

About the Role

We're looking for an experienced **Principal Software Engineer** to lead the design, development, and operation of our **Cloud Ingestion Platform**, which powers data collection, transformation, analytics, and access across major cloud providers and cloud-native environments. You will play a key role in ensuring that high-volume, high-fidelity data from diverse sources is efficiently ingested, normalized, and made accessible for behavioral detections, investigations, and advanced analytics, while also maintaining the performance and scalability of our underlying data infrastructure.

This role blends architectural leadership with hands-on technical depth. You'll design systems that are secure, resilient, scalable, and cost effective, while mentoring teams and guiding the evolution of our cloud data ecosystem.

Responsibilities

- Architect and implement scalable ingestion and transformation platforms in our cloud environment
- Lead technical design for ETL systems that normalize diverse telemetry into clean, queryable data models
- Manage and optimize cloud database infrastructure, including workload tuning, cost optimization, scaling strategies, and automation for high availability and resilience
- Collaborate with data and infrastructure teams to implement observability, capacity planning, and proactive performance management
- Build APIs and services enabling investigative workflows, including federated query capabilities
- Ensure data integrity, security, and governance across cloud-native environments
- Drive technical excellence through mentorship, design reviews, and architectural guidance
- Evaluate and adopt emerging technologies in cloud data engineering, analytics, and infrastructure automation
- Participate in on-call rotations and deploy highly available, production-grade services

Qualifications

Required:

- 10+ years of software engineering experience, with 5+ years in distributed systems or large-scale data platforms

- Strong proficiency in Python, Rust, Java, or Go
- Expertise in ETL, data modeling, and stream/batch processing frameworks
- Deep hands-on experience with AWS, Azure, or other major cloud ecosystems
- Proven ability to manage and optimize cloud databases (e.g., Snowflake, Trino, Delta Lake, Iceberg, RDS, or similar)
- Expertise in workload optimization, scalability, and query performance tuning
- Experience deploying workloads using infrastructure as code (e.g., Terraform, CloudFormation, or Helm)
- Track record of technical leadership, including mentorship and architectural ownership
- Strong understanding of security, privacy, and compliance in cloud data environments

Preferred:

- Experience managing data lakes and federated query engines (Trino, Presto or equivalent)
- Passion for enabling customers to extract value from data we store on their behalf
- Experience building microservices and APIs focused on customer lifecycle management, entitlements, and data acquisition
- Familiarity with security telemetry (CloudTrail, O365, Azure Activity Logs, etc.)
- Experience with observability frameworks (Prometheus, Grafana, Datadog, etc.)
- Background in cybersecurity, threat detection, or security analytics
- Contributions to open-source or cloud data infrastructure projects

Education

- B.S or M.S or Ph.D. in Computer Science (or equivalent experience)

