

Kevin Scaldeferri

Portland, OR

Email: kevin@scaldeferri.com

Phone: (626) 487-1763

URL: <http://kevin.scaldeferri.com/resume/>

Professional Objective

I am not currently looking for work. I am also not interested in relocating. Really.

An individual contributor position where I can use my expertise in large system development, scaling, and optimization; where initiative and self-sufficiency are valued; and with co-workers dedicated to continuous learning and improvement.

Employment History**Lead Software Engineer**

New Relic, Core Data Platform Group

October 2014–Present

The Core Data Platform is responsible for storing and querying all of the data which New Relic collects on behalf of our customers. I've worked throughout this group on multiple major components of our systems, as well as contributing to shared libraries throughout the company.

Projects:

- I lead the design and development of our streaming NRQL evaluation system. This involved adapting the core NRQL execution engine to be suitable for continuous streaming execution, as well as one-off queries; designing a Kafka-based streaming topology for highly multi-tenant workloads; and implementing key distributed consensus algorithms and snapshotting techniques to enable zero-downtime deployments of updates to the system.
- As a member of the NRDB team I worked on a number of projects including improvements to the reliability of queries against leading-edge data; working with other internal teams to transition them off of a legacy HTTP-based insert API to our current Kafka-based API; contributing to the client libraries for both inserts and queries to the system; designing rate limiting systems appropriate for a highly multi-tenant database; and implementing various new features and improvements to the NRQL language
- I am a co-leader of our JVM Community of Practice. This involves organizing monthly meetings and twice-yearly offsites, and assisting in the formation and activities of working groups within the community of practice. As part of this work, I also contribute regularly to our shared tooling and libraries, and help guide teams to keep up-to-date with the current versions and best practices for use of these libraries.
- I served as an internal consultant during the initial development of our continuous delivery system

Principal Software Engineer

Gilt Groupe, Architecture Team

September 2008–October 2014

The Gilt Groupe is a leading online flash sale company focusing on designer clothing, upscale home furnishings, and travel. As a member of the architecture team, I worked on a variety of projects core to the business and technical infrastructure.

Projects:

- I architected a solution for real-time event processing and implemented tools and libraries to support event generation and consumption. This involved researching, prototyping, and evaluation of a variety of tools such as Kafka, Kinesis, Spark, Summingbird, and Avro.
- I designed and built Gilt's continuous delivery system which tests and deploys new versions of hundreds of microservices. I also helped develop an SBT-based build system for our microservices which integrates with the continuous delivery system. I have spoken about this work at multiple conferences.
- For 3 years, I was the lead developer for Gilt's checkout and order management systems, responsible for ensuring the availability and scalability of core e-commerce functions under extreme daily load spikes. Major initiatives included migrating performance critical components to microservices on the JVM, supporting the transition to multiple warehouses and drop-ship vendors, enabling international shipping and payment in multiple currencies, and integrating third-party providers for tax calculations.
- I helped establish Gilt's Portland office, including recruiting and interviewing all engineering candidates, on-boarding new hires, creating a team culture, and managing 4 other developers.

- I helped scale Gilt's original Ruby on Rails application by applying various performance analysis techniques and optimizing based on the results. This work led to contributions to Rails and other open source libraries to improve aspects of their performance.
- Champion for code review, unit testing, continuous integration, and agile development processes.

Senior Developer

*Yahoo! Search Marketing, Right Media Group
September 2007–September 2008*

After the acquisition of Right Media by Yahoo!, I worked as a technical liaison to integrate the Right Media Ad Server (RMAS) system into the Yahoo! infrastructure

Projects:

- I decomposed the monolithic RMAS codebase into several smaller software packages, and converted those packages to use the standard Yahoo! build, packaging, and deployment tools.
- I analyzed performance and memory usage of the RMAS system, and implemented core algorithmic improvements which yielded an order of magnitude improvement in memory usage and churn and a dramatic concomitant reduction in 99th percentile response times.

Senior Developer

*Yahoo! Search Marketing, Content Match Group
July 2005–September 2007*

The Content Match group provides contextually targeted advertisements for websites both within and outside the Yahoo! network. I developed high-volume, high-reliability server applications using C and Perl.

Projects:

- I designed and led the implementation of a system for inserting advertisements into RSS feeds. This required novel modifications of existing ad delivery mechanisms to balance desirable business features such as learning and ad rotation while preserving an acceptable user experience.
- I created the design for geo-targeting advertisements, and specified the required API changes between distributed components as well as the required functional changes in those components. I implemented the geo-targeting component for converting IP addresses to geo-regions, and supervised the implementation of the rest of the changes.
- I led a 4 person team to implement private marketplaces within the Content Match system. This system allows high-quality advertisers to bid for placement on selected high-quality publisher sites.
- I designed and implemented a system for performing live diagnostics on our distributed clusters. This system allows business people and developers to perform detailed queries on the state of a cluster to help diagnose problems in the system.
- I worked on a number of infrastructure improvements to streamline development of our system, including process automation and elimination of boilerplate code. I was also one of a small group of “blessed” code reviewers given the responsibility to approve all code for production use.

Senior Developer

*Yahoo! Search Marketing (formerly Overture Services), Sponsored Search Group
April 2002–July 2005*

The Sponsored Search system at YSM serves hundreds of millions of web requests daily, providing search results directly as HTML to users and via XML feeds to partners including ISPs, portals, and browser plugins. On this team, I developed systems using Perl, Java, and C/C++; running within Apache/mod_perl, Dynamo, and Weblogic; and interfacing with Oracle and BerkeleyDB databases.

Projects:

- I worked on a number of projects to improve search coverage using approximate query matching techniques, including character edit transformations, subphrase matching, and semantic matching based on extremely short query strings. This required integration of fast C libraries with the Perl codebase, and extensive performance and scalability testing. I developed data denormalization schemes to help meet performance requirements, and various heuristics to ensure reliable system operations and acceptable resource utilization while maintaining high coverage and relevance.
- I led the internationalization of the search systems during expansion into France and Japan and conversion of all internal systems from ISO-8859-1 to UTF-8. This required upgrading several base systems, including Java, Perl, and Dynamo; and developing extensive regression tests for all parts of the search systems. I identified frequent trouble spots in the code base and researched and developed internationalization best practices. I specified the external API for character encoding support. I also created wrappers for several non-internationalized third-party modules.

- I was involved in evaluating and integrating third-party libraries for Asian language support, particularly morphological analyzers. I worked closely with developers at an external vendor to provide feedback from our experiences, both in terms of functional correctness and production reliability, which they could use to improve future versions of their software.
- I collaborated with our QA and release management teams on several projects, including enhancing and extending the toolset for functional and regression testing, and consulting on the requirements and design for a load testing framework. I also developed a system for interspersing code launches with more frequent metadata launches.
- I redesigned and refactored large portions of the search pipeline to make better use of object-oriented and metadata-driven programming techniques. In net, I eliminated several thousand lines of largely duplicated code with no negative effect on performance, and significantly improved readability and maintainability.

Senior Developer / Technical Lead

Arsdigita Corporation

January 2000–February 2002

I was a member of the development team for the Arsdigita Community System, a web application framework for developing database-backed community and collaborative web services. I developed software using Java, TCL, SQL, PL/SQL, Oracle, Tomcat, and Aolserver.

Projects:

- I was the technical lead of a 6-person team focused on the development of a rules-based personalization system incorporating dynamic classification of users based on behavior and targeted delivery of content. Responsible for design and code reviews within this team and for other teams.
- I designed and developed several modules to support various forms of web-based collaboration, including: a system for authoring of manuals and other publications incorporating user comments and feedback; a repository for code sharing and reuse with capabilities for feedback, ratings, and bug tracking; a workflow engine based on the Petri net model; a file storage system with fine-grained permissioning and commentability; and a threaded bulletin board system.

Web Developer

Graduate Student Council, Caltech

November 1999–June 2001

I reengineered the web version of the Technique, the graduate student guide for Caltech. When I started this project, the web site was simply a static HTML dump of the print version. I designed a content management system appropriate for the needs of this publication including full-text searching of the contents of the book, reader comments on the text, and a searchable database of restaurant reviews.

Systems Administrator

Theory Group, Physics Dept., Caltech

June 1999–August 2000

I supported approximately 40 users in a mixed environment of Sun servers, thin clients, and PCs. I was responsible for hardware purchasing decisions, software installation and upgrades, backups, and security on the cluster. In addition, I provided documentation of the use of software supported on the systems.

Education

BS in Physics, May 1997. with High Honors
University of Maryland at College Park

BS in Mathematics, May 1997. with Honors
University of Maryland at College Park

MS in Physics, June 2000
California Institute of Technology

Programming Languages

Used professionally: Scala, Java, Ruby, C/C++, Perl, SQL, PL/SQL, XML/XSL, TCL, Fortran, IDL.

Limited Experience: Go, Python, Haskell, Erlang, OCaml, Lisp, MIPS Assembly, R.

Conference Talks

- "Overdoing Microservices: How Small Is Too Small?". *Open Source Bridge*. 2016.
- "Monoids, and Sketches, and CRDTs, oh my!". *Open Source Bridge*. 2016.
- "Continuous Delivery and Large Microservice Architectures: Reflections on Ioncannon". *OSCON*. 2015.
- "Innovating Faster with a Micro-Service Architecture using SBT, Continuous Delivery, and LXC". *Open Source Bridge*. 2013.
- "Enabling Microservice Architectures with Scala". *Lambda Jam, CUFP*. 2013.
- "Beautiful Concurrency with Erlang". *OSCON*. 2008.
- "How to Serve a Billion Requests a Day with Perl". *OSCON, YAPC::NA*. 2005.