

Cedric Specht

Hamburg, 22527 • +49 (176) 20204039 • cedric@specht-labs.de

linkedin.com/in/cekienzl • specht-labs.de • github.com/cedi

Staff Engineer, Platform Infrastructure and SRE

I am a Staff Site Reliability Engineer who writes Go and stays close to production. My work covers Kubernetes operators, distributed systems, observability, deployment safety, incident response, and cloud infrastructure across AWS, Azure, Google Cloud, and OCI.

At Celonis, I build reliability and AI platform infrastructure. Recent work includes a Go operator for declarative SLOs, a multi-cloud inference gateway, and a company-wide Go CLI for managed access to coding agents.

My roles combine hands-on implementation with architecture and cross-team technical direction. I also mentor engineers, review designs, and speak publicly about Kubernetes, SRE, observability, and distributed systems.

Career Experience

Staff Site Reliability Engineer, Celonis, Hamburg, Germany

10/2025 until present

Staff SRE and technical lead in Celonis' Reliability and Developer Experience (RDE) organisation, working closely with Platform Infrastructure and Enterprise Security. Own the reliability platform, SLO automation, on-call automation and production access, and the reliability of the company's AI inference gateway, across 50+ production Kubernetes clusters, roughly 240 catalogued services and 20,000 ArgoCD-managed applications released through Kargo.

- **Contract-Driven SLO Generation:** Wrote the framework that treats a service's OpenAPI or gRPC spec as its reliability contract. Teams annotate an endpoint with the latency and availability they commit to, and the Go Kubernetes operator I built turns that into per-service Datadog SLOs and burn-rate monitors, with no hand-written monitor configuration. Owned it end to end, from concept through formal design review and implementation to GA in June 2026. It is how Celonis intends to hit 90% SLO coverage on Tier 0/1 services in FY27, and the same SLIs gate automated rollout progression in Kargo, so a release only advances when the service still meets its contract.
- **Framework Design and Delivery:** Wrote the design documents behind the SLO and on-call frameworks, took both through formal review, then wrote the SLO operator and shipped the on-call operator to production.
- **Reliability Targets and Incident Command:** Turned a four-nines directive into a staged path, three nines first from today's two and a half with four as the north star, so teams have a reachable next step. Command incidents, own the postmortems, and drove the cross-repo fixes after a Sev-0 so it cannot recur.
- **SLI/SLO Platform on Backstage:** Evolving the contract-driven SLOs into a fully formalised platform on Roadie Backstage, where SLIs are catalog resources with their own lifecycle. Separating SLI generation from SLO instantiation takes reliability-as-contracts beyond API specs, so SLIs can be generated automatically from contracts of any kind and coverage grows without per-team effort, while the SLO operator, Kargo rollout gates and dashboards consume the same entries, so a new consumer no longer means a change to the operator.
- **AI Inference Gateway:** Took the company's LiteLLM-based gateway from a near-empty repository to the production platform that fronts the engineering organisation's agentic coding LLM traffic across AWS Bedrock, Azure AI Foundry, Google Vertex AI and OCI. Own its reliability, with Enterprise Security as the service owner. Added autoscaling, per-engineer spend attribution over OpenTelemetry, and workload-identity federation in place of static credentials, and sized Azure as the second frontier-model provider against 24 hours of production traffic. Engineers get frontier models through one governed path, and the company gets provider redundancy, one security review, and a cost figure per engineer instead of an estimate.
- **Coding-Agent CLI:** Built the Go CLI every Celonis engineer uses for managed access to Claude Code, Codex, OpenCode and Pi, and took it from first build to company-wide GA in five days. It handles Entra ID sign-in, stores credentials in the OS keychain, reconciles config without clobbering user plugins, and self-updates from signed releases on macOS, Linux and Windows. Within weeks it carried 85% of all agentic LLM traffic from engineers, which ended per-engineer setup guesswork and removed the reason to bypass the gateway.
- **On-Call Automation:** Replaced hand-managed Terraform for OpsGenie on-call teams with a Go Kubernetes operator, which treats the Backstage service catalog as the source of truth and reconciles teams, schedules and escalations from it into the alerting backend. Teams declare who responds and when. That made the switch to Datadog On-Call a change of backend, and I designed the Datadog-native framework that completes it.
- **Just-in-Time Production Access:** Working with Enterprise Security to end standing production write access without slowing incident response. Building a Go plugin for Teleport that grants access automatically to responders on an active Datadog incident, scoped to the incident's declared environments and revoked when it resolves, with a rules DSL over incident status, responder role and severity. Responders never wait on an approver, a compromised credential reaches nothing outside a live incident, and no accidental change can land outside the incident's declared scope.

Lead Site Reliability Engineer, Databricks (neon.tech), Hamburg, Germany

06/2025 to 09/2025

Joined the [neon.tech](#) team just days before the acquisition by Databricks to lead a team of 7 engineers with a diverse skillset across reliability and infrastructure engineering during the post-acquisition integration of Neon into Databricks. Driving organisational change, technical migration, and strategic alignment with Databricks' mature platform ecosystem. Representing operational and infrastructure needs in cross-functional leadership forums to ensure stability, maintainability, and clear team purpose during the transition. Role concluded following the post-acquisition reorganisation of the team.

- **Designing and leading the organisational split** of the former Neon SRE/Infra team into two focused teams (Infrastructure & Production Engineering), defining roles, responsibilities, and ownership boundaries for improved execution and clarity within the new Databricks organisation.
- **Authoring a detailed integration plan** and design documents covering migration of observability, GitHub infrastructure, and ZTNA access to existing Databricks solutions, ensuring a successful integration into Databricks.
- **Creation, maintenance and prioritisation** of an infrastructure tech debt backlog to ensure focused improvements during the migration phase and beyond.

Senior Site Reliability Engineer / Technical Lead, Microsoft Azure, Hamburg, Germany

02/2022 to 05/2025

Technical Lead in Azure's Safe Change infrastructure SRE Team, responsible for chaos engineering, resiliency validation, and release infrastructure harmonisation. Led the modernisation of Azure's release infrastructure, migrating 60+ repositories and 600+ pipelines, increasing deployment reliability and speed across multiple critical customer-facing services including, among others, Azure Cosmos DB, Log Analytics, Web Apps & Function Apps.

- **Chaos Engineering & Resilience Validation:** Designed and implemented Chaos Engineering experiments to validate system failure hypotheses covering 80% of high-impact critical customer scenarios and improve resilience strategies and built synthetic monitoring and business validation testing to proactively identify and mitigate reliability risks.
- **Organised multiple internal learning sessions**, developing a 9-part self-guided onboarding tutorial as part of the SRE Academy, enabling new engineers to onboard 75% faster to the new release system.
- **Leadership and Team Management:** Technical lead & Scrum Master for my immediate team of 5 engineers, responsible for setting technical direction, mentoring, and defining strategies and goals for the team as well as the broader department, serving as the technical lead for a newly formed team within the Safe Change Infrastructure SRE organisation, and supporting multiple program managers and teams from across three other SRE organisations in bootstrapping new SRE engagements.
- **Cross-Org Collaboration and Stakeholder Engagement:** Partnered with 10+ service teams across Azure to help them migrate to the new release system, contributing high-quality pull requests to their repositories as best-practice examples driving down change related outages by 20%.
- **SRE Playbook, Wiki & Tech Talks:** Core contributor and committee member for the Azure SRE Playbook, authoring one new SRE pattern with 3 sub-patterns and reviewing 3 further major patterns into it. Maintained the cross-organisation Azure SRE Wiki and spoke at Azure SRE Tech Talks on reliability, deployment strategies, and platform engineering.
- **Recognition & Awards:** Azure Reliability Quality Star - Leadership Excellence Award for sustained high-quality contributions to Azure's engineering culture and reliability improvements.

Tech-lead Manager Kubernetes SRE, German Edge Cloud, Hamburg, Germany

07/2020 to 01/2022

Built and led the Kubernetes SRE Team: Established and scaled a remote team from two to 6 highly skilled SREs, taking full ownership of the company's Managed Kubernetes Platform, spread across 3 availability zones and hosting 50+ customer clusters. Ensured only high-quality changes made it into production by reviewing code, design documents, and architecture changes daily, implementing state of the art GitOps tooling and observability, resulting in a 75% reduction in change related outages over 12 months.

- **Incident & Change Management:** Developed and implemented new incident, change, and problem management processes, improving reliability and operational efficiency, enabling an average 10 minute time-to-engage and reducing time-to-mitigation by several hours on average through more streamlined and efficient incident management processes and standard operating procedures.
- **Cross-Functional Collaboration:** Worked closely with the Service Management team to improve incident response, change reviews, and operational excellence as well as the Infrastructure, OpenStack, and CEPH Storage teams, ensuring seamless integration and optimised performance across compute, storage, and networking layers resulting in 10% increased storage throughput and decrease in etcd commit latencies driving customer satisfaction.
- **Platform & Product Leadership:** Took on the Product and Platform Owner role, shielding the team from unnecessary business complexity while aligning priorities with company strategy and CTO directives.

- **Cloud-Native & Open Source Advocacy:** Fostered a culture of open-source collaboration, contributing improvements back to the cloud-native community and positioning the company's offerings within CNCF certification programs.

Site Reliability Engineer 2, SharePoint Online, Microsoft, Dublin, Ireland

01/2019 to 03/2020

Running Live-Site operations for one of the largest M365 services with over 200 million monthly active users and over 1 exabyte of data, including incident response and management, rapidly diagnosing and resolving critical issues to maintain SharePoint Online's 99,99% SLA.

- **Disaster Recovery & Infrastructure Modernisation:** Led an initiative to improve disaster recovery playbooks using a more resilient storage solution, ensuring recovery procedures remained accessible even during blackout scenarios.
- **Onboarding & Global Expansion:** Played a key role in onboarding and training a new SRE team in China, enabling 24/7 follow-the-sun operations.
- **Community & Knowledge Sharing:** Organised meet-ups for Microsoft Ireland's Open Source Club.

Software Engineer, Network Security, Sophos, Karlsruhe, Germany

01/2017 to 12/2018

- **Network Security & Threat Detection:** Worked on the Synchronised Security Engine, significantly improving network threat detection rates compared to competing vendors.
- **IPSec & Network Protocol Implementation:** Worked on the implementation of IPsec IKEv2 in the Linux Kernel for the Firewall Appliance.
- **Scalability & Load Testing:** Implemented extensive firewall load testing using the Ixia BreakingPoint platform, ensuring performance under high traffic loads. Developed custom load-testing frameworks with Python Mininet SDN, simulating concurrent user traffic.
- **Testing & Release Acceleration:** Expanded the integration test suite for firewall products, leading to faster and more reliable release cycles.

Software Engineer, MARKANT Handels and Service GmbH, Germany

08/2015 to 12/2016

Moved the department from CVS to Git with a CI/CD pipeline, taking deployments from weekly to several times a day, and built operations tooling for time-sensitive trading systems.

Junior Software Engineer, Streit Datentechnik GmbH, Haslach im Kinzigtal, Germany

09/2012 to 07/2015

C++, C# .NET and T-SQL development on Windows, including a disassembler that reads dependencies out of PE and .NET executables.

Education

Software Engineering, 2015 | Gewerblich-Technische Schulen Offenburg, Offenburg im Kinzigtal, Germany
Realschule, 2012 | Realschule Schramberg, Schramberg, Germany

Language Proficiencies

German (Native) | English (Proficient)

Extra-Curricular Activities

DevOps-Lead Kubernetes rc3.world, Chaos Computer Club

12/2021

Supported the rc3.world backend team as part of the rC3 infrastructure team, overseeing and maintaining a Kubernetes cluster and a complex multi-staged Elixir deployment hosting the remote congress experience.

- **Orchestrated a highly available Elixir Application** using Kubernetes, Kustomize, Helm, ArgoCD, ingress-nginx, cert-manager, and the kube-prometheus-stack.
- **Optimised build & deployment process** by enhancing the Dockerfile layout to use a multi-stage build process, ensuring a smaller and more efficient production image.
- **Built and maintained a multi-stage GitLab CI pipeline** to compile, verify, and deploy documentation for the Elixir application powering the event.

Head of Network, GPN18/19, Entropia e.V.

01/2018 to 06/2019

Led the network infrastructure team of over 15 volunteers with a diverse background for one of the largest CCC events in Germany after the Chaos Communication Congress, acting as the primary representative towards event management, handling sponsor acquisition, multi-homed BGP connectivity, budget management, and logistics for a fault-tolerant network.

Personal Projects

- **Home Lab:** Operate two Kubernetes clusters using ClusterAPI on Hetzner, K3s on Raspberry Pis, as well as a Raspberry Pi based CEPH storage cluster, providing a solid testbed for distributed systems research.
- **Deployment & Automation:** ArgoCD/GitOps-based Kubernetes deployments and Ansible for bare metal services.
- **Observability & Monitoring:** Maintaining a full observability platform running on Kubernetes consisting of the Grafana LGTM Stack (Mimir, Loki, and Tempo) and Thanos for metrics, logs, and traces, ensuring real-time visibility into performance.
- **Networking & Security:** Build and manage a zero-trust home network using Tailscale, leveraging ACL-based access control for network segmentation and security.
- **Open Source Contributions:** Actively contribute to multiple open-source projects.

Core Skills

- Go, Kubernetes Operators & Distributed Systems
- Multi-Cloud Infrastructure (AWS, Azure, Google Cloud, OCI)
- Chaos Engineering & Resilience Testing
- CI/CD, Deployment Orchestration, Infrastructure as Code
- Low-Level Systems Programming
- Technical Leadership, Mentorship & Stakeholder Management

Talks

- [Site Reliability Engineering Explained - An Exploration of DevOps, Platform Engineering, and SRE](#)
- [Understanding Alerting - How to come up with a good enough alerting strategy](#)
- [Modern Observability - Scalable Observability with the LGTM Stack: Harnessing the Power of Loki, Grafana, Tempo, and Mimir](#)
- [From 0 to Kubernetes Eine Einführung zur Container-Orchestrierung mit praktischen Antworten auf die häufigsten Fragen wie "warum?" oder "wann?"](#)
- [Congratulations, it's a Kubernetes! But now what?](#)
- [Hilfe, ich bin Manager - und jetzt? Oder auch: Wie man in der IT Karriere machen kann](#)