

*On the job experience:*

- Java since 1997, Linux, SQL, Perl, JSON, XML, Avro, REST APIs, Agile, Puppet
- Elasticsearch, Lucene, SOLR, MySQL, Clustrix, Spring MVC, Hibernate, Kafka, Hadoop, Spark
- NoSQL since 2009: Redis, memcached, MongoDB, HBase, Cassandra, DynamoDB
- Distributed. high-performance, highly available, multi-threaded, loosely-coupled, service oriented, scalable architectures

Jan 2017 – Present: **Leapfire**, Sunnyvale, CA, Elasticsearch Consulting Architect

Training; Development coaching on site or video conference; Data modeling: schema and query development; Performance testing, proof-of-concept and feasibility studies; Making good decisions on tradeoffs for use cases; Cluster topology, configuration and operations; Capacity and growth planning; Monitoring, Alerting, Security, Backup/Restore, Disaster Recovery; Diagnostics and optimization; Leveraging the whole Elastic Stack including Logstash and Kibana; Integration with Kafka; Cross Data-Center replication; Big-data analysis with Spark-Elasticsearch.

Feb 2014 – Dec 2017: **Symantec**, Mountain View, CA, *Senior Principal Software Engineer*

As the Distributed Data Architect for consumer products division Norton Engineering, I consulted for many development groups in various stages of progress.

I designed the *Norton Distributed Data Stack* (NDD Stack) as a cloud-ready reference architecture for cost-effective scaling of big data. NDD Stack, comprising open source components that complement each other, is an excellent choice for many Use Cases. *Cassandra* for cross-data-center Key-Value Store, *Elasticsearch* as an Indexing Engine with unparalleled ad-hoc query performance, and *Spark* as a Compute Engine for sorting, joining, and map/reduce. We use message queues RabbitMQ and Kafka as appropriate.

I collaborated with product software architects to develop their application architecture adapting the NDD Stack to their Use Cases. I did data modeling, NDD training presentations, performance testing, cluster topology/operations/capacity planning, and more. I am taking whole Java-Oracle teams and transitioning them to become good distributed-database programmers. We've deployed NDD Stacks in four internal data centers as well as clouds OpenStack, AWS and Azure. I was Symantec's leading Elasticsearch expert and a member of Elastic's Customer Advisory Counsel.

Previously, in a team of 12 developers building Java back-end for cloud-based file backup. Led the development of Search REST API product features and Event/Activity Stream analysis. Lead for Scaling and Performance. Investigation and Performance Testing of: Search Clustering SOLR versus Elasticsearch, Sharded Relational MySQL Fabric/Pythos versus Clustrix MySQL, Event System in Cassandra and Elasticsearch, Object Store Swift.

Elastic asks "How did Geena become so knowledgeable of Elasticsearch?"

<https://www.elastic.co/blog/elasticsearch-support-an-investment-that-keeps-paying-off-at-symantec>

May 2013 – Jan 2014: **Disney Interactive**, Palo Alto, CA, *Lead Software Engineer*

Developed Java back-end REST APIs for use across many mobile apps in Java with MongoDB.

Jan 2012 – Nov 2012: **Yahoo!, Inc.**, Sunnyvale, CA, *Software Engineer*

Developed Hadoop software in Java for behavioral ad targeting.

Oct 2010 – Aug 2011: **Software Consultant**

Dec 2007 – Sept 2010: **CBS Interactive**, San Francisco, CA, *Software Architect*

Software architect and Lead Engineer for team that replaced the old static web site to make CBS.com a dynamic site with social activities. "During her three years with the network's digital group, revenue increased by 700% and the network's portal, CBS.com, increased its audience tenfold." Designed the CBS.com back-end as a Java Tomcat application with Spring MVC, Hibernate, MySQL, Linux, Perl, JSON, XML, memcached, Redis and Lucene.

Architected and personally implemented several mission-critical feature sets, doing all coding myself:

- Queryable Video Feed (QVF) meta-data integrator and query engine for all videos on CBS.com.
- Configurable Fantasy Games for TV Shows similar to Fantasy Football (e.g., Fantasy Survivor, Fantasy Amazing Race, ...)
- Contests (e.g., to predict Grammy Award winners).
- Integrated CBS.com with Facebook Connect doing all front-end Javascript and back-end Java work.
- CBS.com's first-ever Site Search, using Lucene and exposed as a JSON REST API (Search as a Service).
- Extreme Voting API for high-traffic high-stakes events (e.g., Grammy and Academy of Country Music Awards Shows).

Collaborated on feature development and interaction with marketing product managers and software engineers within our team, outside CBS groups and partner companies such as Grammy, ACM, Victoria's Secret, Sprint, and others.

July 2006 – Nov. 2007: **MetroFi**, Mountain View, CA, *Software Architect*

The municipal WiFi broadband ISP with free ad-supported networks nationwide.

Ad Server in Java, JSPs, JavaScript and AJAX.

Nov. 1999 – Jan. 2006

**Tribal Fusion**, Emeryville, CA, *Chief Contextual Architect*

**Leapfire Technologies, Inc.**, Sunnyvale, CA, *Founder, Software Architect*

*Tribal Fusion acquired Leapfire Technologies in April 2004.*

Principal Technical Lead of the Context Project. Joined when Tribal Fusion had 15 employees; by 2006 grew to over 100 employees and served 11 billion monthly impressions. The *Context Software* implements linguistic text analysis to target advertising to the content of web pages. It is a highly-available distributed multi-threaded Java system using AJAX techniques for ad displays, Java Tomcat application server, and JDBC to Oracle database (first version used MySQL).

At Leapfire Technologies, designed and implemented software that links products, text ads, and relevant information to content of web pages. It is a proprietary linguistics-based text-filtering system built with Java using scalable distributed multi-threaded implementation. Text-ads campaign on Associated Press had an astounding 1.68% Click-Through Rate over six months across a variety of news areas including health, entertainment, business, sports, technology and more.

Leapfire Technologies placed text ads on AP Wire stories on over 400 newspaper web sites nationwide. Designed and implemented proprietary map-reduce application to implement distributed text-analysis algorithms (*e.g.*, to compute document frequencies of two-word phrases over a corpus of Associated Press news articles over a three-year period).

Aug. 1997 – Sept. 1999

**Inktomi Corporation**, San Mateo, CA, *Software Engineering Manager*

**C2B Technologies, Inc.**, San Mateo, CA, *Director of Software Technology*

*Inktomi acquired C2B in Sept 1998*

Designed the first version of the *Shopping Engine*, Inktomi's third product line. Lead software engineer for the *Integrated Order Process* (IOP) designed to integrate portals' shopping applications with merchants' web stores. All projects were in Java.

At C2B, lead software engineer for the *Shopping Engine*, C2B's flagship product, a pioneering *comparison-shopping* query engine. Designed the overall software architecture for this system, detailed design, and much of the software development. Did presentations and technical due diligence for prospective customers and investors.

The C2B Shopping Engine is implemented as multi-threaded Java Servlets and connects to a relational database system through JDBC, and a geographic database through sockets. Invented (independently at a similar time with Jungle) the robot technology known as *Screen Scraping*; developed a spidering and scraping technology used to build the multi-merchant product database and to query merchant web sites in real-time.

#### **US Patents Issued:** \* *Primary Inventor*

7,415,429 \* Providing navigation objects for communications over a network

7,363,248 \* Pre-filling order forms for transactions over a communications network

7,349,867 \* Tracking transactions by using addresses in a communications network

6,385,602 Presentation of search results using dynamic categorization

#### **Major Accomplishments that I Designed the architecture, Developed code, and Launched to production:**

2015: designed topology for the production Elasticsearch clusters for all of Symantec, enterprise and consumer.

2009: first back-end for CBS.com. Grew from 3M visitors per month to 25M in two months.

2001: first Contextual Advertising Server two years ahead of Google Ad Sense; company acquired.

1998: one of the first (of three) Comparison Shopping Engines; company acquired.

#### **Refereed Publications**

- D. Jackson and G. L. Rollins, "A New Model of Program Dependences for Reverse Engineering", 12/1994.
- D. Jackson and G. L. Rollins, "Abstraction Mechanisms for Pictorial Slicing", 11/1994.
- R. Harper, P. Lee, F. Pfenning, and G. L. Rollins, "A Compilation Manager for SML/NJ", 6/1994.
- J. M. Wing, G. L. Rollins and A. M. Zaremski, "Thoughts on Larch/ML and a new application for LP", 7/1992.
- G. L. Rollins, "SourceGroup: A Selective Recompile System for SML", 9/1991.
- G. L. Rollins and J. M. Wing, "Specifications as Search Keys for Software Libraries", 6/1991.
- G. L. Rollins, "A Simple System for Object Storage in Common Lisp", 6/1991.
- P. Lee, F. Pfenning, G. L. Rollins and W. S. Scherlis, "The Ergo Support System: An integrated set of tools for prototyping integrated environments", 11/1988.

Gave talks on some of these publications at major Computer Science conferences.