Kaan Cinar

kaancinar.com | github.com/aydinkaancinar

Email: kaan.cinar@alumni.utoronto.ca

Mobile: +1-437-982-0096

Professional Summary

Software Developer and Platform Engineering Professional with 5 years of experience building **developer tooling and automation** and operating and improving **cloud-based SaaS services and high-availability infrastructure**. Experienced in **Python development, Kubernetes, API integrations, infrastructure-as-code, performance engineering, production debugging, observability, capacity planning, and multi-region architecture**.

Experience

- **IBM** Markham, ON

  Staff Site Reliability Engineering Professional (Band 7 – Promoted October 2025) *October 2025 – August 2026*

  - **Technical Architecture & Client Ownership**: Served as the **primary SRE point of contact** for a strategic enterprise customer generating approximately **85 million API calls per week, peaking at 1.6 million per hour**, leading weekly architecture discussions, authoring technical design documents for client and internal stakeholders, and guiding production reliability decisions.
  - **Production Kubernetes Troubleshooting**: Continued hands-on support, as part of an engineering team, across **50+ Kubernetes environments**; diagnosed application, DB2, node, load-balancer, DNS, TLS, routing, and firewall/allowlist failures using logs, metrics, SQL, and network diagnostics.
  - **Multi-Region High Availability**: Architected a dual-environment design using **IBM Cloud Internet Services (Cloudflare) Global Load Balancing**, health checks, automated failover, and **DB2 HADR**; created architecture diagrams and led design reviews before rollout, enabling cross-region service and data redundancy with near-zero interruption for critical services during regional failover.
  - **Capacity Planning & Performance**: Built capacity-planning dashboards for node, pod, and deployment utilization; analyzed CPU and memory demand, workload placement, and top resource consumers to select **node flavors, worker-pool sizes, and deployment replica counts**. For a customer generating approximately **85 million API calls per week**, translated planned usage forecasts into required node counts, machine profiles, and per-deployment pod capacity while tracking error rates and latency against a **99.9% availability SLO**.
  - **On-Premises-to-Cloud Migration Support**: Supported recurring customer migrations from on-premises deployments to current application versions on IBM Cloud. Data-transfer paths included customers providing **DB2 backups for restoration in the cloud** or migrating their application projects directly; focused engineering work on validating and sizing cloud infrastructure to handle each customer's production workload.
  - **ArgoCD Operations & Change Delivery**: Used **ArgoCD** to deploy application changes, promote images between environments, restore services, and synchronize configuration updates following secret rotations and Kubernetes resource-definition changes; rolled out upgrades to customer environments.
  - **Release Diagnosis & Rollback**: Diagnosed upgrade failures caused by unmet customer-side dependencies and release incompatibilities with large customer projects, using application logs to investigate deployment failures; rolled back releases and coordinated prerequisite dependency upgrades with customers.
  - **PostgreSQL Recovery & Cloud Access**: During an **IBM Cloud disaster recovery exercise**, restored the latest PostgreSQL backup into a new instance, assigned its security group and virtual private endpoint gateway, created an API key, updated the managed secret, and triggered **ArgoCD** to roll out the new credential.
  - **Customer Requirements & Product Collaboration**: Reviewed customer requests daily, implemented infrastructure and YAML configuration changes to enable features, authentication, and load balancing, and worked with support, developers, and product management on product-level requests.
  - **Kubernetes Capacity & Cost Optimization**: Performed preventive capacity maintenance by analyzing workload demand and resizing Kubernetes node-pool flavors; designed a worker-pool migration with **half the CPU and twice the memory** to lower cost, moving workloads safely by draining old nodes and enabling new nodes one at a time.
  - **Observability Engineering**: Built shared **Instana dashboards** for node and pod utilization, top resource consumers, and growth trends, and **Grafana dashboards** for application checks and execution volume by tenant; used distributed tracing to follow requests across services and configured monitoring alerts and thresholds.
  - **Technical Leadership**: Led engineers and coordinated cross-functional teams during **critical incidents and disaster recovery events**; handled approximately **biweekly SEV-1 incidents** and participated in a weekend on-call rotation once every three weeks.
  - **Root-Cause Analysis & Prevention**: Led post-incident reviews and RCAs, aligned internal teams on corrective actions and recurrence mitigation, and presented findings and prevention plans directly to customers.
  - **Mentorship**: Acted as technical lead during North American operating hours, **mentoring junior engineers**, creating runbooks and troubleshooting documentation, and guiding architecture, incident, and production decisions.
  - **Enterprise Integrations & Authentication**: Helped enterprise customers integrate applications using **OAuth and REST APIs** and configure **SAML authentication**, troubleshooting integration behavior and advising on implementation.
  - **Java Client Integration Review**: Guided and reviewed customer **Java client** logic for REST API cookie storage, expiry handling, and refresh; validated correct behavior and balanced traffic through logs during development and test validation before production rollout.

- **AI-Assisted Incident Analysis**: Used **IBM Bob**, an internal VS Code-based AI assistant, to analyze non-sensitive application logs, identify patterns, and inform operational decisions; generated diagnostic reports that resolved multiple customer issues without escalation to product support.
- **AI-Assisted Automation Delivery**: Used IBM Bob to draft **Ansible playbooks and Python automation**; reviewed and corrected generated code, then validated changes in development and test environments before moving them into production.
- **Cross-Account Inventory Transfer**: Personally designed a two-stage automation to transfer inventories between **Ansible Tower instances** isolated in separate cloud accounts and private networks, using **cloud object storage** as the controlled handoff layer for inventory export and import.

- **IBM** Markham, ON
  Site Reliability Engineering Professional (Band 6) *June 2023 – October 2025*
  - **Enterprise Platform Operations**: Supported, as part of an engineering team, **50+ isolated single-tenant Kubernetes environments** and multi-tenant production clusters used for essential healthcare, banking, insurance, and airline operations.
  - **Internal Developer Tooling**: Personally designed and maintained **Python applications in a shared GitHub repository** used by 5+ engineers in production operations to retrieve kubeconfigs and switch Kubernetes clusters across IBM Cloud accounts, eliminating manual overhead and providing a foundation for additional diagnostic tools; documented each tool and created runbooks, automation, dashboards, and troubleshooting documentation used by the **entire SRE team**.
  - **Performance Testing Platform**: Developed and owned a shared-repository **multithreaded Python API testing platform** supporting configurable concurrency, payload size, request volume, and rate, with automated reports to verify critical operations during maintenance and production validation.
  - **Collaborative Code Review**: Contributed to shared automation repositories through merge requests, received peer feedback on changes, and reviewed contributions from other engineers.
  - **Infrastructure as Code**: Built and maintained shared-repository **Ansible automation across development, test, and production environments**, reducing selected operational tasks from **up to an hour of manual work to a few minutes**.
  - **Secrets & Network Security**: Integrated Ansible automation with managed secrets, coordinated secret rotation with **ArgoCD**, and maintained network allowlists for secured application access.
  - **Diagnostic Tooling**: Built tooling that reduced application and infrastructure log collection from **30+ minutes of manual effort to a one-click, 2–3 minute workflow**, accelerating distributed-system debugging.
  - **Production Bash Automation**: Shipped a **Bash cron job** that restarted selected Kubernetes pods for a customer on request, replacing a recurring manual production operation with a scheduled workflow.
  - **Identity Data Reconciliation**: Investigated customer access discrepancies by comparing users and permission groups stored in application databases against **LDAP**; identified missing or inconsistent records and determined whether manual correction was required in LDAP or the database.
  - **Linux Systems Administration**: Administered and troubleshot Linux systems using shell scripting, process and resource inspection, systemd, cron, package management, filesystem and disk diagnostics, and certificate investigation.
  - **Deployment & Disaster Recovery**: Executed **Blue/Green deployments** and multi-region disaster recovery exercises, enabling infrastructure and application changes without customer downtime.

- **Oracle** Toronto, ON
  DevOps Engineer Intern *May 2021 – September 2022*
  - **Oracle Cloud Infrastructure as Code**: Built configuration-driven **Terraform** for Oracle Cloud networks, subnets, security groups, Kubernetes clusters, and virtual machines, using loops and YAML resource definitions to improve modularity and reduce duplication.
  - **Kubernetes Engineering**: Developed **Python automation for production Kubernetes cluster upgrades**, enabling controlled upgrades while maintaining service availability.
  - **CI/CD & Containers**: Wrote **GitLab CI/CD pipelines in YAML** to build and tag **Docker** images and publish them to **JFrog Artifactory**, improving repeatability of software delivery workflows.

Projects

- **Personal Portfolio Platform** Private Repository
  AI-Assisted Software Engineering Project *2026 – Present*
  - **Frontend Engineering**: Using **OpenAI Codex** as an AI-assisted development environment to build a responsive personal portfolio with **React 19, TypeScript, Vite/Vinext, Tailwind CSS**, reusable UI components, light/dark themes, a web-native resume, and an interactive pathfinding game.
  - **Container Delivery & CI/CD**: Containerized the application with **Docker** and implemented **GitHub Actions** workflows that run production builds, validate Terraform, build and health-check the container, and publish versioned images to **GitHub Container Registry**.
  - **Cloud Deployment Automation**: Building an **Oracle Cloud Infrastructure** deployment foundation with **Terraform** and short-lived **GitHub OIDC** authentication, including trust configuration and automated token-exchange validation without storing long-lived cloud credentials.

Education

- **University of Toronto, St. George** — Toronto, ON
  *Honours B.Sc. in **Computer Science**, Minor in Statistics, Graduated with Distinction* — *May 2023*

Additional Information

- **Work Authorization & Location**: Legally eligible to work in Canada without employer sponsorship; able to attend a Toronto office regularly for hybrid work
- **Global Collaboration**: Comfortable with occasional early-morning or evening meetings to collaborate with international teams
- **Regulated SaaS Context**: Supported customer environments with HIPAA requirements and a product that was SOC 2 compliant; no direct audit or compliance-program ownership claimed

Skills

- **Programming & Automation**: **Python, Bash, SQL, TypeScript, JavaScript**, C++ (academic/personal projects), Java, Elixir (learning interest), AI-Assisted Development with OpenAI Codex and IBM Bob
- **Databases**: **DB2, DB2 HADR, PostgreSQL**, Production Database Troubleshooting, Backup Restoration
- **Cloud & Platform**: **Kubernetes, OpenShift, Docker, IBM Cloud**, Oracle Cloud, IBM Cloud Internet Services (Cloudflare), OpenShift Routes, Linux
- **Infrastructure as Code**: **Ansible, Terraform**, Configuration-Driven Provisioning
- **Software Delivery**: **ArgoCD, GitHub Actions, GitLab CI/CD, JFrog Artifactory, GitHub Container Registry**, Git, Jira (Internal Ticketing), Helm-Based Deployments Operated Through Ansible, YAML Pipelines, Image Promotion, Release Rollbacks, Blue/Green Deployments
- **Observability**: **Prometheus, Grafana, Instana**, PagerDuty (Incident Response Use), SLOs, SLAs, Error Budgets, Distributed Tracing, Alert Configuration
- **Architecture & Reliability**: **Distributed Systems, High Availability, Disaster Recovery, Multi-Region Architecture**, Load Balancing, Performance Engineering, Capacity Planning, Cost Optimization
- **APIs, Networking & Enterprise Integrations**: **REST APIs, OAuth, SAML**, Public and Private VLANs, Cookie and Session Handling, API Keys, Secrets Management, Security Groups, Private Endpoints, Network Allowlisting