

Providing practical and innovative technology solutions

SUMMARY

Over the past 20+ years, I have developed a very sought after combination of technical, business, management and people skills. My core competencies include, but are not limited to the following:

Agile Scrum (Certified ScrumMaster), Big Data, NoSQL, Data Architecture, Backend Distributed Systems Design & Development, Data Warehousing, Business Intelligence, ETL, Data Integration, Master Data Management, Data Governance, Database Development & Administration and Database Performance Tuning & Optimization.

EXPERIENCE

Inivata, Inc.
Research Triangle Park, NC

VP of IT and Software Development
10/2017 – Present

Inivata is a global clinical cancer genomics company employing the precision of circulating tumor DNA (ctDNA) to improve personalized healthcare in oncology.

- Providing technical leadership and management for multiple software development teams in RTP, NC and Cambridge, UK, to develop & maintain systems to support high-complexity, genomic lab diagnostic tests.
- Improved Software Development Life Cycle (SDLC) by implementing Agile Scrum framework and best practices.
- Developed three year strategic plan and budget for Global IT organization to include new systems development, enterprise business systems implementation & integration (Salesforce CRM, Data Warehouse, Finance), systems support and infrastructure services.

Distil Networks, Inc.
Raleigh, NC

Big Data Architect, Director of Data Engineering
10/2014 – 10/2017

Distil provides proactive and precise bot mitigation for web applications, mobile applications and APIs. Responsible for the architecture, design and development of our Data Platform to support internal and external products.

- Design, development, implementation and support of a Cloudera Hadoop ecosystem to ingest, enrich and store 2-3 billion records per day. The primary components of the Data Platform include HDFS, Hive, Impala, YARN, Zookeeper, Kafka, Hue and was developed using Java 8. The Data Platform is used by data scientists, analysts, support and engineering for research, analysis and data mining.
- Designed, developed and implemented a near real-time, distributed, streaming data pipeline for our machine learning products to improve bot detection. The system was developed in Java 8 along with Kafka and Redis.
- Designed and built job management system for Data Platform using Python, Fabric, Supervisor and crontab.
- Built functional prototypes against several NoSQL and distributed databases as a proof of concept for an advanced reporting engine.

Java 8, Python, Fabric, Supervisor, Cloudera 5.10, HDFS, Impala, YARN, Kafka, Zookeeper, Hue, Redis, Lua, CrateDB, CitusDB, Cassandra, Couchbase, Riak, PostgreSQL, Redshift, Linux, Ubuntu, Amazon Web Services, OpenStack, Github

Bridge2 Solutions provides marketing and fulfillment solutions to leading financial institutions, loyalty rewards programs and Fortune 500 companies. Contracted to evaluate the existing ETL system and develop a new agile and scalable architecture to support the growing data demands and business requirements.

- Design, development and implementation of new, high-performance Groovy ETL integration platform utilizing Groovy, Spring, GParas, MongoDB, Spock and Gradle as the primary technologies.
- Migrated existing legacy Talend ETL jobs to new ETL platform to include 25+ merchant data feeds and the generation of several client product catalogs.
- Responsible for MongoDB design, development and support to include performance tuning and optimization. Also responsible for Groovy application performance tuning and optimization.
- Creation and support of Rackspace Cloud non-production environments to include ~20 CentOS Linux servers.
- Provide technical guidance, code reviews and support for Alpharetta and off shore (Moldova) ETL development teams.

Groovy 2.2.2, MongoDB 2.6, Spring, Spock, GParas, Gradle, Git, JSON, XML, Rackspace, Linux

The 6fusion SaaS platform provides metering and reporting on both virtual and physical IT infrastructures along with providing a marketplace for both IT infrastructure sellers and buyers. As the data architect, I was solely responsible for the architecture, design, development, administration and support of both the SaaS platform databases and Data Warehouse databases to include, but not limited to the following:

- Design, development and implementation of PostgreSQL Data Warehouse platform. Design and development of custom, high-performance ETL tool written in Ruby on Linux platform. Design and development of custom Reporting Engine written in Ruby to push HTML, CSV and Excel reports.
- Responsible for SQL Server 2008 and PostgreSQL database design, development, back-end/front-end developer support, performance tuning & optimization, database administration and Ruby on Rails ActiveRecord migrations for 6fusion SaaS platform.
- Design and implemented real-time, dynamic change data capture (CDC) solution for PostgreSQL.
- Research and prototype development of Big Data / NoSQL solutions to include Hadoop (Cloudera), Riak, MongoDB and RethinkDB.
- Continuous evaluation and development using alternative technologies to include the following: Amazon Web Services, MySQL, Percona Server, SQL Server Reporting Services (SSRS), SQL Server Integration Services (SSIS), Pentaho, Jaspersoft and Roambi.

SQL Server 2008, SQL Server Reporting Services, SQL Server Integration Services, PostgreSQL 9.2.4, MySQL 5.5, Ruby on Rails 3.2, Ruby 1.9.3, Linux

Part of the Business Intelligence Technology team for a mortgage-servicing subsidiary of IBM. Responsible for the management, data architecture, data integration and the design, development and implementation of the enterprise data warehouse platform to include, but not limited to the following:

- Providing technical leadership and solutions architecture for Business Intelligence - Enterprise Data Warehouse team.
- Technical lead architect of enterprise Data Governance initiative and development of next generation data warehouse platform using Data Vault™ modeling methodology.
- Architecture, development and implementation of a Master Data Management (MDM) solution to integrate data from multiple systems using a transactional queue architecture utilizing SQL Server 2008 and Informatica.
- Gathering end-user data and reporting requirements and performing in-depth analysis of source systems, data and complex business rules in order to build optimal solutions to meet the needs of our data consumers.
- Designing, developing and implementing various components of our enterprise data warehouse to include: staging, integration, operational data store (ODS) and analytics databases. Responsible for all phases of the SDLC including gathering requirements, data/system analysis, data modeling, ETL development, database development, testing, deployment, etc...

Informatica PowerCenter 8.1/9.1, Autosys, SQL Server 2005-2008, SQL Server Reporting Services (SSRS), SQL Server Integration Services (SSIS), Sybase PowerDesigner 15.3, Visio Architect, SharePoint

Technical Consultant

9/1993 - 08/2009

Provided technology services focused in Business Intelligence, Data Warehousing, Database Development and Data Integration. Major projects include, but not limited to the follow:

Nike	Miller-Coors	Roadway Transportation	Ruf Strategic Solutions
CompuCredit Corporation	Qsent, Inc.	WebMD	Electric Lightwave, Inc.
MCI/Worldcom	Bellsouth	Delta Airlines	John Harland
Astra/Merck Pharmaceuticals	Trane Corporation	Interactive Diagnostics, Inc	IT Corporation

EDUCATION

Purdue University
West Lafayette, IN

Bachelor of Science, Industrial Engineering
08/1989 - 05/1993
