

William J. Gilleran

Platform Engineering • Kubernetes • AI Infrastructure
Norfolk, MA
billgilleran.com • linkedin.com/in/billgilleran

Profile

Platform engineer with a decade across enterprise infrastructure, Kubernetes, observability, DevOps, and internal developer platforms. I build reusable capabilities that improve delivery and operability, then validate them against live systems through telemetry, failure analysis, and recovery. My current focus is production-minded infrastructure for AI workloads and agent-assisted operations.

What I Build

Platform	Operations	AI Infrastructure
Kubernetes / k3s	Python & Go	vLLM
GitOps / Argo CD	OpenTelemetry	GPU Inference
CI/CD	Prometheus / Grafana	MCP
Internal Developer Platforms	Loki / Tempo	Agent Tooling
Identity & Secrets	REST APIs	RAG & Vector Search
Gateways & Networking	Automated Testing	Persistent Memory
Cloud Infrastructure	Runbooks & Recovery	Evaluation & Guardrails

Flagship Builds

The Hive — Platform Engineering Lab

Nine-node k3s platform with 25+ applications reconciled through Argo CD, Cilium networking, Longhorn and NFS-backed state, Prometheus/Grafana/Loki/Tempo telemetry, Authentik identity, Vault and External Secrets, and a 102 TB ZFS storage foundation.

Alpha-5 — Private AI Infrastructure

RTX 5090 system with 32 GB of GPU memory serving local models through vLLM. Operates alongside LibreChat, Open WebUI, RAG, vector search, memory, and MCP tooling to exercise capacity, telemetry, persistence, and safe agent access.

TechBearing — Operator OS

Approximately 14 Go and Next.js services exploring live operational context for humans and agents. Automated ingestion maintains current state; MCP over SSE provides a typed tool boundary; durable storage preserves history beyond any model session.

Professional Experience

Port | Technical Success Manager | Jul 2025 – Present

Partner with enterprise engineering organizations adopting internal developer platforms. Built API-driven tooling for health scoring and recurring-issue detection; led technical discovery across Kubernetes, CI/CD, identity, APIs, and cloud infrastructure; and translated field evidence into product direction. Contribute to Port's open-source CLI. Increased platform adoption from 7% to 18% for one strategic account, supporting a customer-initiated renewal, and helped another reach 60%.

RapDev | Datadog Engineer | Feb 2022 – Jul 2025

Developed maintained Python automation for repeatable enterprise migrations, integrated observability into production cloud and Kubernetes environments, and diagnosed distributed-system failures across infrastructure, networking, APIs, and telemetry pipelines.

Connected2Fiber | Support Engineer → DevOps Engineer | Jul 2018 – Feb 2022

Progressed from the company's first support engineer into DevOps, building cloud infrastructure, automation, and operational tooling while turning recurring user problems into durable engineering improvements.

LogMeIn | Technical Account Manager | Feb 2017 – Jul 2018

Served as the technical bridge between enterprise customers and engineering teams, translating production problems and business requirements into actionable technical work.

Earlier: Technical Support Representative, LogMeIn • Technical Support Analyst, Bullhorn (2014 – 2017)

Technical Foundation

Ten CNCF certifications across platform engineering, cloud native, and observability; CKA and LFCS in progress.

Platform Engineering: CNPA • CGOA • CAPA • CBA

Cloud Native: KCNA • KCSA • CCA • KCA

Observability: OTCA • PCA • Datadog APM & Distributed Tracing • Datadog Log Management

Applied AI: Google Generative AI Leader (GAIL)

Cloud: AWS Cloud Practitioner • Microsoft Azure Fundamentals

Portfolio

billgilleran.com/fit/platform-engineering — role-specific evidence and system case studies

Education

B.A. Communications — Western New England University