Danylo Pilovets

w.pilovets@gmail.com linkedin.com/in/danylo-pilovets Website: danylopilovets.com Kyiv, Ukraine
+380 (95) 530-95-93

Platform & MLOps Tech Lead | Senior DevOps / SRE Engineer

High-impact Infrastructure Lead with 5+ years of end-to-end experience across the entire systems stack from bare-metal architecture and cloud platforms to production AI/ML ecosystems. Versatile tech lead who has built, scaled, and maintained every layer of modern infrastructure: Kubernetes orchestration, distributed data stores, high-throughput streaming, GPU compute, GitOps CI/CD, and enterprise observability. Proven track record of taking full ownership of complex, highly ambiguous engineering challenges, bridging the gap between Data Science, Software Development, and Platform Reliability.

Skills

Engineering leadership, Kubernetes, Helm, Docker, Linux, Terraform, SQL/NoSQL, Python, Bash, AWS/GCP, CI/CD, Nginx, Prometheus, Victoria Metrics, Grafana, Loki, Ansible, ArgoCD, FluxCD, Vault, Git, Airflow, Arize Phoenix, Langfuse, LiteLLM, DeepEval, vLLM, Kubeflow, VertexAI, LangGraph

Experience

AI Platform Architect / Lead MLOps at NDA(FinTech) June 2026 - Present

- Delivered a one-year exclusive contract in five months as the only engineer on the project: designed and shipped a production AI platform on GKE from an empty GCP project to 28 GitOps-managed applications serving live workloads
- Single-handedly architected and built “AI-Enabled Internal Developer Platform”. Unified Argo CD, Prometheus, and Alertmanager into a multi-team IDP featuring cluster-wide health metrics, real-time deployment-to-alert correlation, and direct deep-links to logs/dashboards. Integrated in-place Alertmanager management (silencing with audit comments) and an embedded business AI assistant. Designed a plug-and-play configuration template enabling teams to self-onboard their applications and clusters, providing a centralized ecosystem directory with seamless auto-routed links to team-specific Vaults, Grafana dashboards, and internal tooling.
- Engineered agentic execution workflows using LangGraph and consolidated LLM access via an in-cluster LiteLLM gateway with Vertex AI (Workload Identity) and Redis-backed rate limiting, removing raw API keys from applications.
- Authored and published tasks-as-code a git-native backlog framework enabling deterministic task execution and automated context alignment for Human and AI-agent collaboration across 500+ tasks with zero context drift.
- Deployed Observability-as-Code (Prometheus, Grafana, Loki, Alloy) featuring a custom 3-tier noise-reducing alerting ruleset and automated Slack routing.
- Enforced zero-trust security using HashiCorp Vault (KMS auto-unseal), External Secrets Operator, Entra ID SSO with fine-grained RBAC, and internal VPN-only ingress alongside dynamic scoped pod terminals for developers.
- Engineered reusable GitLab CI/CD pipeline templates (8 shared templates, automated OPA policy enforcement, Helm schema validation, and doc-link integrity checks) to standardize build, test, and release cycles across all engineering teams
- Navigated the GKE AI platform through complete ISO/IEC 42001 certification compliance, leading architectural alignments with enterprise tech leads and enforcing strict AI management standards across cluster security, data governance, and operational workflows.

Skills: Kubernetes, Docker, GCP, Vertex AI, Terraform, Helm, ArgoCD, LiteLLM, LangGraph, Langfuse, pgvector, Docling, GitLab CI/CD, Vault, Entra ID SSO, PostgreSQL, Redis, Kafka, Qdrant, Prometheus, Grafana, Loki, Alloy, TypeScript (Hono, React), Python, Bash.

DevOps Engineer at Wix, Kyiv, Ukraine **December 2025 - June 2026**

- Managed and supported large-scale Kubernetes environments running business critical workloads.
- Led migration of monitoring infrastructure from Chef to Ansible, standardization and automation.
- Performed seamless migrations of large Victoria Metrics TSDB instances, deprecating legacy storage clusters and introducing new environments without downtime, data loss or target loss.
- Designed and implemented new monitoring environments, data sources, grafana dashboards, observability workflows.
- Participated in incident response, troubleshooting and RCA.
- Automated monitoring infrastructure provisioning and configuration using IaaC Terragrunt and Terraform.
- Developed custom exporters and custom metrics to improve visibility into platform and application health.
- Designed, developed and maintained Helm charts for monitoring and observability services deployed in Kubernetes

Skills: Kubernetes, Helm, Grafana, VM, AWS, Python, Bash, Chef, Ansible, Java, Terraform

MLOps Engineer at Playtika, Kyiv, Ukraine **May 2023 - April 2026**

- Supported machine learning workflows on Kubernetes with GPU-based workloads in production and training environments.
- Designed and maintained custom Helm charts for machine learning infrastructure and services.
- Implemented a robust log and data rotation strategy across numerous services, preserving terabytes of free disk space and preventing system downtime due to storage issues.
- Built custom alerting and monitoring systems tailored to the needs of ML and data services.
- Worked with a wide range of NoSQL and SQL databases, ensuring efficient data storage, replication, and seamless integration with distributed systems — while maintaining data integrity and preventing data loss across various failure scenarios.
- Ensured 24/7 availability and reliability of Kubernetes clusters, proactively monitoring, troubleshooting, and maintaining infrastructure to support critical production workloads.
- Participated in on-call rotations, promptly responding to incidents and minimizing downtime across production environments.
- Led post-incident reviews to uncover systemic issues, improve observability, and prevent recurrence of critical failures.
- Tackled complex, ambiguous issues that others avoided, driving root cause analysis and delivering long-term infrastructure solutions.
- Optimized Kubernetes resource usage by reducing request over-provisioning by approximately 60% across three different clusters, which significantly improved cluster efficiency and reduced operational costs.
- Migrated multiple operators and applications from kubespray setup to the Palette platform via Cloudstack, ensuring more consistent deployments, streamlined lifecycle management, and increased overall reliability.

Skills: Kubernetes, Helm, Aerospike, Hadoop, MySQL, Kafka, Airflow, Kubeflow, RunAI, Nginx, Prometheus, Grafana, ArgoCD, FluxCD, Jenkins, Ansible, Cloudstack, MCP, vLLM, Arize Phoenix

SRE Tech Lead at Oxtech, Tel Aviv, Israel **April 2021 - May 2023**

- Debugging applications based on linux.
- Managed and maintained bar-metal Kubernetes. (also worked with Rancher)
- Migrated projects between Kubernetes clusters, resulting in a 20% cost savings by leveraging existing vendor-rented hardware.
- Created custom Helm charts.

- Developed and implemented monitoring architectures, fostering a culture of observability-driven development, reducing Kubernetes alerts by 80%, and making the system more stable and user-friendly for Kubernetes newcomers.
- Created and maintained GitLab CI/CD pipelines for continuous integration and deployment, ensuring efficient and automated workflows.
- Provided technical assistance in various DevOps-related services, supporting backend developers and SRE colleagues with infrastructure automation, CI/CD pipelines, and monitoring tools.
- Worked with various databases, ensuring efficient data management, optimization, and seamless integration with applications.
- Automated cluster-wide infrastructure deployments using FluxCD, ensuring GitOps-driven consistency, reducing manual operations, and speeding up environment provisioning.
- Optimized Kubernetes resource usage by reducing request over-provisioning by 40%, which significantly lowered cluster load and improved cost efficiency.
- Implemented GitOps practices to streamline and automate the deployment process.
- Migrated CI/CD execution from bare-metal GitLab Runners to Kubernetes-based runner pods, increasing scalability, fault tolerance, and simplifying runner maintenance.

Skills: Kubernetes, IaaC SaltStack, Docker, Gitlab CI/CD, Redis, MongoDB, RabbitMQ, Vault, Helm, Linux, Bash, Prometheus, Grafana, Loki, FluxCD, ArgoCD

DevOps Engineer at Atlassian, Remote, Ukraine **January 2020 - April 2021**

- Worked with cloud infrastructure (primarily AWS, with partial GCP usage) to support internal platforms and services.
- Supported and improved internal CI/CD pipelines, ensuring stable and predictable delivery workflows across multiple teams.
- Collaborated with engineering teams to maintain and troubleshoot Kubernetes-based services, focusing on reliability and performance rather than full ownership.
- Contributed to infrastructure optimization initiatives, including resource tuning and basic cost visibility improvements.
- Participated in maintaining workflow orchestration (Airflow) within containerized environments.

Skills: Kubernetes, IaaC Terraform, Docker, Helm, Ansible, Airflow, AWS/GCP, ArgoCD, Grafana, Bash, Python, Kafka

Education

CENTRAL UKRAINIAN STATE UNIVERSITY **2018-2024**
M.S.

Pet projects

Horizontal (Github**:** fintech-event-pipeline-aws) - many decentralized services on a single layer, connected through events. Breadth of the platform.

Vertical (Github: mlops-view) - a single domain (LLM), but built end-to-end across all layers of the stack: cloud → platform → ML pipeline → observability. Depth of the stack.