 200+ accounts), HCP Terraform/Sentinel, Packer, Ansible, Python, Bash, PowerShell

CI/CD & Delivery: GitHub Actions (matrix builds, OIDC auth), GitLab CI, ArgoCD

Security & Governance: IAM, Zero Trust, SCPs, Checkov, Trivy, Policy-as-Code (Sentinel)

Monitoring & Observability: Prometheus, Grafana, Loki, OpenTelemetry, CloudWatch, Azure Monitor

Experience

Royal Caribbean, AI Cloud Platform Engineer

Miramar, FL
Sept 2025 – present
10 months

- Architected an agentic AI account-vending platform using LangGraph on AWS Bedrock (Claude), operating in Microsoft Teams with a custom tool suite that validates intent against live AWS Organizations state, auto-assigns CIDRs, and triggers HCP Terraform run pipelines with human-in-the-loop approval. Combined with Control Tower OU/SCP automation, collapsed account provisioning from 3+ days of ClickOps to under 10 minutes of self-service for 10+ engineering teams
- Built production RAG platform on Amazon Bedrock Knowledge Bases (Aurora pgvector, Titan v2 embeddings) with advanced retrieval. Claude Haiku handles fast query reformulation and sub-query decomposition, hybrid semantic+keyword search retrieves candidates, Cohere Rerank v3.5 reranks via cross-encoder, and Claude Sonnet produces grounded answers with citations. Includes a Comprehend-based PII sanitizer with DynamoDB audit trail. Serves 10+ engineers over 5,000+ internal docs via Microsoft Teams
- Built CI-integrated RAG eval harness with golden Q&A sets, retrieval metrics (recall@k, MRR), and LLM-as-judge scoring, gating Knowledge Base updates on quality regressions and surfacing per-component performance drift across embeddings, reranker, and generator
- Built 5-agent LangGraph "doc curator" pipeline (crawler → cosine+LLM dedup → cross-doc contradiction detection → quality scoring → supervisor) with human-in-the-loop review and Microsoft Teams Adaptive Card notifications, keeping the RAG corpus fresh and self-healing
- Built internal MCP (Model Context Protocol) server exposing AWS Organizations, Terraform state, and HCP run APIs as typed tools for downstream LLM agents, enabling consistent, auditable tool calling across the vending agent and future AI workflows
- Led org-wide IaC adoption across 5 cloud platforms (AWS, Azure, GCP, OCI, Hyper-V), bootstrapping HCP Terraform with centralized state, Sentinel policy-as-code governance, Entra ID SSO, and RBAC. Established the standardized delivery layer that downstream automation (including the agentic vending platform and AI workload provisioning) executes against
- Built parallel Terraform import tool in Python (asyncio, uv workspaces) on GitHub Actions matrix pipelines, discovering 45k+ live resources across 200+ AWS accounts and auto-generating import blocks + scaffolded HCL to migrate the org from ClickOps into version-controlled IaC

Coalfire, Cloud Engineer II

- Managed secure multi-account AWS infrastructure with AWS Organizations, Transit Gateway peering, SCPs, and RAM across production environments
- Automated multi-cloud provisioning (AWS and Azure) with reusable Terraform modules, cutting deployment time by 40%
- Deployed GitLab on Azure with Entra ID SSO and multi-stage CI/CD pipelines with Trivy security scanning, improving developer velocity by 30%
- Produced architecture diagrams and technical documentation in LucidChart, communicating infrastructure designs and security boundaries to client stakeholders

Remote
Jan 2025 – Sept 2025
9 months**Hash Studios LLC**, Cloud Engineer

- Deployed production applications on Amazon EKS with ArgoCD GitOps, built CI/CD pipelines with GitHub Actions, Trivy scanning, and blue/green deployments
- Automated infrastructure provisioning with Terraform, reducing deployment times by 50%
- Implemented Prometheus/Grafana observability stack with real-time metrics, alerting, and dashboards across production workloads

Plantation, FL
June 2024 – Dec 2024
7 months**I-Payout**, System Administrator

- Spearheaded hybrid cloud migration to AWS and Azure (20+ workstations). Designed Site-to-Site VPN and enforced PCI-DSS compliance
- Automated employee onboarding with PowerShell and managed Active Directory / Group Policy (GPO) for user lifecycle, access control, and security baseline enforcement, reducing setup time by 80%
- Built Python, Bash, and Ansible automation tooling to streamline IT operations and Linux patching, reducing manual intervention and configuration drift

Fort Lauderdale, FL
June 2022 – Oct 2024
2 years 5 months

Projects

Autonomous AI Agent Fleet (9 Agents on AWS Lambda)

- Deployed 9 autonomous AI agents on AWS Lambda using Strands SDK with Bedrock (Claude Sonnet/Haiku) and Gemini fallback. Built shared common library with model tier routing, DynamoDB dedup, Discord delivery, and CloudWatch observability. 100% Terraform-managed with reusable IaC modules and GitHub OIDC auth. Python monorepo with uv workspaces

Multi-Cloud Resource Inventory Platform

- Python service scanning 37+ AWS and 24+ Azure resource types org-wide using parallel processing, OIDC auth, and GitHub Actions automation with S3 upload and Teams notifications

Education

AA **Broward College**, General StudiesDavie, FL
Aug 2020 – May 2022