

# HARSHITH UMESH

Boston, MA | +1(857) 799-0566 | umesh.h@northeastern.edu | [github.com/Harshith-umesh](https://github.com/Harshith-umesh) | [linkedin.com/in/harshith-umesh](https://linkedin.com/in/harshith-umesh) | [www.harshithumesh.com](http://www.harshithumesh.com)

## EDUCATION

---

### Northeastern University

Master of Science in Computer Science

- Relevant Coursework: Program Design Paradigm, Database Management Systems, Cloud Computing, Algorithms

Sept 2023 - May 2025

Boston, Massachusetts

### PES University

Bachelor of Technology in Computer Science and Engineering

- Relevant Coursework: Data Structures, Computer Networks, Information Security, Machine Learning, Web Development

Aug 2017 - May 2021

Bangalore, India

## SKILLS

---

**Languages:** Python, Java, Golang, Bash, C, SQL, HTML, CSS

**Cloud Technologies:** Docker, Kubernetes, Openshift, Podman, AWS, RabbitMQ, ArgoCD, Ceph, Ansible

**Data & Visualization:** Grafana, Kibana, Pandas, Matplotlib, Prometheus

**Tools:** Linux, Git, Apache airflow, Jenkins, Flask, Junit, pytest, Postman, Selenium, Wireshark, Nmap

**Databases:** MySQL, PostgreSQL, MongoDB, Elasticsearch, SQLite

## WORK EXPERIENCE

---

### Teaching Assistant, Khoury College of Computer Sciences

Jan 2024 - May 2024

- Assist in the instruction of course materials for *Networks and Distributed Systems* to a class of **160 students**.
- Provide guidance and support to students, including grading and offering feedback on assignments and projects.

### Red Hat

Software Engineer

Performance and Scalability - Openshift - R&D

Bangalore, India

Jan 2021 - Aug 2023

- Involved in the **Research and Development** of tools to improve the performance of Red Hat OpenShift.
- Contributed significantly to the development of two inhouse **CNCF** sandbox projects “Kube-burner” and “Arcaflow”.
- Revamped the OpenShift Performance CI by migrating from a legacy system to a **GitOps model** powered by Apache Airflow.
- Architected and implemented an ETL pipeline to move data and metrics stored on elasticsearch to an RDBMS, which is consumed for **data science, machine learning and data visualization** applications.
- Authored and maintained scripts written in python and bash, which use boto3 to identify and delete idle or zombie OpenShift clusters and all associated resources created on AWS, helping **reduce operational costs by 45%**.
- **Lead maintainer** of the tooling and infrastructure used by the team, deployed argoCD using helm to monitor the health of all services used by the team. Integrated it with slack to notify of failures and downtimes.
- Conducted **analysis and benchmarking** of **OCP and Kubernetes** control and data plane components across **10** minor releases.
- Developed an end-to-end CI/CD pipeline using Jenkins to streamline the OpenShift installer development process; any major code change triggers an installation of an OpenShift cluster on **AWS, Azure or GCP** and returns a **pass or fail** report.
- Developed and deployed an internal object storage file server “Snappy” using podman containers to store internal data as an alternative to AWS S3 **reducing cloud spending costs by 20%**.
- Designed and deployed **Analytics dashboards** using Grafana and Kibana to help visualize the performance of benchmarking tools such as FIO, Uperf, HammerDB, Sysbench and Prometheus on OpenShift.

## PROJECTS

---

### DBaaS for Rideshare - [Github link](#)

- Built a highly available database as a service(DBaaS) for a carpooling ‘RideShare’ application using REST APIs and python-flask.
- SQLite is the database used. The application is run inside docker containers hosted on AWS EC2 instances. AWS load balancer and zookeeper were used to maintain high availability.
- Docker SDK and RabbitMQ were used for scaling up/down containers, and Postman for testing the API endpoints.

### Setting Up and Exploration of Security in a Hybrid Cloud - [Research Publication link](#)

- Built a hybrid cloud infrastructure for the document storage of a university. AWS S3 is used as the public cloud storage, and Ceph storage deployed on bare-metal servers as the private cloud storage. Nginx load balancer connects the two clouds.
- Faculty, students and outsiders each have different access control permissions(RBAC). Designed and automated a script to perform penetration tests, identifying security vulnerabilities in the hybrid cloud.

### Four Layer Authentication with Honeypot and Cloud Data Encryption - [Research Publication link](#)

- Created a secure private cloud interface using Django, featuring a multi-layered authentication process including password, OTP, ReCaptcha, and RSA certificate.
- Multiple authentication failures redirect to a honeypot, a dummy file store where the intruders’ IP address is captured. Successful authentication redirects users to a file store where files are AES 256-encrypted.

### Image Processing Java Application - [Private university github link](#)

- Developed an image processing application in Java, with similar functionalities as Photoshop. The application was built using the Model-View-Controller (MVC) design pattern.
- The application features both a CLI and GUI. Java Swing was employed to create the GUI. Comprehensive testing was conducted using the JUnit framework.