ROBERT HAFNER

ROBERT@TEDIVM.COM — 413.222.6909 — CHICAGO, IL

LinkedIn · GitHub · Open Source Portfolio · Blog · Terraform in Depth

PROFILE

Distinguished Engineer who turns AI from experiment into enterprise capability. Two decades of infrastructure and developer experience, now applied to building the platforms, security programs, and education that let organizations adopt AI safely at scale. Filed a patent on AI attack detection, authored *Terraform in Depth* (Manning), and publishes public guides like "Beyond the Vibes: Coding Assistants and Agents" to help developers move from experimentation to production. He's the engineer you bring in when AI adoption needs to move from pilot to production, securely and at scale.

EXPERIENCE

DISTINGUISHED ENGINEER

Comcast; Chicago, IL — January 2025 – Present

Promoted to Distinguished Engineer to lead AI platform strategy, where I built the infrastructure, security, and education programs that turned AI from experiment to enterprise capability.

- Formed and led the AI SDLC team as Tech Lead, designing the AI platform architecture combining LLM Gateways, Context Gateways, and agent-to-agent workflows, then presented to the CISO to secure executive buy-in.
- Launched the LLM Gateway with legal and security guardrails for enterprise-wide AI adoption, and designed MCP-based Context Gateway features for secure, governed AI data access.
- Designed a multi-agent AI framework with persistent workspaces and DAG-based task graphs, used to build AI-orchestrated developer tools and deep research agents for organizations such as legal.
- Built AI-orchestrated developer tools that made teams faster: an agentic security analysis tool that outperformed external pentesters, a language-agnostic test-generation tool, and an automatic documentation system following the Diataxis Framework.
- Designed, built, and launched the first Open Weight Model Cluster at Comcast, working across DevX, Legal, and Security to get it from concept to production.
- Filed a patent on detecting AI contextual rug-pull attacks via embedding models, establishing IP for novel AI safety techniques.
- Led security across the organization: joined the Securing Source Code and AI in the SDLC workgroups alongside GCS executives where I introduced the dependency standard, built the system used to evaluate supply chain attack impact, participated in incident response sessions that changed how we respond, and regularly joined threat modeling sessions for AI applications company-wide.
- Championed developer education across the organization: founded an AI security community, led monthly AI learning sessions, created AI skills and bootstrap guides, gave Knowledgefest talks, and taught some of the highest-rated internal classes.

SENIOR PRINCIPAL ARCHITECT

Comcast; Chicago, IL — February 2023 – January 2025

After years of startups, I wanted to test what I could do at a larger company. The result: infrastructure that saved hundreds of thousands of developer hours, security that kept the platform safe, and programs that built trust across the engineering community.

- Accelerated developer velocity across the organization: reduced ML pipeline onboarding from 3–5 days to 5 minutes, achieved parity between public and hosted CI/CD runners (10,000+ hours saved), and built a self-service developer support bot in under one week (~8,000 hours saved).
- Eliminated long-lived credentials by leading the OIDC initiative connecting CI/CD pipelines to cloud and secrets management, saving an estimated 40,000 developer hours per year.
- Built the infrastructure tooling that scaled the organization: centralized Terraform state management at a fraction of commercial cost, a shared module library with 28+ reusable modules, a standardized module template, and a progressive automation framework for 99.99%–99.999% uptime.
- Kept the platform secure by discovering two CVSS 9.9 vulnerabilities in the CI/CD pipeline, designing policies that prevented impact from a widespread GitHub Actions exploit, and building the repository analytics platform for full code visibility.
- Developed an automated PR framework for CI/CD adoption reporting and targeted modernization outreach, integrated with the internal developer portal.
- Built developer trust through a knowledge-sharing community, mentoring across multiple engineering organizations, and serving as SME for the top-rated internal course on infrastructure as code.
- Joined the OPEA Technical Steering Committee for the Linux Foundation's Open Source Enterprise AI initiative, shaping open source enterprise AI tooling.

PRINCIPAL ENGINEER

Aptible; Remote — March 2022 – November 2022

Joined Aptible to apply my background in AI, security, and development toward making cloud infrastructure safe and accessible for developers of all skill levels.

- Architected and led development of the next-generation Aptible Platform, a SaaS interface that let customers provision AWS resources securely without managing infrastructure.
- Built the platform's core: account management via AWS Organizations, a Terraform abstraction layer for exposing cloud resources as platform primitives, and a Connection system for secure, explicit resource linking.
- Created open source tooling including a Python Terraform wrapper, and fixed a critical authentication bug in the Aptible Managed Terraform Provider to reignite customer adoption.

PRINCIPAL ENGINEER

Explosion AI; Remote — July 2021 – March 2022

Left Rad AI to join Explosion AI, where I helped resuscitate a stalled product by refactoring the codebase, building out the cloud infrastructure, and establishing the processes the team needed to move forward.

- Resuscitated a project stalled for three years by refactoring the codebase and refactoring key components for faster development.
- Built the AWS environment from scratch with an IaC-first approach, including dozens of reusable Terraform modules, AWS SSO, and multi-account strategies through AWS Organizations.
- Established the company's first proper development lifecycle: staging environments, deployment pipelines, a knowledge base, and a structured proposal process that gave the team a repeatable path from code to production.

PRINCIPAL ENGINEER

Kenna Security (acquired by Cisco); Chicago, IL — April 2021 – June 2021

In just two months, eliminated recurring data export failures entirely, halved export times while also using cheaper infrastructure, and defined the company's Terraform adoption standards. Left when Kenna was acquired by Cisco.

PRINCIPAL ENGINEER

Rad AI; Berkeley, CA — March 2020 – April 2021

Transitioned from VP of Engineering to Principal Engineer as the company scaled, trading people management for focused work on platform architecture, security, and the infrastructure that powered our product.

- Architected the Product, Research, and QA platforms with disaster recovery as a first-class concern. The entire platform could be relaunched in under four hours, critical because radiologists depended on it to maintain patient care workloads.
- Scaled the MLOps infrastructure to production: expanded DGX training clusters with InfiniBand, built the full inference platform, and designed mock models for testing without GPUs.
- Created the infrastructure tooling that let the company move fast, including reusable Terraform modules for rapid prototyping and a CI/CD pipeline with automated rollbacks for safe, rapid deployments.
- Acted as Security and Privacy Officer, building out the compliance program and guiding the company through SOC2 Type 2, HIPAA+, and penetration testing.

VICE PRESIDENT OF ENGINEERING (FOUNDING ENGINEER)

Rad AI; Berkeley, CA — April 2018 – March 2020

As VP of Engineering and founding engineer, I built the organization from the ground up: hiring, research infrastructure, product platform, and the internal tools that made it all work.

- Joined as the founding engineer and hired across full-stack, ML, data science, and Windows development.
- Designed the first research platform, enabling HIPAA-compliant ML research across on-site DGX and AWS Cloud GPUs. Critical for stretching our first funding round without blocking model training.
- Architected the SaaS platform behind our first product, building an API and analytics layer capable of multiple deployments per day with less than an hour of total downtime over three years and the reliability that gave developers confidence to move fast.
- Created internal research and validation tools from scratch, including model storage, NLP experiment validation for radiologists, and a browser-based annotation system, because no off-the-shelf options existed for our use case at the time.
- Partnered directly with sales to convince cautious hospitals and radiologists to share data with a startup, addressing security and privacy concerns that were blocking deals in a highly regulated environment.

VICE PRESIDENT OF ENGINEERING

Vicarious AI; Union City, CA — November 2014 – April 2018

Led engineering and infrastructure for a deep learning research company, building the team, the platforms, and the culture that let researchers work without bottlenecks.

- Built the engineering team, hiring top talent and supporting company-wide recruiting efforts.
- Rebuilt infrastructure with configuration management and infrastructure as code, then architected hybrid-cloud HPC clusters for infinite scaling, a headless rendering farm for synthetic data, and distributed GPU training.
- Eliminated infrastructure bottlenecks by building self-service tools that let researchers and developers provision, configure, and scale clusters on their own.
- Audited and hardened company infrastructure, managing third-party security audits and improving defenses across all systems.

STARTUP ADVISOR AND NOMADIC DEVELOPER

Vacation Time; The World — January 2014 – October 2014

Took a break to travel the world, advise startups, and contribute to open source. Organized a 45-person Burning Man camp, spent a month cat-sitting in Paris, and rebuilt my personal datacenter.

VICE PRESIDENT OF IT

Malwarebytes Corporation; San Jose, CA — October 2008 – January 2014

Brought on as a consultant to rescue infrastructure during a traffic surge, then stayed through multiple promotions as the company scaled. Built and led four departments (Server Side Development, Malware Intelligence, Dev/Ops, Corporate IT), architected analytics and update platforms serving millions of daily downloads, and managed the IP blacklist powering the product's network protection — all while keeping infrastructure secure under constant attack. Scaled the organization from one person to over 25 worldwide.

DEV/OPS CONSULTANT AND COFOUNDER

SolunaNet / Soluna Enterprises; Chicopee, MA — December 2005 – April 2010

Started my career cofounding two web hosting and Dev/Ops firms. Managed hosting for 600+ sites, performed security audits and emergency incident response, built fault-tolerant LAMP applications, and led web development projects for clients ranging from startups to established businesses.

SKILLS

Languages	Python, JavaScript/Node, PHP, Bash, HCL/Terraform, Go — Familiar with TypeScript, Java, Perl, C, Scheme
AI / ML / MLOps	LLM Gateway design, MCP Gateways, multi-agent AI frameworks, DAG-based task graphs, RAG architecture, model training and inference platforms, DGX cluster management, hybrid-cloud ML environments, prompt engineering, AI security, vLLM, HuggingFace
Infrastructure	Terraform / OpenTofu (module design, state management, custom providers, testing, GitOps), Puppet
Systems	Linux (Ubuntu/Debian), MacOS, Puppet/PuppetDB, HPC (slurm, htcondor), NVIDIA GPU, DGX, VMWare ESXi
AWS	EC2, S3, Kinesis, SageMaker, IAM, VPC, Route53, EFS, RDS, Redshift, DynamoDB, Lambda, SQS, Comprehend, API Gateway, ELB/ALB/NLB, CloudWatch, CloudTrail, GuardDuty, SSM, SNS, Lustre FSX, AWS SSO/Organizations, ECR, EKS, EventBridge, STS, Secrets Manager
Orchestration	Docker/OCI, ECS (EC2 and Fargate), Kubernetes/EKS, Nomad, multi-architecture image builds
Data	MySQL, PostgreSQL, MongoDB, Redis, Elasticsearch, Hadoop, ActiveMQ, SQS, Kafka, ORM
CI/CD & DevOps	GitHub Actions, CircleCI, Jenkins, Git Flow, OIDC authentication
Security	HIPAA, SOC 2 Type 2, Vulnerability Assessment, Risk Assessment, AI Security, Supply Chain Security
Networking	Firewalls, IPSec and SSL VPNs, Subnetting, NAT, DNS, DHCP, InfiniBand/Mellanox
Development	Agile Processes (scrum, kanban), API Design, Continuous Integration, Scalability

EDUCATION

WORCESTER POLYTECHNIC INSTITUTE Worcester, MA — 2004-2006

SPRINGFIELD TECHNICAL COMMUNITY COLLEGE Springfield, MA — 2006-2007

REFERENCES

AVAILABLE ON REQUEST.