Marharyta Halamuzdava

Performance & Software Quality Engineer — large-scale distributed systems

mhalamuzdava@gmail.com marharyta-halamuzdava

Summary

Performance and software-quality engineer with 9+ years testing and optimizing large-scale distributed systems. Deep across system architecture, network, and databases, with a broad monitoring/profiling toolset — and a habit of automating the analysis itself to cut investigation time and sharpen insight.

Experience

Palo Alto Networks, Staff Performance Engineer — Dec 2022 – Apr 2025, 2 years 5 months

- Optimized database and search performance (OpenSearch, Cassandra, BigQuery) through data-node configuration and query optimization.
- Implemented Kubernetes Horizontal Pod Autoscaling (HPA) in production, yielding significant infrastructure cost reduction.
- Guided major platform migrations — application move to GCS and BigQuery, AWS Keyspaces → Cassandra, AWS MSK → hosted Kafka, and Spring Boot / Java / MSK upgrades — analyzing regression risk and preventing critical performance issues through early detection and proposed fixes.
- Designed and implemented end-to-end performance regression and integrated it into the CI pipeline (including Python data-generation and cleanup for high-volume environments), cutting level-of-effort and making the suite runnable by other teams.
- Performed JVM memory profiling and optimization — resolving leaks, cache mismanagement, and high-utilization issues.
- Established integrated monitoring across Splunk, OpenSearch, Kibana, and Grafana; built advanced performance dashboards with Prometheus and InfluxDB and a comprehensive metrics-tracking and alerting system.

VISA Inc., Senior Performance Engineer — July 2019 – Apr 2021, 1 year 10 months

- Developed capacity models through systematic analysis of application performance patterns and resource utilization, and architected monitoring strategies from detailed study of customer applications and infrastructure diagrams.
- Analyzed execution plans to find MSSQL query bottlenecks and optimized them through rewriting and indexing of complex queries, reducing overall latency by 40%.
- Fixed uneven traffic distribution by tuning load-balancer configuration and health-check parameters — reducing latency, improving resource utilization across service pods, and increasing system stability.
- Eliminated OutOfMemory errors by tuning generation ratios, expanding metaspace limits, and sizing the heap to usage patterns; optimized garbage collection by benchmarking multiple GC algorithms and applying the best configuration.
- Built automated performance data-collection and analysis frameworks in Java and Bash, significantly reducing manual effort.
- Applied CPU, memory, and bus-architecture knowledge to performance characterization and software optimization, improving system performance for hardware components (CPU / memory / disk).

EPAM Systems, Senior Performance Analyst — Feb 2016 – Sept 2018, 2 years 8 months
Client: Wolters Kluwer

- Transformed the performance-testing process — realistic system-load scenarios plus an automated testing framework — cutting release cycles from ~10 days to ~3 days with

higher quality and improved reliability, and a stronger ability to predict and prevent performance issues.
- Mentored and onboarded a junior performance engineer across a stream of 10+ complex independent products, structuring the knowledge-transfer and mentoring strategy and growing them to an independent contributor within 3 months.
- Continuously expanded testing scope to new applications and features, defined non-functional performance requirements, and authored technical test plans surfacing benefits and risks.
- Led performance-results overviews — technical summaries for the engineering team and historical trend/improvement summaries for the business team.
- Automated test-data generation in Java and built batch scripts for integration testing of Java applications; developed and maintained JMeter scripts and ran distributed load generators for high-load tests at 1K+ TPS against AWS and on-premise environments.

EPAM Systems, Performance Analyst — Consultant — Nov 2015 – Dec 2015
Client: Thomson Reuters — 2 months

- Consulted on improving the existing performance-testing strategy using domain best practices — reviewing PoCs, Confluence structure, result reports, tooling, and methodology — and delivered a detailed process-overview presentation with a list of recommendations.
- Tracked adoption of the recommendations, shared constructive feedback, and supported effective cross-team communication.

EPAM Systems, Performance Analyst — Dec 2014 – Aug 2015
Client: a customer in the travel domain — 9 months

- Defined performance requirements from Google Analytics production usage, designed and built JMeter workload models/simulations to test XML-based web services, established performance and scalability benchmarks, analyzed results to identify bottlenecks, and led results overviews for dev/DB engineers and managers.
- Found the root cause of the slowest transactions through MS SQL database profiling and recommended procedure and index optimizations.
- Ran demo sessions and shared performance-testing knowledge with the development team, cutting feedback time.

EPAM Systems, Jr. Performance Analyst — Sept 2014 – Oct 2014 — 2 months

- Developed a parameterized JMeter script to load a SOAP web service and identified system metrics and service capacity.
- After spotting high CPU usage, performed .NET code profiling to locate hot spots; the resulting algorithm optimization increased service capacity 10×. Also compared profiling tools and shared trade-offs with the team (adopting ANTS Profiler).
- Identified functional and web-service configuration issues during performance testing and collaborated closely with developers on detailed, reproducible reports — roughly halving the issue-fixing time.

Skills

Performance Testing: Apache JMeter, Locust, Pytest, LoadRunner, Custom Python load tooling, Selenium WebDriver

Profiling Diagnostics: JVM memory profiling & GC tuning, Java Flight Recorder (JFR), Java Mission Control (JMC), YourKit, AppDynamics, Heap dump analysis, MSSQL Profiler / SQL Profiler, ANTS / dotTrace (.NET profilers)

Observability Monitoring: Prometheus, Grafana, Splunk, ELK (Elasticsearch/Logstash/Kibana), InfluxDB, Kubernetes Dashboard, SCOM / SquaredUP, Perfmon, Sysdig, Fiddler

Cloud: AWS (EC2, EKS, S3, RDS, Athena, Glue, Keyspaces, MSK, CloudWatch, OpenSearch), GCP (GKE, BigQuery, Pub/Sub, Cloud Monitoring, GCS)

Databases Data: Apache Cassandra, OpenSearch / Elasticsearch, Google BigQuery, MS SQL Server, Apache Kafka / MSK

CI CD Infra: Docker, Kubernetes (incl. HPA), Helm, Jenkins, Terraform, Gradle, Git, Linux (systems administration)

Languages Scripting: Python, Java, Bash, SQL, Regular Expressions, XPath / CSS selectors

Methodology: Performance regression (E2E, CI-integrated), Capacity planning & modeling, Bottleneck analysis (CPU/memory/GC/I/O/DB/network), Non-functional requirements & test planning, High-load / distributed load generation (1K+ TPS)

Education

Belarusian National Technical University, Bachelor's Degree in Industrial Engineering