

Pinglei Guo

Vancouver, BC, Canada plguo002@gmail.com [linkedin.com/in/at1510086](https://www.linkedin.com/in/at1510086) github.com/at15 at15.dev

WORK EXPERIENCE

- AWS:** Software Engineer - Amazon CloudWatch and X-Ray Vancouver, Canada Sep 2020 - Present
- Led a team of 6 engineers to implement W3C trace ID support in AWS X-Ray. Refactored data sharding and indexing logic to remove X-Ray trace ID timestamp dependency, unblocking 20,000 AWS accounts and enabling OpenTelemetry integration.
 - Automated on-call investigation using AWS Bedrock by indexing runbooks and incidents into OpenSearch for RAG and creating an agent to execute common runbook commands. Reduced average investigation time from 20 to 5 minutes.
 - Collaborated with Lambda and Application Signals teams to integrate OpenTelemetry in X-Ray ingestion API. Reduced Python Lambda function cold start time by 60% by removing OpenTelemetry collector from the Lambda layer.
 - Designed and implemented conversion logic between OpenTelemetry trace spans, X-Ray segments and CloudWatch transaction search log events, reducing customer trace data storage cost by 70% and enabling 100% trace sampling.
- AWS:** Software Engineer - CloudWatch Agent and OpenTelemetry distro Vancouver, Canada Sep 2020 - July 2021
- Maintainer of open source projects on GitHub: Amazon CloudWatch Agent and Amazon Distro of OpenTelemetry Collector.
 - Developed the ECS extension for OpenTelemetry Collector, enabling Prometheus discovery with metadata from ControlPlane.
 - Created Kubernetes operator to shard Prometheus scrape targets dynamically based on load instead of static hashmod in hackathon. Contributed to OpenTelemetry Collector Operator as Target Allocator and used in EKS add-on.
- Google:** Software Engineer - Batch job on Kubernetes Sunnyvale, CA May 2019 - Nov. 2019
- Eliminated 5% job loss rate during controller deployment by optimizing the sequence of database updates and job submissions.
 - Reduced debug session startup latency from 2 minutes to 40 seconds using on-demand artifact loading from object store.
- PayPal:** Software Engineer - Multi cluster container orchestration platform in Go San Jose, CA May 2018 - May 2019
- Developed a Go-based monitoring tool to correlate logs, metrics, and traces during deployment failures, automating root cause identification and reducing on-call debug time from 10 minutes to 2 minutes.
 - Introduced a canary deployment strategy and dependency readiness checks, reducing deployment rollback ratio by 20%.
- Dongyue Web Studio at SJTU:** (Part-time) Full stack web developer & Tech lead Shanghai, China Sep. 2013 – Jan. 2016
- Led web and mobile teams to redesign the online event booking website tongqu.me, implementing a REST API and single-page application using AngularJS. Reduced codebase size by 60% and improved page load time by 80%.
 - Implemented a Redis and MySQL based queue system to handle ticket rushes, preventing website crashes and ticket over-selling, while supporting 5x more traffic.

PROJECT EXPERIENCE

- Distributed database benchmark tracker** github.com/benchhub UCSC 2017 – 2018
- Designed configuration for running database benchmark for different RDBMS and TSDB in distributed environment.
 - Built a CI service to store benchmark results in databases for performance regression detection
- Distributed Time Series Database** github.com/xephonhq/xephon-k UCSC 2016 – 2018
- Implemented a distributed time series database on top of Cassandra in Go. Support both JSON and Protobuf over HTTP/2.
 - Designed a columnar storage engine modeled after Parquet and InfluxDB with high compression and less write amplification.
 - Created a benchmark suite for OpenTSDB, KariosDB, InfluxDB with a generic interface to support different TSDB.
 - Published a comprehensive report on TSDB design and implementation [awesome-time-series-database](https://github.com/xephonhq/awesome-time-series-database).

EDUCATION

- | | | |
|-----------------------|-------------------------------------|-----------------------|
| MS. Computer Science | University of California Santa Cruz | Sep. 2016 – Mar. 2018 |
| BS. Materials Science | Shanghai Jiao Tong University | Sep. 2012 – June 2016 |

SKILLS

- | | |
|-----------------------|-------------------------------------------------------------------------------|
| Programming Languages | Go, Java, TypeScript, JavaScript, Python, Dart, Rust, SQL |
| Cloud and Infra | AWS, GCP, Cloudflare, Kubernetes, Docker, CDK, Terraform, Prometheus, Grafana |
| Databases | DynamoDB, Cassandra, MySQL, PostgreSQL, Elasticsearch, OpenSearch, Redis |
| Frameworks | gRPC, React, Next.js, Spring |