

# Jordan Moore

<https://www.linkedin.com/in/jordanmoorerhit/>  
moorejordan@outlook.com | 217.622.5993

## EXPERTISE

### CLOUD

#### AWS

EC2 • S3 • VPC • IAM • Route53 •  
EMR • EKS • ASG • CloudWatch

#### GCP

GCE • AppEngine • GKE

### HADOOP ECOSYSTEM

HDFS • YARN • Spark • HUE  
Hive • PrestoDB • Oozie • Ambari

### CONFLUENT PLATFORM

Kafka • Schema Registry • Avro  
• Kafka Connect

### MICROSERVICES

Docker • Mesos+Marathon •  
Nomad • Consul

### DATABASES

MySQL • Neo4j • Cassandra • ELK  
Stack

### CI/CD

Ansible • Puppet • Terraform •  
Jenkins • Vagrant

## SKILLS

### ENVIRONMENTS

Backend • APIs • Tool Development  
• Data Analysis • Architecture

### PROGRAMMING

Java • Scala • Maven • JAX-RS •  
Dropwizard • Spring-MVC •  
Hibernate • JDBC

Python • BeautifulSoup • Flask

#### Familiar:

Android • Go • PHP • JavaScript

## EDUCATION

### ROSE-HULMAN INSTITUTE OF TECHNOLOGY

#### BSC IN SOFTWARE

#### ENGINEERING '15

Double Major in Software  
Engineering and Computer  
Science

Minor in Mathematics and  
Computational Science

Cum. GPA: 3.2 / 4.0

## EXPERIENCE

### INSPIRED INTELLECT | FORMERLY, AVALON CONSULTING, LLC

July 2015 – Present

- Working with clients to enhance and develop their overlapping Big Data, Event-driven architectures, microservices, and/or Kafka needs.
- In between projects, tasked with creating repeatable sales demos using Vagrant, Terraform, or Docker orchestration, in combination with Ansible and Bash.
- Demonstrated Data Analytical and Log Parsing abilities through the use of Elasticsearch, Logstash, and Kibana.

### CLOUD DEVOPS + RELIABILITY ENGINEER June 2019 – Present

- Using Java and Event-Driven Architecture around Apache Kafka and Cassandra to develop an AZ-Aware multi-cloud, system capable of abstracting popular Docker schedulers including Marathon, Nomad, and Kubernetes.
- Migrating Apache Mesos workloads to Nomad, and Nomad Docker services to Kubernetes

### HADOOP + KAFKA ADMINISTRATOR Mar 2017 – June 2019

- Administering and maintaining a hybrid cloud environment between a 100+ node on-premise Hadoop cluster, multiple 5-20 node Kafka clusters, and several cloud environments for a vacation rental company.
- Providing on-call support for infrastructure and developer questions, maintenance notifications, and debugging sessions.

### HADOOP EDW OFFLOADER Feb – Mar 2017

- Wrote Apache Pig scripts to convert plaintext Oracle database dumps into Parquet
- Improved runtime of an on-prem Oracle Pro\*C ETL process from 5 hours to within a 30 minute window using a combination of Apache Impala and the Parquet data format on AWS EC2.

### APACHE SPARK STREAMING DEVELOPER Sept - Nov 2016

- Built a Spark Streaming ETL pipeline for processing payment transaction documents made available on HDFS. Data became readily available to querying via Kafka, and into Hive, MemSQL, and ElasticSearch.
- Wrote custom Hadoop InputFormats with regex support to handle parsing semi-structured document files into Spark RDD Row events.

### GRAPH DATABASE ANALYST Aug - Sept 2016

- Researched differences in Neo4j and OrientDB graph databases while translating relational, tabular data into a format more suitable for a graph.
- Used both technologies to query and identify connected components of an online retailer's web traffic metadata for purposes of improving targeted advertising.

### APACHE SPARK MLLIB DEVELOPER May - Aug 2016

- Refined an iterative Random Forest algorithm to categorize textualized call-center issues reported through a help-desk system.
- Provided an ML model with over 80% accuracy against the provided dataset after manual parameter tuning.