Yuri (Iurii) Trukhin

Director of Infrastructure · Cloud, GPU & Developer Platform Leadership
Amsterdam, Netherlands
yuri@trukhin.com | trukhin.com | blog.trukhin.com | linkedin.com/in/trukhinyuri

Summary

I lead global infrastructure across CloudLinux and TuxCare, covering hybrid data centers, public cloud, multi-site networks, build systems, observability, and security compliance. One recent change moved I/O-heavy QA builds from shared Ceph storage onto local enterprise NVMe, cutting their build time to roughly a third on hardware we already owned. Before CloudLinux, I led Docker, Kubernetes, Dev Containers, and terminal product work on JetBrains IDEs, and earlier built the cloud platform business at MTS from its first cloud engineering hire to a 140+ person organization with direct cloud P&L and EBITDA accountability.

Experience

Director of Infrastructure | CloudLinux (serving CloudLinux & TuxCare) | Aug 2025 - Present | Amsterdam, Netherlands (remote)

Reporting directly to the CloudLinux CEO, directing global infrastructure across hybrid data centers, public cloud, multi-site networks, CI/CD, observability, security, and incident response.

- **Organization & 2026 strategy:** wrote the 2026 infrastructure plan, split reactive operations into platform, storage, IaaS, and automation teams with clear SLOs and incident reviews, and brought in leads for networking, databases, and Kubernetes.
- **Storage tiering & bare-metal KubeVirt:** benchmarked I/O-heavy QA builds and moved them off shared Ceph onto local enterprise NVMe (roughly 3x faster elapsed build times on existing hardware, using stateless rebuilds if a local disk fails); started a bare-metal KubeVirt evaluation.
- **Multi-site networks & DR:** upgraded multi-site network links and routing topology (5x inter-site throughput), moved BGP and IPsec configuration into version control, verified sub-30s automatic failover in live link-cut tests, automated DR restore checks, and unified logs, metrics, and traces on Grafana, Loki, Mimir, and Tempo.
- **FinOps & AI guardrails:** with Finance and the CEO, set up 12-month capacity forecasts, 1-year cloud commitments, competitive hardware RFQs, AWS cost attribution, and guardrails for AI coding tools (sandboxes, test gates, key isolation, token limits).
- **Interim Chief Security Officer (dual role, Sep - Nov 2025) & audit delivery:** kept 24/7 SOC operations staffed, redesigned the CSO role, and hired the permanent CSO; rolled out mandatory 2FA across core infrastructure portals; my team delivered the infrastructure evidence for our SOC 2 Type II examination on schedule within existing budget; mapped controls against ISO 27001, GDPR, DORA, NIS2, and EU CRA; and designed ephemeral build runners generating signed per-build provenance toward SLSA Build Level 3.

Product Lead (IntelliJ Platform) | JetBrains | May 2023 - Aug 2025 | Amsterdam, Netherlands (relocated Sep 2023)

- **Cloud, containers & AI workflows:** led product direction for Docker, Kubernetes, cloud, and AI workflows on IntelliJ Platform IDEs (used by 15M+ developers), using user interviews and product telemetry with engineering, design, and GTM to grow Docker and Kubernetes feature adoption 3x while NPS rose 25 points in my product area.
- **Dev Containers & rebuilt IDE terminal:** owned WSL, Docker, and Dev Containers tooling (contributed merged specification change devcontainers/spec#381), documented where modern CLIs failed in the legacy IDE terminal, and led the product and business case for the rebuilt terminal architecture shipped as default in 2025.2.

Head of Cloud Development Center | MTS Web Services (#CloudMTS / MWS) | Oct 2018 - May 2023

First hire in MTS's new cloud engineering department (*Progression: Product Lead 2018-2020 -> Head of Engineering Department 2020-2021 -> Head of Cloud Development Center, Platform & Services Engineering 2021-2023*).

- **Organization & cloud P&L:** hired and led 140+ people across product, platform engineering, and operations, owning cloud pricing, budgets, and annual revenue and EBITDA targets. MTS's 2020 annual report ranked #CloudMTS #2 by IaaS revenue in its home market; within 30 months of launch, the cloud business passed EUR 40M+ in annualized cloud revenue run rate (EUR equivalent) across 13 cloud services.
- **Commercial V100 GPU service (2019-2020):** launched MTS's commercial NVIDIA DGX-1 Tesla V100 and vGPU service in July 2019 (1 PFLOPS FP16 Tensor Core aggregate peak) and scaled fleet capacity 4x to 4 PFLOPS FP16 peak by November 2020, owning pricing and customer deployments.
- **0-to-1 platform & commercial ops:** brought early multi-tenant architecture ideas into Cloud Hub and Cloud Billing, authored Open Cloud Standard v1/v2, and set up how customers bought, activated, paid for, and received support across compute, storage, Kubernetes, IoT, and AI/ML services.

CTO (2017-2018) · Earlier Cloud Strategist & Evangelist (2013-2017) | Infobox (InfoboxCloud) | Nov 2013 - Oct 2018

- **Cloud growth & multi-DC expansion:** owned cloud strategy, product development, and direct cloud P&L/EBITDA; launched the first IaaS, PaaS, and SaaS lines, tripled annualized cloud revenue run rate over 18 months, expanded across three availability zones including an Amsterdam data center (InfoboxCloud reached #4 globally among Jelastic PaaS hosting partners), and launched Office24 (hosted office service), with Infobox receiving 2x Microsoft Regional Hosting Partner of the Year (2015, 2016).

Researcher to Lead Software Engineer & Architect | CNIP GIS | 2007 - Nov 2013

- **Distributed geological data systems:** designed a corporate geological data bank in Java and .NET and led 10 engineers building distributed analysis plugins that accelerated oil-and-gas reserve calculations.

Independent R & D

Personal, non-commercial software architecture projects where I test infrastructure and developer-tooling ideas directly in code.

- **TRUKHIN #CloudRING & OCSv3:** open cloud portability architecture exploring workload movement across providers without single-vendor lock-in (**pkg/dataact** export tooling, provider-neutral **cloudringctl** GitOps/upgrade/rollback, Go **sdk/ocsv3** and zero-dependency Rust **ocs-sdk-rs**, bare-metal KubeVirt experiments); business case defended during non-degree Executive MBA coursework at RAU (Yerevan, 2022-2023).
- **MetaBot & Frontier (MetaBot Apr 2025 · Local AI workspace):** MetaBot began on April 18, 2025 as a Python prototype automating personal finance workflows. Frontier is a separate local engineering workspace built afterwards to run coding agents in isolated Git worktrees inside bubblewrap/seccomp sandboxes, paired with a SQLite task coordinator, read-only-by-default integration gate, negative-control regression checks, and human review before release.

Upstream Code, Writing, Education & Recognition

- **Merged Upstream Code:** Linux kernel mainline **ac9de62b1902** (authored arm64 default **ARCH_MMAP_RND_BITS** ASLR hardening, **Reviewed-by: Kees Cook**); **devcontainers/spec#381** (contributed); **ceph/ceph#72706** (authored BlueStore allocator fragmentation-score calculation fix); **vllm-project/aibrix#3078** (authored Kubernetes LLM inference routing state bug fix).
- **Writing & Talks:** Personal engineering blog **blog.trukhin.com** (BYTEBURST notes on running coding agents); JetBrains Blog ("Using Dev Containers in JetBrains IDEs", 2024, author; "JetBrains Terminal: A New Architecture", Apr 2025 architecture announcement, product lead); Debugging Kubernetes with IntelliJ IDEA (2024, presenter); JavaOne Speaker (2013).
- **Education:** Tver State Technical University - 5-Year Engineering Specialist Degree (Specialist Diploma, Computer Systems and Networks, 2010); Postgraduate coursework (non-degree, 2010-2014); Executive MBA coursework (non-degree, RAU Yerevan, 2022-2023); VZMSh Correspondence Mathematics School (secondary program, 2005).
- **Recognition:** 2x Microsoft Regional Hosting Partner of the Year (awarded to Infobox, 2015, 2016); Microsoft MVP (2011-2012); JavaOne Speaker (2013). References available upon request via yuri@trukhin.com.