

Ujjawal Kumar

[LinkedIn](#)

[Github](#)

Email : ujjawal4046@gmail.com

PROFESSIONAL SUMMARY

Experienced software engineer with 7+ years building and operating **petabyte-scale database and storage infrastructure**. Leads architecture for HBase/Phoenix data management at Salesforce, with expertise spanning backend systems, distributed storage, replication, backup and disaster recovery, cloud infrastructure, data integrity, multi-tenant security, and performance engineering.

EXPERIENCE

• Salesforce

Hyderabad, India

Lead Member of Technical Staff

Jun 2023 - Present

- **Technical Leadership:** Lead technical direction for a **10-engineer** team and define the technical roadmap for the HBase/Phoenix data management platform, driving architecture across storage, database, and cloud migration teams for initiatives including CDC-based backups, real-time backups via Flink, multi-cloud integration, and distributed restore checkpointing
- **Petabyte-Scale Migration:** Spearheaded the live, zero-downtime migration of **20+** production petabyte-scale clusters from AWS EKS to Salesforce's managed Kubernetes platform, designing data-copy, live replication transfer, rollback/failover, and production rollout strategies.
- **Tenant Lifecycle:** Redesigned the tenant data lifecycle, replacing resource-intensive MapReduce batch scans with a custom compaction-based garbage-collection coprocessor, saving **\$100K** in annual compute costs
- **Data Security:** Designed a multi-tenant data-at-rest encryption architecture at the backup file-format layer, providing strict tenant isolation and BYOK.
- **Ransomware Resilience:** Architected a ransomware-resilient backup platform using S3 Object Lock and cross-account replication to isolated recovery buckets, establishing it as the default security standard for new production deployments.
- **Reliability and Performance:** Uncovered and fixed critical correctness bugs in the replication protocol that masked silent data loss. Diagnosed glibc allocator behavior and reduced container memory usage from 36GB to 16GB through MALLOC_ARENA_MAX tuning.

Senior Member of Technical Staff

Aug 2021 - May 2023

- **Object-Storage Backup Architecture:** Redesigned backups around object storage, reducing infrastructure cost by **80%** while decoupling backups from database clusters and addressing correctness and performance differences between HDFS and object storage.
- **Mentorship:** Mentored interns and junior engineers on distributed storage systems, including hotspot mitigation, bulk-load strategies, and production storage monitoring.

Member / Associate Member of Technical Staff

Jun 2019 - Jul 2021

- **Disaster Recovery:** Built automated backup migration and parallel restore workflows, reducing Recovery Time Objective (RTO) by **60%** for critical production databases. Contributed to Kubernetes deployment, CI/CD, and production monitoring for distributed infrastructure.
- **Service Hardening:** Hardened resilience for mission-critical storage services (SOR) across 40 global data centers, resolving high-severity incidents involving MapReduce memory leaks, secondary-index corruption, and data recovery.

EDUCATION & RESEARCH

• Indian Institute of Technology Kharagpur

Kharagpur, India

Bachelor of Technology in Computer Science and Engineering

Jul 2015 – Jun 2019

- **Research Assistant:** Developed an automated Differential Fault Analysis framework for AES using frequent itemset mining; published in the Journal of Cryptographic Engineering (2019).
- **Awards:** Full Ride Scholarship, Mayukh Mitter Foundation & JBNSTS.

SKILLS

- **Software & Systems:** Java, C++, Python, Rust, Ruby, distributed databases, storage systems, replication, backup/DR, SQL/NoSQL
- **Cloud & Infrastructure:** AWS (S3, EKS, EC2, EBS, IAM), GCP, Kubernetes, Docker, Terraform, Helm, CI/CD, Apache Flink, Spark
- **Data Platforms:** HBase, Apache Phoenix, Hadoop, HDFS, MapReduce, Kafka