Miscs

- Real-Time Support Ticket Intelligence Pipeline
- Delivery AI Software Engineering Senior Manager

Real-Time Support Ticket Intelligence Pipeline

Scenario

You're joining "TicketPulse," a SaaS company handling ~50,000 support tickets/day across email, chat, and social. Support agents currently triage everything manually. Your job: build the backbone of a system that ingests tickets, uses an LLM to classify and enrich them, and makes them searchable and filterable for agents and analysts in near real time.

Required stack

- **FastAPI** : ingestion endpoint + query API
- **Kafka** : event backbone between ingestion and processing
- **Celery** : async enrichment workers
- **An LLM** : (any provider, or a mocked/local stand-in that mimics real latency/failure behavior) classification and reply drafting
- **Elasticsearch** : indexed store for search, filtering, and basic analytics

Core requirements

1. **Ingestion API** : `POST /tickets` accepts a raw ticket (customer id, channel, subject, body, timestamp) and publishes it onto a Kafka topic. Should return immediately, not block on processing.
2. **Enrichment pipeline** : consumes from Kafka, and for each ticket:
   - Deduplicates (e.g., same customer + similar content within a short window)
   - Calls the LLM to classify: category, urgency (low/med/high), sentiment
   - Generates a short suggested-reply draft
   - Handles LLM failures without losing the ticket (retry, then fall back to a "needs manual triage" state, never silently drop)
3. **Indexing** : enriched ticket lands in Elasticsearch with a mapping that supports both full-text search on the body and structured filtering (category, urgency, sentiment, date

range).
4. **Query API** : `GET /tickets/search` supporting full-text search + filters; bonus if it supports "similar past tickets" for a given ticket id.
5. **Idempotency** : replaying the same Kafka message twice must not double-index or double-notify.

Explicitly out of scope

- No frontend/UI, API responses are enough
- No auth/multi-tenancy
- No actual load testing at 50k/day, reason about it, don't prove it
- No fine-tuned model, prompting an off-the-shelf LLM (or a stub) is fine

Sample data

Use a public dataset like Kaggle's "Customer Support on Twitter" for realistic ticket text, or generate synthetic tickets with an LLM. Either is fine, just note which one and why in your README.

Deliverables

1. **Working repo** : `docker-compose up` should bring up Kafka, Elasticsearch, Celery worker(s), broker, and the FastAPI app.
2. **README** : how to run it, and the key design decisions you made (not a commit-by-commit history).
3. **ARCHITECTURE.md** : short written answers to:
   - How does this scale to 10x ticket volume? What's the first thing that breaks?
   - How do you control LLM cost at this volume?
   - What's your dead-letter / poison-message strategy when a ticket keeps failing enrichment?
   - What would you monitor in production, and what would page someone at 2am vs. just show up on a dashboard?
   - What did you deliberately cut given the time box, and what would you build next?

Optional stretch

- Dead-letter queue with a visible retry/inspection path

- Cost-aware LLM routing (cheap model for obvious cases, escalate only when uncertain)
- Semantic "similar tickets" via embeddings instead of ES `more_like_this`
- Basic metrics/tracing

Submission

Send a link to the repo (or a zip) with the README and ARCHITECTURE.md included. Reach out if any requirement is ambiguous, how you handle ambiguity is also part of what we're looking at.

Delivery AI Software Engineering Senior Manager

About the job

Who We Are

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we help clients with total transformation-inspiring complex change, enabling organizations to grow, building competitive advantage, and driving bottom-line impact.

To succeed, organizations must blend digital and human capabilities. Our diverse, global teams bring deep industry and functional expertise and a range of perspectives to spark change. BCG delivers solutions through leading-edge management consulting along with technology and design, corporate and digital ventures—and business purpose. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, generating results that allow our clients to thrive.

We Are BCG X

We’re a diverse team of more than 3,000 tech experts united by a drive to make a difference. Working across industries and disciplines, we combine our experience and expertise to tackle the biggest challenges faced by society today. We go beyond what was once thought possible, creating new and innovative solutions to the world’s most complex problems. Leveraging BCG’s global network and partnerships with leading organizations, BCG X provides a stable ecosystem for talent to build game-changing businesses, products, and services from the ground up, all while growing their career. Together, we strive to create solutions that will positively impact the lives of millions.

What You'll Do

As part of BCG X, you will collaborate closely with consulting teams across a wide range of advanced topics. In this role, you will apply software development methodologies to deliver value to BCG Consulting teams, BCG X case teams, BCG X Product teams, and Practice Areas, acting as a subject-matter expert in software engineering and execution.

You will combine deep software engineering expertise with modern AI-native development

practices, leveraging platforms such as Harness, cloud technologies, AI coding and engineering agents, automation and modern DevSecOps approaches to accelerate the journey from idea to production.

You will lead and enable teams across the full product lifecycle—from discovery, architecture and rapid prototyping through development, testing, deployment and ongoing optimisation—delivering solutions that are secure, scalable, resilient and production-ready. Assignments will range from rapid AI-enabled Proofs of Concept and MVPs to enterprise-grade products, with a focus on increasing engineering velocity, quality and developer productivity through intelligent automation and emerging AI capabilities.

Our Expert Consulting Track

The Expert Consulting Track (ECT) is vital to BCG’s ability to successfully meet our clients’ demand for deep expertise and advanced technical capabilities. Experts deliver next-level impact for BCG's most complex and cutting-edge client challenges. Experts focus on developing BCG's thought leadership, commercial capabilities, and intellectual property assets; and their work is essential to our delivery model. Working as part of a multidisciplinary team to bring the ‘best of BCG’ to our clients, Experts come from diverse backgrounds, with expertise often gained outside of BCG in industry, academia, or specialized consulting.

Notre filière d'experts-conseils

La filière d’experts-conseils (FEC) est essentielle pour permettre à BCG de répondre avec succès à la demande de ses clients en matière d’expertise approfondie et de capacités techniques avancées. Les experts ont une incidence de premier ordre sur les défis les plus complexes et les plus pointus des clients de BCG. Les experts se concentrent sur le développement du leadership intellectuel, des capacités commerciales et des actifs de propriété intellectuelle de BCG, et leur travail est essentiel à notre modèle de prestation de services. Travaillant au sein d’une équipe multidisciplinaire pour offrir le « meilleur de BCG » à nos clients, ces experts viennent d’horizons divers et possèdent souvent une expertise acquise à l’extérieur de BCG dans un secteur d'activité, le milieu universitaire ou les services-conseils spécialisés.

What You'll Bring

Technical Skills

Must have strong experience:

- AI-native software engineering: Hands-on experience using Generative AI, AI coding assistants and/or agentic development tools to accelerate software design, development, testing, documentation and delivery; understanding of LLM-based applications, APIs, RAG and AI-enabled engineering workflows.
- Modern software engineering: Strong proficiency with Python and JavaScript/TypeScript, Unix/Linux development environments, Git and modern software development practices.

- Full-stack development: Experience building scalable, responsive applications using modern frontend frameworks such as React, Angular or Vue, and backend frameworks such as FastAPI, Flask, Django or Node.js.
- Cloud & platform engineering: Strong experience with AWS, Azure or Google Cloud, containerisation with Docker/Kubernetes, Infrastructure as Code such as Terraform, and modern platform engineering practices.
- Modern CI/CD & DevSecOps: Experience with enterprise CI/CD and automated software delivery platforms such as Harness, GitHub Actions, Jenkins or CircleCI, including automated testing, security, quality gates, deployment strategies and continuous delivery.
- Data platforms & APIs: Strong SQL and database experience (NoSQL, Graph, Vector…), together with designing and integrating REST APIs and scalable data services.
- Quality engineering: Experience with automated testing frameworks such as Pytest, Jest or Mocha, alongside modern approaches to test automation, continuous validation and AI-assisted software quality.
- Engineering architecture & scalability: Track record of designing, developing and supporting secure, resilient and highly scalable software applications and services, with an understanding of modern application architecture, observability and production operations.
- Developer productivity & automation: Experience using automation, reusable engineering patterns, developer platforms and AI-enabled tooling to improve engineering velocity, software quality and developer experience.

Job Requirement

- Bachelor’s or master’s degree in computer science, Engineering, Technology or a related discipline, or equivalent practical experience.
- Significant experience in software engineering and engineering management, including leading 20+ cross-functional resources across geographically distributed teams and delivering enterprise-scale products, platforms and rapid MVPs.
- Strong technical credibility with hands-on software engineering experience, combined with the ability to lead teams adopting AI-native development practices, modern engineering platforms and emerging technologies.
- Demonstrated experience driving the adoption of Generative AI, AI-assisted development, automation and agentic engineering to improve software delivery velocity, quality and developer productivity.
- Strong understanding of engineering culture, talent development, performance management and organisational leadership, ideally within a large-scale consulting, technology or software engineering environment.
- Proven ability to operate effectively in a dynamic, ambiguous environment, balancing multiple priorities while maintaining a strong focus on outcomes, quality and execution.
- Credible and inspirational technology leader with excellent communication, presentation and stakeholder-management skills, able to engage confidently with engineers, product leaders, consultants and senior executives.

- Strong understanding of modern Agile product and engineering practices, including Scrum, Kanban, DevSecOps and continuous delivery, with a track record of delivering complex programmes with excellence across quality, cost, speed and operational performance.
- Experience establishing engineering standards, governance and AI/technology guardrails while enabling teams to move rapidly and experiment with emerging technologies.
- A continuous-learning mindset, with the curiosity and leadership to evaluate and operationalise new AI, software engineering and delivery technologies that create measurable business value.

Other Skills (+)

Technical & Delivery Oversight

- Thought partnering as a trusted advisor with stakeholders / consulting teams on software development, delivery, and architecture topics, working closely with lead engineers and business stakeholders
- Taking ownership of technical envisioning, feasibility, budgeting, timelines, and objectives/plans for enterprise grade software applications
- Responsible for periodic design and code reviews, coaching on new technologies, and sharing constructive feedback with team members
- Solve complex business problems with technical solutions, identifying roadmap dependencies and collaborating with cross functional teams
- Identify and build necessary technical capabilities within the team through upskilling, training, and setting best practices and standards for development, integration and testing

People Management, Collaboration & Outreach

- Lead software and AI engineering teams across the full people lifecycle, including hiring, team design, talent development, succession planning, performance management and retention, aligned to organisational priorities and investment.
- Build a high-performing, inclusive engineering culture that encourages innovation, experimentation, continuous learning and responsible adoption of AI, while maintaining high standards of engineering excellence.
- Provide regular, meaningful and constructive feedback, creating an environment where engineers and leaders feel empowered, supported and accountable for outcomes.
- Establish yourself as a recognised technology and AI engineering leader for consulting and topic teams, acting as an ambassador for modern software engineering, AI-enabled delivery and BCG X capabilities.
- Communicate complex technical and AI concepts with clarity to both technical and non-technical stakeholders, synthesising insights into compelling narratives, presentations and

decision materials.
- Contribute to and shape BCG’s engineering IP, accelerators, reusable assets, delivery methodologies and best practices, helping scale successful approaches across teams and engagements.

Problem Solving, Technical Leadership & Mentoring

- Bring strong solution-architecture and hands-on engineering expertise across the end-to-end software and AI delivery lifecycle, from discovery and prototyping through production and optimisation.
- Lead teams in solving ambiguous and complex problems, balancing business value, technical feasibility, speed, quality, security and long-term maintainability.
- Coach engineers, technical leads and engineering managers, developing future-ready technical leadership and AI engineering capabilities across the organisation.
- Identify opportunities to apply AI, automation and intelligent engineering agents to improve delivery velocity, quality, developer experience and operational effectiveness.
- Champion a culture of continuous improvement, using delivery and engineering insights to identify opportunities to improve how teams build, test, deploy and operate software.

Additional info

Please apply for the positions equivalent to your previous work experiences. The seniority will be evaluated according to your experience and career structure at BCG X during the process.

Boston Consulting Group is an Equal Opportunity Employer. All qualified applicants will receive consideration for employment without regard to race, color, age, religion, sex, sexual orientation, gender identity / expression, national origin, disability, protected veteran status, or any other characteristic protected under national, provincial, or local law, where applicable, and those with criminal histories will be considered in a manner consistent with applicable state and local laws.

BCG is an E - Verify Employer. Click here for more information on E-Verify.