

Best Technology Talent For Companies Around The World



Stanislav

Senior Java Developer

Summary

Senior Java Software Engineer with 9+ years of experience building backend, microservices, and event-driven distributed systems. Strong hands-on expertise in Java, Spring Boot, Kafka, PostgreSQL, Redis, Elasticsearch, Cassandra, AWS, GCP, Kubernetes, and Docker, with particular strength in performance optimization, database/query tuning, and high-load systems.

Experienced in owning the full development lifecycle, from system design and data modeling to implementation, testing, production delivery, and performance tuning. Has a strong track record of measurable improvements, including migrating 30+ services from Java 8 to Java 21, doubling concurrent request capacity, reducing memory footprint by 28%, cutting production incidents by 40%, and reducing infrastructure costs by 30%. Hands-on experience with Kafka-based architectures, gRPC/Protobuf communication, resilience patterns, observability with Micrometer/Prometheus/Grafana, and performance testing with Gatling.

Previously worked extensively with high-volume data and streaming systems, including pipelines processing 5M+ events per day, GCP Dataflow and Apache Beam, Kafka optimization, Google Spanner, Cassandra, and Elasticsearch. Has also led a team of 3 backend engineers and mentors developers to mid and senior levels.

Additionally, actively uses AI-assisted development tools, including Claude Code, MCP servers, custom skills, and LLM agents. Built internal MCP tooling that automated engineering activities such as test scaffolding, release notes, and code-review checklists, saving the team 20+ hours per sprint.

TECHNICAL SKILLS:

- Core: Java, Spring Boot, Spring Data, Spring Cloud, Spring Security, Hibernate / JPA, Microservices, Event-Driven Architecture

- AI-Assisted Development: Claude Code, Model Context Protocol (MCP) servers, custom skills, agent workflows and harnesses; integrating LLM coding agents into the SDLC
- Messaging & Streaming: Apache Kafka, Confluent Schema Registry, Apache Avro, RabbitMQ, gRPC, Protobuf
- Data Stores: PostgreSQL, Redis, MySQL, Cassandra, Elasticsearch, Google Spanner
- Testing & Performance: JUnit, Mockito, Gatling, Micrometer, JProfiler, VisualVM, JVM/GC tuning
- Cloud & DevOps: AWS, Google Cloud Platform (GCP), Kubernetes, Docker, Jenkins, CI/CD
- Big Data: Apache Beam, GCP Dataflow
- Tooling: System Design, Design Patterns, UML, Git, Maven, Gradle, Liquibase, Flyway, JIRA

Professional Career

Senior Software Engineer

April 2021 – Present

Project info:

Backend distributed system built on Java and Spring Boot with 15+ microservices, Kafka-based communication, cloud infrastructure, observability, performance testing, and AI-assisted product and development capabilities

Responsibilities:

- Designed and implemented an AI image-to-text and speech-to-text recognition feature that cut data-entry time by 4x and attracted 20% month-over-month new users
- Led migration of 30+ services from Java 8 to Java 21, reducing memory footprint by 28% and enabling virtual threads that doubled concurrent request capacity
- Reduced security vulnerabilities by 80% by integrating SAST/DAST scanning with SonarQube and OWASP Dependency-Check and remediating critical CVEs
- Reduced mean time to resolution from hours to under 30 minutes by instrumenting services with Micrometer, Prometheus, Grafana, and distributed tracing
- Reduced production incidents by 40% quarter-over-quarter by adding Resilience4j circuit breakers, timeouts, and retry budgets
- Cut CI feedback time from 40 to 12 minutes by parallelizing Jenkins stages, caching builds, and moving integration tests to Testcontainers
- Reduced backend-to-backend request time by 20% by migrating inter-service communication to gRPC and Protobuf
- Reduced infrastructure costs by 30% while maintaining production stability through Gatling performance testing and capacity right-sizing

- Built internal MCP servers and Claude Code skills that automated test scaffolding, release notes, and code-review checklists, saving 20+ hours per sprint

Software and technologies:

Java, Spring Boot, Spring Cloud, PostgreSQL, gRPC, Protobuf, AWS, Kafka, Kubernetes, Docker, Keycloak, Elasticsearch, Redis, Jenkins, Git, Maven, Micrometer, Grafana, Prometheus, Gatling, Claude Code, MCP, LLM agents

Senior/Lead Software Engineer

March 2019 – April 2021

Project 1 info:

High-volume event-driven streaming platform processing 5M+ events per day using Java, Kafka, Apache Beam, GCP Dataflow, and Google Spanner

Responsibilities:

- Redesigned the database schema and implemented asynchronous event-driven processing, reducing server response time from 15 seconds to 400 milliseconds
- Built Apache Beam jobs on GCP Dataflow with autoscaling to process 5M+ events per day while keeping processing cost flat as volume doubled
- Reduced Dataflow costs by 25% through pipeline profiling, shuffle elimination, and worker right-sizing
- Designed the Google Spanner data model with interleaved tables and secondary indexes, keeping read latency under 10 ms at peak QPS
- Reduced Kafka storage by 30% and increased network throughput by 40% by adopting Confluent Schema Registry with Apache Avro
- Reduced memory consumption by 80% by re-architecting bulk processing into streaming and chunked batch jobs
- Improved message-processing throughput by 40% through targeted proofs of concept on the streaming pipeline
- Led a team of 3 backend developers and coordinated delivery against architect and stakeholder commitments

Software and technologies:

Java, Spring Boot, Spring Cloud, GCP, Google Spanner, Dataflow, GKE, Apache Beam, Kafka, Kubernetes, Docker, Jenkins, Git, Gradle, Gatling

Project 2 info:

High-load Java backend and event-processing environment using Elasticsearch, Cassandra, Redis, Kafka, AWS, and JVM performance tooling, including peak Black Friday traffic

Responsibilities:

- Reduced Elasticsearch query execution time by 40% and cluster load by 15% through query and indexing optimization
- Cut GC pause times by 50% and CPU utilization by 20% by analyzing and tuning G1 garbage collector behavior under load
- Improved system response time by 20% by migrating 10 GB of cached JSON data from disk to Redis
- Stabilized Cassandra under peak load by fixing hot-partition keys and tuning compaction, reducing read timeouts by 90%
- Kept end-to-end event latency under 200 ms at peak by tuning Kafka consumer batch sizes, rebalancing partitions, and adding backpressure handling
- Enabled the system to handle 5x Black Friday traffic with no incidents through Gatling load testing, Redis caching, and horizontal pod autoscaling

Software and technologies:

Java, Spring, Guice, Guava, Freemarker, Jetty, Tomcat, MySQL, Cassandra, Elasticsearch, AWS, Zookeeper, JProfiler, VisualVM, Jenkins, Git, SVN, Maven, Bash

Java Software Engineer

August 2018 – March 2019

Project info:

Low-latency Java system connecting to 30+ concurrent exchanges and processing real-time communication through non-blocking pipelines

Responsibilities:

- Built fully non-blocking I/O pipelines with Netty and RxJava, sustaining sub-50-ms decision latency across 30+ concurrent exchange connections
- Reduced SQL query execution time by 30% by introducing Hibernate second-level caching

Software and technologies:

Java, Spring, Hibernate, Netty, RxJava, WebSockets, JUnit, Mockito, Git, Maven

Java Software Engineer

April 2017 – August 2018

Project info:

Full-stack Java web application with RESTful APIs, asynchronous communication, and a real-time chat system

Responsibilities:

- Built and integrated RESTful APIs with frontend applications

- Increased system responsiveness by switching from synchronous to asynchronous communication
- Designed and implemented a real-time chat system covering frontend and backend

Software and technologies:

Java, JavaScript, Spring, Hibernate, MySQL, Redis, RabbitMQ, WebSockets, Thymeleaf, JUnit, Maven, Git

Education

Dnipro University of Technology

Master's Degree in Computer Science

Languages

English: Upper-Intermediate

Ukrainian: Native