

Summary

- A creative, hands-on technical leader specializing in Big Data applications, Cloud Computing, and highly scalable systems
- Over 25 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

Big Data

Hadoop
Hive
Pig
Flume
Storm
Splunk

Languages

Perl
Ruby
SQL (T-SQL, PL/SQL, pgSQL)
Unix Shell
C
LAMP Stack

Databases

HBase
Cassandra
Greenplum DB
Amazon DynamoDB
Vertica
Postgres

General

Teaching and Public Speaking
Amazon AWS Stack
Database Administration
Linux Systems Administration
Cloud Computing
Functional Programming

Experience

Digital Moonlight LLC (July 2014 - Present)

Principal Systems Architect

- Client: Space Dynamics Laboratory
 - Lead effort to select next-generation technology stack for real-time processing of satellite imagery and metadata
 - Implement and benchmark proof-of-concept model using Apache Kafka and Storm
- 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

- Provide general technology and business consulting services to multiple small businesses
- Client: Boston University CDIA
 - Teach web programming and database development courses
- Client: NetApp
 - Gather requirements for PoC project to replace Oracle data warehouse with Hadoop
 - Install and tune 40-node Cloudera Hadoop cluster
 - Develop Hive queries and streaming MapReduce jobs in Perl
- Client: Catalist
 - Develop PoC to demonstrate migration of Oracle data warehouse to Hadoop ecosystem
 - Develop ETL tools to move data from Oracle to HBase and Hive
 - Develop analytics application using Streaming MapReduce
- Client: GuestMetrics
 - Lead effort to re-implement an unscalable client-server system using Amazon Elastic MapReduce, S3, DynamoDB, and other linearly scalable cloud-based services

America Online (December 2005 - March 2009)

Senior Systems Engineer

Senior Database Administrator

- Develop tools to generate XML sitemaps for millions of user-created web pages
- Lead database designer for AIMPAGES.com, AOL Member Directory 2.0, and the tagging & commenting infrastructure used by multiple AOL properties
- Develop ETL, administrative, and reporting tools in Perl, C, and Unix Shell
- Extend and maintain ORB, AOL's "master control" system for systems operations
- Write and tune stored procedures for Sybase, MS-SQL, MySQL and Postgres
- Administer 50+ replicated Sybase, MS-SQL, Postgres, and MySQL servers supporting the AOL.com portal site as well as other AOL branded properties
- Architect sharded RDBMS clusters based on analysis of projected and actual loads

Rapier Technology (November 2003 - September 2005)

Principal / Senior Database Engineer

- Client: National Center for Biotechnology Information
 - Administer exabyte-scale sharded RDBMSs supporting ENTREZ and PubMed, the world's largest public repositories of genomic data and medical journal abstracts
 - Develop tools to automate database and system administration
- Client: The Rouse Corporation
 - Develop ETL and reporting tools to port data from legacy ERP/BI system to next-generation replacement

Benelogic (March 2002 - November 2003)

Senior Database Engineer

- Lead database designer for company's namesake product, a web-based benefits enrollment system
- Designed generic security and authentication module for web-based applications
- Create disaster recovery plans and coding standards for database systems
- Administer and tune production, test, and development databases

Rapier Technology (February 1999 - November 2003)

Principal / Senior Database Engineer

- Client: Verizon
 - Administer intranet training site serving 200,000+ employees
 - Develop administrative and reporting tools
- Client: Dakota Imaging
 - Develop Transact-SQL stored procedures
 - Conduct performance tuning and troubleshooting of deployed product installations
- Client: Bell Atlantic
 - Port Oracle PL-SQL stored procedures and schema to Transact-SQL; optimized performance of ported code
- Client: Small Business Administration
 - Lead a team that created a web-based reporting interface to an existing Data Warehouse
 - Serve as point of escalation and technical resource for operational DBAs

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