

Kevin Scaldeferri

Portland, OR

Email: kevin@scaldeferri.com

Phone: (626) 487-1763

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

Professional Objective

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**Senior Developer**

Yahoo! Search Marketing, Content Match Group

July 2005 – Present

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 the 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: Perl, C/C++, Java, SQL (Oracle), PL/SQL, XML/XSL, TCL, Fortran, IDL.

Limited Experience: Lisp, OCaml, Haskell, Erlang.

Award, Honors, and Certifications

- **Overture Star Performer**, 2004
- **Sun Java Certified Programmer**
- **Graduate Research Fellow**, National Science Foundation
- **Phi Beta Kappa**
- **Best Honors Thesis**, University Honors Program, UMCP
- **Best Undergraduate Research Project**, Physics Department, UMCP

Publications

- "How to Serve a Billion Requests a Day with Perl". *OSCON, YAPC::NA*. 2005.

