

Steven Hicks

📍 Chicago, IL ✉ hackernews@mrmonksy.com

Summary

Driven by a passion for software development that spans over three decades (15 years professional), I focus on delivering reliable, scalable, fault-tolerant, and maintainable systems. I've consistently demonstrated impact across diverse environments, from startups, scale-ups, and large enterprise companies. I'm committed to fostering a learning environment within my teams, enabling them to consistently improve our products and reduce long-term costs.

Experience

Principal Software Engineer

2020 - May 2025

Rally Health/Optum

- Architected, managed scoped/team, and contributed to a total redesign of the legacy Health Activity Rewards pipeline leading to a reduction of software complexity, operational cost with Go, SNS, and Lambda. (Reducing infrastructure spend from \$50k+ to \$8 a year)
- Led, architected, contributed to, and developed a culture for a team to build an GraphQL service (Using Caliban, ZIO, and Scala) from scratch in 2 months handling 300RPM+ reliably (Reducing client integration time from 3 months to 3 days)
- Mentoring throughout the organization in functional programming, system-design, event driven architecture, and Scala concepts: massively overall reducing code complexity and time to delivery
- Reduced high-priority developer support calls from 26/year to 2/year by optimizing the legacy Health Activity Rewards pipeline monolith for resiliency
- Led, architected, and contributed to many complex systems though decomposing complex business requirements into clear delivery plans for speed, flexibility, maintainability, long term value and cost savings, and stability
- Reduced the risk of \$500k+ HIPAA fines by deploying and maintaining an instance of Mapbox on a large Kubernetes environment
- Championed the adoption of a proactive and reactive maintenance culture, leveraging Root Cause Analysis (RCA) to systematically identify and eliminate service failures. Resulted in improved operational efficiency and decreased development time.
- Led, and developed culture for, initiatives to safely upgrade/replace existing services with 0 downtime
- Contributed to, peer reviewed, presented, and authored Architectural Decision Records
- Caught memory leaks, reducing downtime in production and preventing data loss
- Resolved and led P1/2 war rooms with production deep dives spanning many teams

Tech Used:Scala, Play, Akka, Slick, Postgres, MongoDB, Datadog, ZIO, Caliban, Play Framework, RabbitMQ, Scalatest, AWS, AWS, ElasticSearch, RDS, Kubernetes, Splunk, Dynamo DB, Lambda, SQS, SNS, Golang, SNS, Terraform, Github Actions

Lead Software Engineer

2018 - 2020

Zebra Technologies

- Managed a team of 8 engineers to deliver solutions for the SmartPack project
- Created a generic data library and mongo DSL to reduce development time and simplify test creation
- Implemented automated quality checks
- Led initiatives to decrease the time to production for new services via Giter8 templating
- Mentored others learning Scala and best practices across various timezones
- Assisted with creating a Lunch and Learn series
- Found and fixed memory leaks, in one case, reducing memory footprint from 6gb to 250mb (96% savings)
- Reduced overall data pipeline memory footprint from 28gb+ to 12gb (60%+ savings)
- Led initiatives to upgrade the Kafka infrastructure from version 0.11 to 2.3 for better reliability and improved developer experience
- Contributed to the effort to migrate from Docker Compose to Kubernetes

Tech Used:Scala, JSON4s, Mongo, Kafka Streams, Scalatest, Docker, K8s

Contract Engineer

2018 - 2018

SpotHero

- Reduced size of weather data in Redshift, by 99.9% from 2.8TB to 5GB without data loss
- Implemented query and data optimizations that reduced data analysis time for business analysts from hours to minutes
- Improved Airflow jobs by adding new functionality and improving maintainability and stability
- Performed profiling and use case analysis to improve long-running Redshift query run times by 80%

Tech Used:Python, Airflow, AWS Redshift, Database Performance Optimization

Senior Software Engineer

2015 - 2018

HERE

- Drastically reduced cycle and delivery time for features via functional, performance, and stability improvements, increasing unit test coverage and implementing bugfixes
- Reduced time from code release to deployment from 2 days to 1 hour
- Led significant initiatives to increase testing coverage and code reuse
- Implemented alerting to improve stability of a data pipeline
- Served as the communications/education sub-group lead for the Code Review Working Group
- Maintained, contributed to, and enhanced the streaming data pipeline
- Mentored co-workers in working with Scala and ScalaTest

Tech Used:Scala, Redshift, Akka, AWS, EC2, Docker

Software Engineer III

2014 - 2015

GoHealth Insurance

- Automated import, auditing, and validation of data from Excel documents, reducing time to market from 2 weeks to under 1 hour
- Performed deep dive live-site debugging in production systems with AngularJS and Grails

- Researched and resolved performance issues involving multiple services and APIs
- Automating the manual build process with Ruby, Jenkins, and Mercurial, resulting in a cost savings of 104+ engineering hours per year

Tech Used:Groovy, Grails, Postgresql, AngularJS, Excel COM Automation

Solutions Engineer

2011 - 2014

Avnet Technologies/Seamless Tech.

- Built, maintained, and enhanced cloud systems via IT automation with HP automation suite and VMWare-related products
- Successfully tackled many challenges arising from rigid network and security policies, high pressure demands, and tight deadlines at many medium to large organizations
- Led Lunch and Learn sessions to update coworkers about new products within HP suite
- PayChex: Reduced VM deployment time from 1 month to under 1 hour.
- AT&T: Replaced the non-formalized and manual VM build process with automation, reducing process time from 1 week to 30 minutes per VM

Tech Used:HP Server Automation, HP Operations Orchestration, HP Cloud Service Automation, VMWare vSphere, Python, Java, Groovy, Bash

Skills

Languages: Scala, Groovy, Java, C++ (Visual, Managed, GCC), C, C# (.NET/MONO), Golang

Scripting/Markup Languages: Python, Ruby, PHP, Javascript, Bash, CSS, XML, XSLT, XPath, MATLAB, ASP.net, HTML, Prolog, Erlang

Frameworks/Libraries: Kafka Streams, Kafka, ScalaTest, Akka, PicoContainer, Jackson, Win32, JasperReports, JUnit, Spock/BDD, ScalaTest, Grails, PowerMock, GraphQL, Hadoop, Dropwizard, Spock, JSON4S, ZIO, KakfaStreams, Caliban, ElasticSearch

Databases/Data Storage: XML, JSON, MySQL, YML, MS-SQL, Derby, MongoDB, Redshift, DynamoDB, Postgres SQL, Couchbase, SQL Server, HERE Datastore

Tools: Maven, New Relic APM, Airflow, Datadog, CI/CD, SBT, Jenkins, GIT, Mercurial, SVN, CVS, Tomcat, Gitlab, Github, SBT, Gradle, Maven, HPOO, HP-CSA

Cloud Tools Proficiency OpenStack, VMWare vSphere+ESX, AWS ECS, EC2, S3, CloudFormation, Cloudwatch, Redshift, RDS, Docker Compose, Google Cloud Platform (GCP), Kubernetes, Lamdas, RDS Serverless, DynamoDB, Graviton2 Lambda, Terraform

Education

Masters of Science: Computer Science

2008 - 2010

University of North Carolina - Charlotte

Bachelors Of Science: Computer Science

2004 - 2008

Elon University

Certifications

- CompTIA Linux+ Professional
- HP0-M33: HP Operations Orchestration (2012)

Projects

- **Awesome-kafka** : Maintain a list of Kafka resources. github.com/monksy/awesome-kafka - Creator
- **Meetup Broadcaster** : Scala based service that monitors for new events on Meetup, and then broadcasts on Twitter and Reddit. -Creator
- **Financial Market and Strategy Simulation System** : Web backend, data stores, custom test framework, distributed computing, highly concurrent code, backtesting of strategies and visualization previous runs. - Creator
- **Cross Platform EBook Library System** : EBook Library Management System: Aggregating, data scraping, and clustering for better indexing and retrieval. - Creator
- **404Browser**© : Popular Web browser with over 100k downloads (Released in ~2002)- Creator

Presentations

- Streaming Architectures Intro and Best Practices
- Graph QL and You (Intro)
- Caliban Walkthrough
- Innovation Sprint Talk: Maven Transitive Dependency Conflict Resolution (Tools)
- Lunch and Learn: Kafka Streams (A POC and Introduction)
- JUG Talk: I didn't know you could do that with Groovy
- JUG Talk: Testing With Spock
- Hack Week Talk: Akka Clustering
- JUG Talk/Lunch and Learn: Refactoring the Monolith: Alternatives to Microservices
- Provided feedback and reviewed *Programming Kotlin* by Venkat Subramaniam
- Group session: Led group session on Digital ownership and Privacy at jCrete 2019