

Faiz Abidi

fabidi89@vt.edu | 540.998.6636

EDUCATION

VIRGINIA TECH

MS IN COMPUTER SCIENCE

May 2017 | Blacksburg, VA

GPA: 3.87

JAMIA MILLIA ISLAMIA

BS IN COMPUTER ENGINEERING

May 2013 | New Delhi, India

LINKS

Website: www.faizabidi.com

LinkedIn: [in/faizabidi](https://www.linkedin.com/in/faizabidi)

Github: [faizabidi](https://github.com/faizabidi)

COURSEWORK

GRADUATE

Operating Systems

Information Storage & Retrieval

Virtual Environments

Cloud Computing

SKILLS

- C++, Python, Bash
- TCP/IP protocols, UDP, ICMP, TLS, IP packets, DNS, OSI layers, MAC addressing, firewalls, load balancing
- Data storage protocols including NFS, CIFS, iSCSI, FC, SAS
- MySQL, Hive, HBase, Pig
- Machine learning and data mining using Apache Hadoop and Spark
- AWS, GDP, Azure, OpenStack, Docker
- Linux (RedHat, Centos, Ubuntu), OSX

RESEARCH

GRADUATE RESEARCH ASSISTANT

Visionarium Lab, Virginia Tech

August 2015 – May 2017

- Built a 3D rendering pipeline using HPC servers at Virginia Tech, Turbo VNC, and VirtualGL
- Work done using C++, Python, and Bash. ParaView used for 3D data visualization

PUBLICATIONS

- Remote High Performance Visualization of Big Data for Immersive Science - SpringSim'18, Baltimore, USA
- Cloud Libraries: A Novel Application of Cloud Computing - IJ-CLOSER, Aug 2012
- Cloud computing and its effects on healthcare, robotics, and piracy - WCST'11

EXPERIENCE

FACEBOOK | PRODUCTION ENGINEER

Aug 2020 – Present | Seattle, USA

NETAPP | MEMBER OF TECHNICAL STAFF (NAS) - III

Aug 2017 – July 2020 | Sunnyvale, USA

- Worked as a member of a team to build a proprietary data migration tool in the NAS ecosystem. Speed-ups of upto 2x seen compared to rsync and similar tools
- Architect solutions to process big data and develop optimizations. Overall average job performance improved 10%-20%
- Wrote and published several technical papers showcasing custom solutions and best practices to run big data jobs on NAS instead of HDFS

NETAPP | OPENSTACK DEVELOPMENT INTERN

June 2016 – Aug 2016 | Raleigh, USA

- Scripted the procedure to virtualize Mellanox drivers to add InfiniBand support
- Got Cinder driver to run in a Centos devstack environment (internal CI system)
- Fixed a couple of bugs in the Manila shared file system project. Code merged with the upstream master branch - [LINK](#)

GRIDMARKETS PTE. LTD | PLATFORM OPERATIONS ANALYST

Sep 2013 – July 2015 | New Delhi, India

- Built a cloud rendering platform using un-used compute capacity making our solution cheaper than any other cloud service provider
- Work done using Bash and Python and Puppet for configuration management
- The stack consisted of KVM, OpenNebula, MySQL, Celery, Redis, and Splunk

SMURFERS.COM | CO-FOUNDER AND CEO

July 2014 – June 2015 | New Delhi, India

- Conceptualized and built an e-commerce company that sold merchandise inspired by popular TV shows
- Responsible for handling all the technical aspects of the company including the website, load balancing, and optimizations
- Created marketing strategies and ads to attract more traffic to the website. Also responsible for managing the company funds

PROJECTS

INCREMENTAL UPDATE FRAMEWORK TO PROCESS TWEETS

- Built an incremental update feature using Python and Bash and open source tools to transfer 1.2 billion tweets from MySQL to HDFS
- My framework improved the system performance by decreasing the server load by almost 50%

SENSOR DATA ANALYSIS USING FOG COMPUTING

- Dockerized a Raspberry Pi cluster using Docker Swarm. Optimized the cluster for energy consumption using DVFS (20% energy savings)
- Used Spark's mlib library to run Decision Trees and Logistic Regression on the cluster

APACHE SPARK VS. R | A COMPARATIVE STUDY

- Big data provided by Eastman Chemical company. Implemented 9 machine learning algorithms using Spark's mlib to run on an AWS cluster
- Predicted use-cases when R would outperform Spark w.r.t run time