

John Siracusa

Senior Software Engineer | Associate Manager

hello@johnsiracusa.dev

github.com/siracusa5 · linkedin.com/in/siracusa5

Senior software engineer and associate manager at impact.com, promoted through four roles since 2018. I lead a small group of engineers on a larger data team and stay hands-on with architecture, design reviews, delivery, and production support. My core work is distributed systems and high-throughput data pipelines on Google Cloud, from first build through migration and upgrade. More recently I have focused on data-platform observability and bringing AI developer tools into everyday engineering work.

Experience

Senior Software Engineer, Associate Manager

impact.com | Mar 2026 - Present

- Manage several engineers on a larger data team: one-on-ones, performance reviews, promotion cases, delivery planning, and hiring.
- Set technical direction for the creator data platform's schema, architecture, and production data, serving 5 core teams and up to 10 partner teams.
- Evaluate AI development tools and advise senior leadership on adoption. Built reusable AI skills and trained 10+ engineers.
- Led the migration of 20+ dimensions from legacy batch processing to the streaming platform, then retired the old jobs.
- Drove the Grafana dashboards and incident-response tooling that became the data group's shared observability framework, mentoring the engineer who built much of it.
- Remain on the production on-call rotation for the data platform.

Data Engineer II

impact.com | Feb 2022 - Mar 2026

- Led batch-to-streaming migrations, including an action-enrichment rebuild that cut latency from 2+ hours to under 5 minutes.
- Ran and extended high-throughput data infrastructure behind client reporting, internal analytics, and data-science workloads at platform scale.
- Built 7+ creator data pipelines and datasets feeding brand and publisher dashboards, internal analytics, and standard and custom client reports.
- Designed a historical data lake with years of change history for accurate point-in-time reporting.
- Built production auditing that helps teams hit service-level objectives and cut time on manual checks and investigations.

Data Engineer I

impact.com | Feb 2020 - Feb 2022

- Migrated Identity, Attribution, and Data Platform services from on-prem Hadoop and Cloudera to Google Cloud, then decommissioned the legacy systems.

Junior Software Engineer

impact.com | May 2018 - Feb 2020

- Built metrics workflows, streaming alerts, and operational dashboards in Scala and shell that surfaced failures earlier.
 - Joined the production on-call rotation early and handled live incidents.
-

Education

The Ohio State University

B.S. Electrical and Computer Engineering | 2018

Earlier Experience

Embedded Systems Engineer Intern

Lexmark International | Jan 2017 - May 2017

- Designed and built a DC power meter to evaluate printer power-supply efficiency, with a custom PCB that met form-factor constraints across several printer lines.
-

Independent Work

Mind Your Blanks

Games for any group

Eight live and solo games with team leaderboards, host controls, content tools, analytics, and an admin dashboard, operated with metrics, alerting, synthetic probes, nightly backups, and automated rollback.

ContextCake

Context infrastructure for AI agents

A dependency-free context cascade that merges company, team, and personal context with clear precedence, shipped as a signed Mac app that connects to AI clients over MCP.

2 Girls 1 Parlay

Sports picks with a public record

A sports prediction site with automated grading, public results, a daily picks hub with cross-game parlays, a player-prop research desk, a live leg-by-leg board, and payments and fulfillment through Stripe and Printful.

Hometown Media

Website for a Tennessee sports media team

Site design and front-end work built around photography, clear services, and an easier way for teams, athletes, and coaches to get in touch.

Technical Skills

Languages: Go, Python, TypeScript, Scala, Java, SQL, Bash/Shell

Data and streaming: Kafka, Spark, SingleStore, MySQL, Bigtable, Google Cloud Storage, data quality, observability

Cloud and delivery: Google Cloud Platform, Dataproc, Cloud Logging, Docker, GitHub Actions, CI/CD

AI and developer tools: Anthropic API, Claude Code, MCP, Ollama, embeddings, GitHub Copilot, Cursor