

Sarah Chen

📞 (415) 555-0182 ✉ sarah.chen@email.com 🔗 [linkedin.com/in/sarahchen](https://www.linkedin.com/in/sarahchen) 📍 San Francisco, CA

PROFESSIONAL SUMMARY

Senior Software Engineer with 8+ years of experience building scalable distributed systems and data platforms. Expert in backend architecture, cloud infrastructure, and leading cross-functional engineering teams. Proven track record of reducing system latency by 60% and infrastructure costs by 35% while maintaining 99.99% uptime across production services.

TECHNICAL SKILLS

Languages: Python, Go, TypeScript, Java, SQL, Rust **Frameworks:** Django, FastAPI, React, Node.js, Spring Boot, gRPC
Cloud & Infra: AWS (ECS, Lambda, S3, RDS, DynamoDB, SQS), Kubernetes, Docker, Terraform, Helm, Jenkins, GitHub Actions **Data:** PostgreSQL, MongoDB, Redis, Apache Kafka, Elasticsearch, Snowflake, Apache Spark **Observability:** Prometheus, Grafana, OpenTelemetry, Datadog, PagerDuty

WORK EXPERIENCE

Vertex Systems

San Francisco, CA

Staff Software Engineer

Mar 2022 – Present

- Architected a real-time event processing pipeline handling 2M+ events per second across 12 microservices, reducing end-to-end latency from 800ms to 320ms.
- Led a team of 7 engineers to migrate a monolithic billing system to an event-driven microservice architecture on Kubernetes, cutting monthly infrastructure costs by \$120K.
- Designed and implemented a multi-region disaster recovery strategy achieving RPO of 30 seconds and RTO under 5 minutes for all critical path services.
- Established observability standards across 40+ services using OpenTelemetry and Prometheus, reducing mean time to detection (MTTD) for incidents by 45%.
- Drove adoption of Rust for performance-critical data ingestion components, resulting in 4x throughput improvement and 70% lower memory footprint.

Nimbus Labs

Remote

Senior Backend Engineer

Jun 2019 – Feb 2022

- Built and maintained the core recommendation engine serving 15M daily active users, achieving a 22% increase in user engagement through collaborative filtering improvements.
- Redesigned the data ingestion layer from batch processing to streaming (Kafka + Flink), reducing data freshness SLA from 6 hours to under 2 minutes.
- Implemented a feature flagging service with gradual rollout capabilities, used by 12 product teams and enabling safe deploys with automated rollback triggers.
- Mentored 4 junior engineers through structured code review processes and weekly architecture deep-dives, with 3 promoted to mid-level within 18 months.
- Optimized PostgreSQL query performance across the analytics suite, reducing average query time by 58% through strategic indexing and materialized views.

Apex Financial

New York, NY

- Developed a high-throughput trade reconciliation engine processing \$2B+ in daily transactions with zero-tolerance accuracy requirements, passing SOC 2 Type II audits.
- Built RESTful APIs and gRPC services handling 50K requests per second for the core trading platform, maintaining p99 latency under 100ms.
- Automated the deployment pipeline using Jenkins and Docker, reducing release cycle time from 2 weeks to daily deployments with canary rollout strategies.
- Created a comprehensive integration test suite achieving 92% code coverage, catching 14 critical production regression bugs before release in the first quarter of adoption.

EDUCATION

Master of Science, Computer Science

Stanford University

Sep 2014 – Jun 2016 GPA: 3.92/4.0 | Focus: Distributed Systems, Machine Learning

Bachelor of Science, Computer Engineering

University of California, Berkeley

Sep 2010 – May 2014 GPA: 3.85/4.0 | Tau Beta Pi Engineering Honor Society

PROJECTS

Cascade | *Open-source distributed task scheduler*

Starred 2.4K+ on GitHub

- Designed a horizontally scalable task orchestration engine in Go with pluggable backends (Redis, Kafka, PostgreSQL), supporting DAG-based workflows with retries and dead-letter queues.
- Built a gRPC API and CLI tool, with comprehensive documentation site and contributor guide. Adopted by 3 organizations for production batch processing workloads.

SignalWatch | *Real-time system anomaly detection*

- Created an unsupervised anomaly detection system using statistical models and streaming window functions, achieving 94% precision with sub-second detection on live production telemetry.
- Integrated with PagerDuty and Slack for intelligent alert routing based on anomaly severity and team on-call schedules.

CERTIFICATIONS

-
- AWS Solutions Architect Professional – Vertex Web Services, 2021
 - Certified Kubernetes Application Developer (CKAD) – CNCF, 2020
 - Nexa Cloud Professional Data Engineer – Nexa Cloud, 2019