

Contact

jobs@conway0.com

www.linkedin.com/in/kevin-conway-68a56222 (LinkedIn)

github.com/kevinconway (Portfolio)

kevin.conway0.com (Website)

Kevin Conway

Principal Software Developer

Austin, Texas, United States

(Printed on June 17, 2025)

Experience

Vanta

1 year 3 months

Staff Software Developer March

2025 - Present (4 months) Remote

Staff Software Developer

April 2024 - February 2025 (11 months)

Remote

I worked with the team that manages Vanta's data ingestion pipeline to achieve a set of product scaling goals. This included:

- Developing a software architecture north star to achieve long term scaling goals
- Planning and sequencing the short and mid term work to make incremental progress towards the north star
- Introducing auto-scaling options for core systems
- Helping the team introduce Kafka for event streaming workloads - Mentoring developers

Freelance

Freelance/Consultant

November 2021 - Present (3 years 8 months)

Remote

I work with companies of all sizes, from ground floor startups to larger enterprises, as a freelance principal developer or consultant. Each engagement is unique but the common problems I solve are:

- Advise on and drive execution of large architectural changes - Migrate from legacy systems or databases
- Improve operational excellence and observability practices
- Interview and qualify candidates for software teams or early stage heads of engineering
- Permissions architecture and data management

- Initial system designs and technology choices for new products

SailPoint

Principal Software Engineer

December 2021 - May 2023 (1 year 6 months)

I was primarily contracted to help facilitate the decomposition of a core, legacy system into new, independent services. I worked closely with technical leads and a small IC team to design and build a bidirectional and cross-DB engine data migration framework that was used to synchronize and progressively roll out or back new services. The migration system leveraged change-data-capture and Kafka to both bootstrap new systems with data and provide guaranteed synchronization between new and legacy systems. The bidirectional feature of the migration allowed for progressive replacement of features even while the legacy system remained under use.

I worked with Sailpoint on a part time basis, between 5 and 15 hours a week. My role was primarily to partner with other technical leads to workshop, debate, and document system design and implementation plans. As part of this I helped define system performance expectations and integrated monitoring to measure performance against those expectations. I occasionally provided insight and investigative deep dives into issues related to Kafka, event processing, and system performance both as part of my primary project and as needed for other projects.

InVision

Principal Software Engineer

December 2019 - November 2021 (2 years)

Remote

I was hired into a new team focused on internal improvements to development tooling, developer platform services, and overall system architecture of the product. The team was small and included only senior level and above developers. While I occasionally contributed code, the bulk of the IC work was accomplished the team.

- Took on administration of the company's architecture committee
- Worked with other leaders to expand the use of RFC by teams to solicit outside feedback on planned work
- Helped introduce and facilitate an RCA process for incidents

- Removed multiple weeks of effort from the creation of every new service by standardizing developer tooling, creating template repositories, and automating manual processes
- Introduced change-data-capture as a tool for migrating data between services and mitigating dual-writer related issues
- Developed a managed, ordered event streaming feature based on Kafka to simplify onboarding and expand use of asynchronous event patterns
- Established an API contract and data schema registry that provided backwards compatibility verification, client and server code generation, and API documentation generation

Atlassian

Principal Software Developer

September 2015 - November 2019 (4 years 3 months)

Austin, Texas

I worked on multiple products and teams at Atlassian. I started as a developer and transitioned to technical leadership. Principal was my title at exit. I was promoted during my time on the Stride product.

Security:

Founding member of a new security development team. My role was primarily enablement and advisory focused.

- Set standards for code quality, test automation, logging, telemetry, SLIs/SLOs, and architecture
- Moved development to an open source model
- Collaborated on system design/architecture with members of the teams
- Created a compatibility layer to unify development and operations of serverless and server based systems

Stride:

My role transitioned from IC to technical lead during my work on Stride. I still performed IC work regularly but focused on solving systemic technical issues, developer velocity, and operational maturity.

- Partnered with SRE to develop an observability and monitoring framework for new Go services
- Introduced application level chaos to test resiliency

- Standardized service resiliency patterns
- Subject matter expert for Go and concurrency
- Paired with network engineers to improve service to service latency by orders of magnitude
- Worked with a peer to invent and open source novel improvements to circuit breaking tools
- Frequently wrote or presented on microservice optimizations, testing, and operational maturity

HipChat:

My focus on HipChat was applying my subject matter expertise of async Python and targeting long standing performance bottlenecks and complex bugs.

- Subject matter expert for async Python and the Twisted framework
- Developed context propagation for logging and tracing in Twisted
- Part of a specialized team that solved message delivery issues and product limiting bottlenecks
- Formed an ad-hoc team to debug an asyncio performance issue that resulted in a patch to cpython
- Worked to unfork and upgrade our open source tools
- Regularly wrote about Python and concurrency

The Home Depot

Senior Online Software Engineer

November 2014 - June 2015 (8 months)

I joined a small team within the company that was introducing Python as a development language and cloud infrastructure as a replacement for on-premise servers. This role primarily included improving the development and operations of Python systems with a secondary focus on evaluating cloud infrastructure offerings.

- Reduced multi-day Python system deployment times to under 5 minutes by introducing new packaging and deployment techniques
- Introduced software development best practices including style guide, unit testing, code documentation, packaging/deployment automation, and pluggable architectures
- Pair programmed with developers to drive learning progression and adoption of best practices

- Developed a small framework for rapid development of new REST API micro-services
- Developed an API and browser based GUI for modifying the traffic throttle rules used to roll out new site changes

Rackspace, the Open Cloud Company

Software Developer

June 2013 - November 2014 (1 year 6 months)

I worked with a team of software developers and operations engineers to commercialize an open source database as a service platform. This role included a combination of developing new product offerings, participating in a large open source community, and improving operations.

- Added support for MariaDB, Percona, and MySQL 5.6+ to our existing MySQL 5.5 product line
- Contributed initial support for PostgreSQL to our open source project
- Scaled usage and billing event service by 1,400x messages per day (ask me why that's funny)
- Identified and solved a concurrency issue that improved our Redis provisioning failure rate from over 60% to less than 1%
- Became subject matter expert for solving Python green-thread concurrency issues
- Created proof-of-concept packaging pipeline that eliminated dependency conflicts between services
- Fixed hashing algorithm used by company built, Java based reverse-proxy
- Openstack contributor

Texas State University

4 years 5 months

Programmer Analyst

June 2011 - June 2013 (2 years 1 month)

San Marcos, Texas

I worked in a full stack development and business intelligence reporting role with a team of other business intelligence analysts. Our primary focus was on generating interactive reports on any and all aspects of the business. I was given the opportunity to build a greenfield, proof-of-concept, interactive reporting tool that was later replaced with a commercial solution.

- Designed REST API for querying RDBMS that abstracted underlying SQL engines

- Oracle, MySQL, PostgreSQL, Sybase, and SQL Server
- Performed cross-database engine data migrations
- Developed rich browser frontend for visual data discovery
- Worked in tight feedback loop with customers to choose development features
- Worked with my team to investigate and make available new data sources
- Drafted a proposal for a student employee career development program
- Optimized SQL queries

PC Device Management

February 2011 - July 2011 (6 months)

Service Desk Representative

February 2009 - February 2011 (2 years 1 month)
