

Jonathan Camp

Infrastructure Engineer/Software Developer

About Me

Infrastructure Engineer/Software Developer with 18 years building and operating distributed systems in production on AWS, Azure, and on-premises data centers. Decomposing services, drawing API boundaries, building event-driven pipelines; I help organizations scale their systems without a rewrite. I own my systems in production, using telemetry and root-cause analysis to harden the parts that show strain. Strong in Python, Java, and Go, with a track record of learning whatever the problem calls for.

Contact Information

📍 Philadelphia, PA
📞 941-564-9430
✉ jon@skyreach.llc

Skills & Expertise

- **Programming languages:** Python, Java, Go, C#, JavaScript/TypeScript, PHP, SQL
- **Distributed systems:** Service decomposition and API boundaries, event-driven architecture, gRPC, pub/sub (SNS, SQS, EventBridge), event streaming (Kafka, Kinesis), caching and replication strategies, distributed locking, graceful degradation
- **Cloud & operations:** AWS (ECS Fargate, Lambda, DynamoDB, Aurora PostgreSQL, S3, Athena, App Mesh, GovCloud), Azure, Docker, 24/7 production operations
- **Infrastructure as code & delivery:** Terraform, OpenTofu, CloudFormation, GitLab CI, CircleCI, Jenkins, GitOps workflows

Professional Experience

Platform Engineer at Prallax Consulting

November 2024—Current

Platform engineering for healthcare technology under FedRAMP Moderate and FIPS compliance requirements, from cloud infrastructure to data pipelines and internal tooling.

- Cut event volume 95% with a real-time change data capture pipeline (Django signals, Kinesis Firehose, S3, Athena, dbt) using incremental MERGE on Iceberg tables, keeping load off production APIs.
- Deployed a FIPS-compliant analytics platform (Apache Superset) in AWS GovCloud with ECS Fargate behind load balancers in auto-scaling groups, encrypted Aurora PostgreSQL, automated secret rotation, custom SSO integration/user provisioning, and CloudWatch observability, all provisioned with Terraform.
- Built a proof-of-concept AI operations platform bridging Claude to enterprise monitoring via MCP on a serverless AWS stack (Lambda, AppSync, DynamoDB, Bedrock).
- Replaced manual workflows with production services built in Go, Pocketbase, and React; operate Keycloak SSO infrastructure with custom provisioning for downstream systems.

Senior Software Developer at Revcontent

September 2019—November 2024

Cross-functional development across distributed cloud infrastructure, from reliability and cost work to customer-facing features.

- Led migration from a legacy Java monolith to microservices, with Go services speaking gRPC over AWS App Mesh and EventBridge, SNS, and SQS handling async work. Extracted services incrementally along API boundaries; cloud costs fell by over 50% by the time the monolith was retired.
- Maintained and optimized the widget-keyed Kinesis partitioning behind billions of daily content recommendations, keeping widget×geo×device aggregate updaters on bounded, ordered state per shard while raw events flowed intact into Snowflake; the aggregates fed advertiser dashboards, bid pacing, and accounting.
- Migrated the event pipeline from Kinesis to Kafka for cost and operational ownership, after managed-service failures left us with no remediation paths within budget. Moving consumers onto Airflow made the data flow visible, down to individual latencies, side effects, and where things stalled.
- Ran low-latency serving state on a replicated, sharded Aerospike cluster, with ElastiCache caching content-recommendation results, and tuned CloudFront CDN caching to cut data-transfer costs.
- Implemented distributed locking on the ad-delivery path with DynamoDB conditional writes, so concurrent spenders could not push campaigns past budget and a stale lock holder could not commit over newer state.
- Instrumented the microservices with Prometheus metrics, Grafana dashboards, and OpenTelemetry traces into Datadog, then used what the telemetry exposed to raise throughput and cut latency and error rates.
- Load-tested services with Gatling and failure-tested in production via live region shunting, rerouting traffic away from a degraded region and verifying the system held under the shifted load.
- Converted legacy infrastructure to infrastructure as code (Terraform, CloudFormation); modernized a legacy PHP UI with React microfrontends via Module Federation.

- **Data stores & pipelines:** DynamoDB, Aerospike, Aurora PostgreSQL, SQL Server, ElastiCache, Snowflake, Kafka, Kinesis, Airflow, dbt, Apache Iceberg, Solr
- **Telemetry & observability:** Prometheus, Grafana, OpenTelemetry, Datadog, Dynatrace APM, CloudWatch, Splunk, SIEM log aggregation
- **Security & compliance:** FedRAMP Moderate, FIPS, Keycloak SSO (OAuth, JWT), penetration-test remediation, forensic analysis

Senior Software Developer at Syllint

September 2018—September 2019

Security monitoring systems and incident response in a managed security services environment.

- Built out a distributed, event-driven security monitoring system ingesting reports and data into dashboards for analysts, working in Docker on Azure with ASP.NET Core, Angular, Solr, and SQL Server.
- Conducted forensic analysis, root-cause analysis, remediation, and reporting for client security breaches, malware attacks, and vulnerabilities; supported expert-witness testimony in intellectual property litigation.

Lead Engineer at Adrizer

August 2016—September 2018

Led a team of five building a serverless advertising automation platform handling \$50k–\$100k in daily ad spend.

- Re-architected the platform from an Apache Storm cluster to a serverless AWS stack (Lambda, API Gateway, Step Functions, DynamoDB), saving thousands monthly in operating costs.
- Implemented immutable infrastructure with CloudFormation, Secrets Manager, and CircleCI deployment automation, enabling identical staging, QA, and production environments.
- Provided 24/7 operational support with root-cause analysis and systematic prevention of recurring issues.

Senior Software Developer at Ned Davis Research

January 2013—August 2016

Financial analysis platform built on a domain-specific language for quantitative research, using Java, Groovy, Spring, RESTful APIs, and automated testing.

- Replaced an outsourced PDF report-generation system with an in-house Java API, Angular frontend, and D3 visualizations, saving tens of thousands annually with page-for-page parity.
- Championed Dynatrace APM across a distributed architecture after identifying observability gaps, building the proof of concept that won buy-in and training the team on investigative workflows.
- Collaborated with quantitative analysts to implement mathematical models using appropriate software design patterns.

Software Developer at McKesson Automation

November 2011—December 2012

previously – Product Support Specialist II August 2008—November 2011

Agile development on medical automation systems in C#, ASP.NET, VB.NET, and Java.

- Developed services on a team building a medication delivery network platform (ASP.NET MVC, service-oriented architecture, domain-driven design), with testing from unit through performance and end-to-end.
- As Product Support Specialist II (2008–2011), diagnosed and resolved complex production issues across multiple languages and paradigms through root-cause analysis and reproduction test cases.

Achievements

B.S. in Computer Engineering

 Drexel University – College of Engineering, Pennoni Honors College – 2008