

Arihant Sethia

• +1 669 770 5007 • arihantsethia07@gmail.com • arihantsethia.com • in arihants

Summary

Engineering leader with over 10 years of experience developing stateful, real-time, business-critical platforms in both startups and large-scale environments. Skilled in distributed systems, multi-region architectures, and autonomous platform design using declarative intent. Demonstrated success building platforms from inception and evolving critical systems. Committed to operational reliability, streamlined architecture, and scalable, secure system primitives.

Experience

CloudKitchens

Tech Lead Manager / Software Engineer

Mountain View, CA

Dec 2020 – Present

- Led the transition from store-centric to multi-entity integration modeling, enabling business expansion across third-party logistics, enterprise APIs, and related products.
- Developed a sharded orchestration service with a durable distributed cache and real-time invalidation for merchant and credential data. This removed critical-path dependencies that caused multiple Sev-1 outages, increasing order import reliability from 99.9% to 99.99%.
- Led design of a merchant account platform managing over 1 million accounts with proactive session management and usage lifecycle. Unified the account model to eliminate recurring rate-limit and authentication issues.
- Co-designed a sharded service for managing core merchant entities. Developed change feeds to support cache invalidation across the entity graph, enabling downstream services to build custom views independently.
- Designed a websocket/mqtt-based streaming platform used by over 20,000 stores, reducing order discovery latency from 30 sec to 2 sec.
- Designed and built the Integration DevKit, enabling external engineers to develop integrations. Defined the security model and artifact pipeline, and led the adoption of observability libraries across a 30+ engineer organization.
- Rearchitected the monolithic gateway into containerized bundles with automated deployments, reducing deploy time from days to <2 hrs.
- Led engineering efforts to improve store connectivity from 53% to 82%, securing 6% of ARR for Otter. Collaborated with Product and Operations to define KPIs, implemented monitoring, and automated merchant account recovery workflows.
- Mentored over 10 senior and staff engineers and set a multi-quarter platform roadmap. Defined core primitives, abstractions, and success metrics to align execution across teams. Participated in over 200 interviews and led hiring panels.

Amazon Web Services – Honeycode

Cupertino, CA

Senior Software Engineer (L6)

Jun 2020 – Dec 2020

- Led the architectural redesign of workbook assets to enable sandboxed edits and draft-publish workflows for apps and automations.
- Drove CloudTrail integration by isolating audit query processing from latency-critical workbook operations.

Software Engineer II

Jul 2018 – Jun 2020

- Optimized a memory-intensive (400GB+) service, increasing per-instance throughput from 500 to 1,800 TPS through scenario-based load testing, JVM profiling, heap analysis, and configuration optimization.
- Achieved 1,400x data reduction by identifying less than 1% relevance in the analytics pipeline. Redesigned the system with explicit metrics and asynchronous logging, improving per-pod throughput by 1.8x.
- Redesigned the load-testing setup to reduce state initialization time and bootstrap memory profile, allowing a 15x increase in test scale.
- Co-designed a successive leader election mechanism to allow fast handoff during graceful shutdown with consistent journal replay.
- Co-invented automatic application generation from user data and launched server-side deep copy to reliably manage large app hierarchies.
- Implemented multi-user undo/redo for collaborative spreadsheets/workbooks.

Amazon – Identity Services

Bangalore, India

Oct 2016 – Jun 2018

Software Engineer II

- Built the Identity Change Data Stream (ICDS), a DynamoDB CDC framework that replaced Oracle materialized views. Adopted by over 40 engineers across multiple Tier-1 services, to enable real-time search and data warehouse ingestion.
- Developed the Oracle to DynamoDB migration strategy and multiple components for data reconciliation and quality measurement. Used across Tier-1 services handling over 1 million QPS.
- Established best practices for indexing, encryption, pre-provisioning, and write sharding. Built foundational libraries for encryption at rest and symbol-table-based serialization to optimize payloads. Led DynamoDB scaling and operational readiness for peak traffic events.
- Partnered with the engineering manager to scale the team from 3 to 18. Developed interns and junior engineers through multiple levels and fostered a culture of ownership to raise technical standards.

Software Engineer

Jul 2015 – Oct 2016

- Led a zero-downtime migration of the Session Service from Oracle to DynamoDB, making it one of teams' 1st Tier-1 systems on DynamoDB.
- Extended a distributed rate-limiting library to over 600 host fleets by adding dynamic per-customer rules, auto-expiry, and real-time observability. The solution was adopted across multiple organizations and remained in production for over five years.
- Built a cross-region replication data validator for encrypted payloads on DynamoDB.

Education

Indian Institute of Technology, Guwahati

B.Tech, Computer Science & Engineering

India

2011 – 2015

Patent & Awards

Patent: U.S. Patent 11,663,199 (2023): Application Development Based on Stored Data

Awards: IIT-JEE (Top 0.3%, 2011) | AIEEE (Top 0.05%, 2011) | KVPY Scholar (2010)