

Sean Levorse Boston, MA

[linkedin.com/in/smlevorse](https://www.linkedin.com/in/smlevorse) | seanlevor.se

Summary

Senior Data Infrastructure Engineer II at HubSpot with seven years of experience building and operating distributed datastores, caching systems, and internal platforms. Known for migrations that span hundreds of teams, incident response, and growing engineers through hard technical work. Diving head first into agentic development by building custom MCP tooling and agents.

Technical Skills

- **Distributed systems and datastores:** Zookeeper, Memcached, HBase, MySQL, Kafka, Hadoop, HDFS, Redis, Vitess, S3, SQS
- **Infrastructure and orchestration:** Kubernetes, Docker, AWS, Puppet, Ansible
- **Services and languages:** Java, Go, Python, Rust, REST, gRPC
- **Build and tooling:** Maven, Gradle, OpenRewrite
- **Web infrastructure and security:** DNS, TLS, NGINX, Cloudflare
- **AI tooling:** MCP tool design, LLM agent design and evaluation, agentic workflows

Professional and Leadership Skills

Technical design, project management, incident response, learning new systems, communicating technical concepts, team leadership, conflict resolution, on-call with empathy, mentorship, managing social capital across teams

Experience

HubSpot, Cambridge, MA | 2019 to present

Senior Data Infrastructure Engineer II, Caching and Coordination Team | Feb 2026 to present

- Lead large-scale projects for the team, from investigation and planning through delivery
- Own outcomes for team goals across Zookeeper, Memcached, and related datastores
- Serve as on-call for team systems and consult product teams on data system design

Technical Lead I, Caching and Coordination Team | Jul 2025 to Feb 2026

- Acting tech lead through a period when the team shrank from six engineers to three
- Drove the team's reliability-focused Feature Freeze, reducing on-call interruptions by 90%

Earlier roles: Senior Software Engineer I (SQL and Caching teams, 2021 to 2025) and Software Engineer (Big Data and SQL teams, 2019 to 2021). Earlier internships at Microsoft, HubSpot, Quality Vision International, and General Electric.

Education

Rochester Institute of Technology, B.S. Software Engineering, *summa cum laude*, 2019

Selected Projects

Agentic Cluster Recommendations

Problem: Balancing noisy neighbor risks with resource costs for isolated clusters.

Solution: Built an agent that checks each team's cluster access and existing usage, then recommends an existing cluster or a new one.

Skills developed: MCP tool design, LLM agent design and evaluation

Memcached Auth and Encryption

Problem: TLS throughput was too low for production Memcached usage and old auth system doesn't scale across regions.

Solution: Client-side encryption with service to distribute credentials and crypto keys only to authorized clients, with an ACL model for cluster owners. The proof of concept showed under 1% throughput degradation compared to 40% with TLS.

Skills developed: security design across teams, validating designs with performance data

JDBI Upgrade

Problem: Over 800 backend repos were using JDBI 2 long after EOL.

Solution: Created OpenRewrite automation for individual modules, catalogued edge cases automation couldn't handle, wrote a migration guide, and opened hundreds of automated PRs.

Skills developed: large-scale automated refactoring, build tooling, balancing automation with social capital to get teams to act, writing guides that other engineers can follow

MySQL Online Schema Change Service

Problem: Schema changes on MySQL tables are risky and hard to manage.

Solution: Maintained and extended an internal tool with schema guardrails that block risky migrations, execution through PTOSC or GhOST, a UI for tracking and rollback, and deploy gates for dependent migrations.

Skills developed: MySQL internals, developer tooling and UX, designing guardrails that prevent incidents before they happen

MySQL Cross-Region Replication

Problem: Cross-region replication with stock MySQL replication is expensive and unreliable.

Solution: Designed an S3-based system to replicate tables from North America to other regions, providing durable, cost-effective cross-region reads where standard MySQL replication was too restrictive.

Skills developed: AWS cost tradeoffs, MySQL internals, Kubernetes

HBase, HDFS, and Hadoop Ansible to Puppet Migration

Problem: Pushing cluster configs from developer laptops is insecure and doesn't scale.

Solution: Implemented pull-based Puppet configuration management gated behind mandatory code reviews.

Skills developed: configuration management at fleet scale, building CI/CD pipelines, Puppet