

## Summary

- A dynamic, hands-on technical leader, public speaker, and technology evangelist specializing in Big Data and Data Science
- Over 30 years experience in software development, systems operations, and business ownership gives a unique perspective on how to create and deploy highly reliable systems that deliver a positive return on investment
- Extensive career commitment to staff development, team building, and knowledge transfer

## Key Skills

### AWS

Elastic MapReduce  
DynamoDB  
Lambda  
Kinesis  
Redshift  
Athena  
ElasticSearch  
ElasticCache

### Big Data

Hadoop  
Hive  
Spark  
HBase  
Presto  
Storm  
Splunk  
ELK Stack

### Languages

Perl  
Python  
Node.js  
SQL (T-SQL, PL/SQL, pgSQL)  
Unix Shell  
C  
LAMP Stack  
HTML/CSS

### Databases

Cassandra  
Redis  
Greenplum DB  
Vertica  
PostgreSQL  
MySQL  
MS-SQL  
Sybase

## Experience

### Amazon Web Services (September 2015 - present)

#### Big Data Team Lead

- Hands-on technical leader for team of engineers supporting Elastic MapReduce, DynamoDB, Kinesis, RedShift, and other AWS Big Data products
- Help grow team from 3 engineers to over 30 by screening, interviewing, and recruiting candidates
- Develop and deliver training for support operations; mentor junior team members in supported services
- Act as subject matter expert for DynamoDB and Elastic MapReduce; serve as escalation point for critical, high-profile support cases
- Develop Machine Learning application to rate candidate resumes based on historical interview outcomes

### Digital Moonlight LLC (July 2014 - September 2015)

#### Principal Systems Architect

- Client: *RocketFuel*
  - Lead effort to develop tools to automate administration of multiple production Hadoop clusters with a combined total of 2000+ nodes
- Client: *Kickback Rewards Systems*
  - Evaluate end-to-end data flow of legacy system
  - Provide recommendation for technology stack and target architecture within Hadoop ecosystem
  - Create proof of concept application to demonstrate feasibility of recommended architecture
- Client: *Clearedge IT Solutions*
  - Build and administer Hadoop / Accumulo cluster to test integration of Zipkin and HTrace with Accumulo

### ClearEdge IT Solutions (September 2013 - July 2014)

#### Principal Solutions Architect

- Create Accumulo client for Perl using Thrift proxy
- Create tool to migrate HBase installation to Accumulo
- Client: *Laureate Education*
  - Install, configure, and tune Hortonworks 2.0 Hadoop cluster
- Client: *Xcel Energy*
  - Lead team to create PoC of round-trip ETL processing between legacy RDBMS systems and Hadoop ecosystem
  - Create Bayesian classifier to demonstrate Natural Language Processing (NLP) analysis of call center notes
- Client: *Purdue University*
  - Create ETL process to anonymize a diverse set of legacy data and load into Greenplum DB and Hadoop

**Impetus Technologies** (June 2011 - September 2013)

**Senior Systems Architect**

- Create and deliver client-facing Big Data and NoSQL learning products for consulting practice
- Conduct ongoing research into new developments and best practices in Big Data, Hadoop, and NoSQL
- Client: *Bridgewater Associates*
  - Evaluate solutions for monitoring legacy applications
  - Identify technology stack and define architecture
  - Create PoC using Splunk, Perl, and Python
- Client: *EMC/Greenplum*
  - Install and configure 24-node Hadoop cluster for competitive PoC project for Global Electronics Manufacturer
  - Write and tune Hive queries for specific use cases
- Client: *Heartland Payment Systems*
  - Lead project to identify highly scalable MPP database product to replace MS-SQL data warehouse
  - Gather requirements, solicit vendor responses, evaluate products, and deliver recommendations
- Client: *American Express*
  - Lead team creating MapReduce jobs to produce similarity matrixes from financial transaction data
  - Act as on-site liaison between client and offshore developers
  - Tune and configure 70+ node MapR Hadoop Cluster
  - Develop ETL and admin tools
- Client: *SAND Technology*
  - Lead project to integrate client's Analytic Database product with Hadoop
  - Identify and design specific components, provide oversight of offshore developers

**Digital Moonlight LLC** (March 2009 - December 2011): **Principal Systems Architect**

**America Online** (December 2005 - March 2009): **Senior Systems Engineer, Senior Database Administrator**

**Rapier Technology** (November 2003 - September 2005): **Principal / Senior Database Engineer**

**Benelogic** (March 2002 - November 2003): **Senior Database Engineer**

**Rapier Technology** (February 1999 - November 2003): **Principal / Senior Database Engineer**

**Maryland Insurance Administration** (May 1997 - February 1999): **Lead Database Administrator**

**Automation Consultants** (July 1996 - February 1997): **Senior Software Engineer**

**Rite Aid** (February 1996 - July 1996): **Senior Systems Engineer**

**Automation Consultants** (September 1995 - February 1996): **Senior Software Engineer**

**University of New Mexico Hospital** (December 1994 - August 1995): **Software Engineer**

**Techsource Consulting** (November 1993 - August 1994): **Senior Programmer / Analyst**

**United States Air Force** (October 1989 - July 1993): **Programmer / Analyst, Software Tester**

## Education

**University of New Mexico**

**Computer Engineering**

8/1993 - 5/1995

**Community College of the Air Force**

**Computer Science**

10/1989 - 2/1993