

Eric Wadsworth

Enterprise Software Architect / Principal Software Engineer

Washington DC	eric@wadhome.org	+1 801-703-8438	linkedin.com/in/wad4ever/
---------------	------------------	-----------------	---

Some technologies I'll admit to having used

Java	GitLab	OpenAPI	MySQL	Kubernetes	GCP
Spring Boot	Maven	Swagger	PostgreSQL	Jenkins	AWS
Spring JPA	Gradle	Liquibase	Docker	Javascript	ElasticSearch
Linux	jOOQ	Hadoop	HBase	Pig	LLMs

Some problems that I have solved

"Our large organization has a huge critical legacy system that is being modernized and re-architected, requiring hundreds of engineering teams over multiple years, and we need guidance from expert enterprise architects and software engineers."

- Leverage extensive experience in enterprise software architecture, cloud-based systems, and software engineering to construct a project roadmap for the project and write a lot of documentation.
- Author and present multiple Architecture Decision Records (ADRs) to guide the project towards best practices.
- Attain subject-matter familiarity to enable recommendation of various specific technologies and methodologies for software development teams.
- Build software libraries and reference implementations to standardize the organization on API usage, application logging, data persistence, and other core services.
- Assist in architecture around application of strangler pattern for adapting legacy software services into the new system.

"Our budding company has a new enterprise software system that is emerging from development. We need to convert it into a reliable, fault-tolerant, scalable production system."

- Shatter the monolithic database, separating it out into individual database servers for development, test, and production, migrating terabytes of data with no service interruption.
- Introduce automatic version-controlled database schema management.
- Rework back-end software services to apply a decoupled 3-tier architecture, enabling quicker feature development.
- Add automated integration testing, fostering a culture of test-driven development.
- Upgrade ElasticSearch integration.
- Apply dev-ops experience to rework build pipelines.

"We are having scalability and reliability issues with our back-end data processing software."

- Architect and develop a replacement back-end data processing system that is scalable, fault-tolerant, event-driven, and microservice-based.
- Introduce a software service that wraps a MySQL database, consolidating use of the persisted data to decouple the back-end architecture.

"Our tens of thousands of servers and petabytes of data are all in a single datacenter, but we need to move to the cloud."

- Lead a team moving all company data and compute into AWS.
- Architect, design, and build many back-end infrastructure components (IaaS and IAC) to support automated resource management in the cloud, giving some control of the infrastructure to each of hundreds of development teams.
- Integrate operations tooling for hundreds of development teams into the new cloud enterprise infrastructure system.

"Critical data from customers arrives at our servers in the data-center closest to them, but it needs to be rerouted to the proper site for that customer."

- Design and build a highly-available high-use tier-one guaranteed-consistency data location service cluster spanning six datacenters, depended upon by multiple other services, and ultimately the entire company, with 100.0% uptime that lasted for many years.
- Build a sophisticated back-end tool to migrate customer data between data centers, keeping it synchronized and available, even while data continually is ingested during the migration.

"Our data persistence layer is barely functional, and overly coupled, impeding development."

- Build a new Java-based data persistence layer, in the form of containerized services managed by Kubernetes.
- Write encryption library in Java to standardize access to strong AES-256 encryption across multiple layers of the back-end.

"Static portions of the front-end of our web application would benefit from dynamic, customer-specific content."

- Design and build a customizable API-based user-interaction feature that enables easy front-end content updates, by editing rules stored in back-end database tables.

"Our data warehouse isn't sufficiently scalable."

- Build a Hadoop-based data warehouse.
- Build multi-threaded ETL software tools, and various data migration tools.
- Retrofit portions of legacy Oracle-based data warehouse, resulting in a 60x performance improvement.

"Our application was so successful that we are receiving more data per day than we can process in 24 hours!"

- Design and build dataflows for multiple Hadoop clusters, using map-reduce.
- Build a high-speed URL normalization tool.
- Write an HBase application layer.

"Our legacy system accepted billions of bad data records, and we're building a new replacement system, but we need a way to handle a data migration that can deal with the bad data that the new system will reject."

- Build a data migration tool that auto-categorizes moving data into buckets by error type, with auto-retry capabilities once handling of each error type was implemented.

"The core of our back-end enterprise application is the domain layer, and it needs a lot of work."

- Build many back-end services.
- Design and build a vertical feature for capturing end-user-contributed content, connecting persistence all the way to the front-end.

Professional experience in the past 20 years

Enterprise Software Architect with the US Navy Sea Systems Command

Washington, DC (Jan 2024 - present) [navsea.navy.mil]

Worked with, and provided guidance to, dozens of software engineering teams to drive progress on a large enterprise-level software project. Built multiple foundational software libraries and back-end software services used by multiple teams.

Lead Managing Consultant at Oteemo (consulting)

Reston, VA (Jan 2024 - present) [oteemo.com]

Led team of software architects, working with multiple clients (mostly with the US Navy).

Director of Engineering at Ethic (fintech)

Manhattan, NY (2022 - Oct 2023) [ethic.com]

Led and implemented efforts around restructuring the back end to implement engineering best-practices, with a focus on the data layer. Designed and implemented several major software features.

Lead Software Engineer at Marqeta (fintech)

Oakland, CA (2020 - 2022) [marqeta.com]

Built back-end payment card fulfillment software services. Designed and implemented a microservices architecture system.

Team Lead, Senior Software Engineer at Ancestry (genealogy)

Lehi, UT (2016 - 2020) [ancestry.com]

Moved compute and data from a local datacenter into the cloud. Built back-end software services.

Senior Software Engineer at Qualtrics (surveys)

Orem, UT (2014 - 2016) [qualtrics.com]

Built back-end software services, encryption libraries, and datacenter synchronization tools.

Senior Data Warehouse Engineer at Commission Junction (affiliate marketing)

Santa Barbara, CA (2011 - 2013) [cj.com]

Designed and built map-reduce-based data warehouse software systems using Hadoop.

Senior Software Engineer at Tynt (big data)

Draper, UT (2010 - 2011) [33across.com]

Built map-reduce systems for processing vast amounts of Big Data in Hadoop clusters.

Software Engineer at FamilySearch (genealogy)

Salt Lake City, UT (2005 - 2010) [familysearch.org]

Built data migration tools, domain-layer software services, data transformation software, and vertical features.

Education

Bachelor of Science in Computer Science

Clearance

United States Department of Defense: Secret