

Adam Petrovic

Senior Software Engineer – Storage Platforms, Cloud Infrastructure & Reliability

Sydney, Australia

 adam@petrovic.com.au  [linkedin.com/in/adampetrovic](https://www.linkedin.com/in/adampetrovic)
 github.com/adampetrovic  adampetrovic.me

Summary

Senior software engineer with 14+ years building high-scale cloud infrastructure, storage platforms, observability systems, and reliability tooling. Current work focuses on Atlassian's internal S3-compatible Object Store platform managing 60PB+ of mixed object data, including GCP/Kubernetes migration, data-integrity verification, storage analytics, compliance, and cost/reliability improvements. Previously led company-wide observability and metrics platform work, including a \$2M+ annual vendor migration and 180M metrics/minute pipelines. Focused on turning ambiguous infrastructure risk into architecture, RFCs, production systems, and tooling adopted across teams.

Skills

- **Languages:** Kotlin, Go, Python, Java, TypeScript, SQL, Bash
- **Cloud & Infrastructure:** AWS, GCP, Kubernetes, Docker, Spinnaker, Helm, CI/CD
- **Storage & Data Systems:** S3, GCS, DynamoDB, PostgreSQL, MySQL, Redis, sharding, backup/restore
- **Reliability & Observability:** SignalFx, Datadog, Prometheus, Grafana, Splunk, StatsD, Sentry, synthetic monitoring, SLOs
- **Data & Performance:** Kafka, Apache Flink, SQS, Databricks, Athena, Locust, high-volume telemetry pipelines

Experience

Atlassian (11+ Years)

Senior Engineer, Transactional Data Platform (Object Store)

Sydney, Australia

August 2023 – Present

- Designed Kotlin-based architecture for Atlassian's S3-compatible Object Store platform, managing 60PB+ of mixed object data and layering encryption, compression, data residency, compliance, and backup/restore over cloud storage for product and platform teams.
- Owned architecture, migration strategy, and implementation for moving Object Store from a traditional microservice to a Kubernetes-native workload on GCP, establishing a migration pattern reused by partner teams while balancing speed with data-integrity safeguards across runtime, credentials, verification, and rollout to deliver a company-wide OKR four weeks early.
- Drove Object Store's technical direction through RFCs and decision records for GCS checksums and seekable compression, turning ambiguous migration, cost, and performance trade-offs into roadmap decisions.
- Built daily S3 Inventory and DynamoDB export pipelines into Databricks, enabling self-service queries over 60PB+ of mixed object data for ownership, cost, integrity, incidents, compliance, and risk; turned multi-day data-integrity investigations into production verification jobs with alerting and delivered \$25K/month in object-tag savings.
- Created a Bitbucket deployment-config diff pipeline that flags risky changes before merge, adopted by multiple platform teams to catch blast-radius issues before production.
- Built repeatable Locust load tests for staging, letting engineers validate feature changes under production-like traffic, compare results across runs, and build confidence before deployment.

Engineering Manager, Observability

October 2020 – August 2023

- Led company-wide metrics migration from Datadog to SignalFx for Atlassian engineering, saving more than \$2M annually while keeping critical metrics workflows stable for thousands of engineers during transition.
- Managed a 10-person team across roadmap, stakeholders, incidents, hiring, coaching, performance development, and operational trade-offs.
- Led architectural shift from StatsD-based metric delivery to OpenTelemetry on Kafka/Flink, proving the company-wide Kafka use case that later supported a dedicated Kafka platform team.

Senior Engineer, Observability

June 2017 – October 2020

- Built and operated Atlassian’s metrics infrastructure at 180M datapoints/minute, supporting production monitoring, alerting, and incident diagnosis across cloud services and internal platforms.
- Developed Pollinator, a Go-based synthetic monitoring and post-deployment verification platform used across regions to catch reliability and performance regressions in release flows before customers did.
- Extended GoStatsD with Prometheus-style histograms and moved telemetry toward Kafka/Flink, improving percentile metrics, stream-processing headroom, and SLO diagnosis for service teams.

Senior Site Reliability Engineer, Jira Cloud

March 2015 – June 2017

- Led a cross-continental SRE group delivering Jira Cloud’s European infrastructure ahead of a public launch milestone, supporting enterprise customers, regional growth, and product commitments.
- Reworked Jira Cloud site imports and introduced Wheel of Misfortune training, reducing failed customer imports, on-call toil, and incident-response gaps through rehearsed failure scenarios.

Freelancer.com (3.5 Years)

Sydney, Australia

Software Engineer

August 2011 – February 2015

- Built platform services for authentication, messaging, notifications, email, analytics, and fraud systems in a global marketplace using Python, Go, Redis, RabbitMQ, Thrift, and async workers.
- Helped move critical paths away from PHP monolith patterns into service-oriented APIs and async processing, improving request latency, failure isolation, and delivery speed for core marketplace flows.

Education

The University of Sydney

Sydney, Australia

Bachelor of Information Technology, Computer Science

2008 – 2012

- Dean’s List with a distinction average.