

SUMMARY

I am a backend software engineer with 10+ years experience, currently working at IBM. I work on kubernetes operators and developer platform and tools.

EXPERIENCE

IBM

San Jose, CA, USA

Software Engineer

Oct 2015 - Present

CloudPak Engineering:

- * Developed Kubernetes controller for Redis, which is used as shared dependency across IBM CloudPaks.
- * Developed general-purpose service broker operator for kubernetes (using OSB API).
- * Developed solution and tooling for disconnected installs of IBM CloudPak on Openshift clusters.
- * Developed internal developer CLI (go) for packaging and delivering IBM CloudPaks.
- * Developed certification framework for delivery of IBM CloudPaks on RedHat OpenShift (RHOCP).
- * Modernized existing CI/CD pipeline using gitops principles, to reduce complexity and build time.
- * Implemented code & image signing for CI/CD pipeline, to ensure integrity and avoid supply chain attacks.
- * Developed certification framework containing security best practices and patterns that IBM products must follow when developing production grade applications on Kubernetes.
- * Maintained inner source components, to promote resusability across IBM products.
- * Participated and contributed in external open source communities like operator-sdk, operator-lifecycle-manager.

Watson Natural Language Understanding:

- * An API as a service platform for natural language understanding tasks.
- * Integrated sentiment service with existing stack (Typescript).
- * Developed helm charts to deploy the product in IBM Private cloud.
- * Automated migration of standalone databases to IBM Cloud Database instances.

Analytics Engine:

- * A IaaS compute platform running Apache Hadoop and Spark.
- * Developed CLI(Go) tool to interact with the cluster, launch spark jobs and retrieve logs.
- * Implemented the webHDFS REST api as file system commands in the CLI.

BigInsights on Cloud:

- * A big data platform running Apache Hadoop on VM's and baremetal machines.
- * Developed Chef recipes to scale clusters by adding nodes, backup and restore.
- * Implemented parallel delivery of security fixes to clusters.
- * Implemented GDPR complaince, by adding Vault support and encrypting disks.
- * Encryption was challenging task due to large size of disks (4 TB X 8 disk X 5-8 nodes). It's typically done by backing up data to temp disk storage and then encrypting it, but instead I proposed to use Hadoops self-healing and rack-awareness to handle data loss during encryption which cut down time from weeks to days.

Spark as a Service:

- * A multi-tenant platform for running notebooks and batch jobs with Apache Spark, running on top IBM Spectrum and Apache Mesos.
- * Created a containerized integration test framework using Cucumber for the CI/CD pipeline.
- * Improved the platform security by fixing vulnerabilities identified from external penetration testing.

Analytics NextGen Workbench:

- * A platform for data scientists to design and develop predictive models and execute with SPSS backend.
- * Developed a scheduler microservice using Akka/Scala for the platform.

University of South Florida*Research and Teaching Assistant*

Tampa, FL, US

*Aug 2014 - May 2015***Research Assistant:**

- * Analyzed customer software subscriptions data provided by Wharton customer analytics initiative (wcai) research group.
- * Involved in cleaning and preparation of data set, feature extraction, data visualization and identifying research questions.

Teaching Assistant:

- * Distributed Information Systems (ISM 6225) and Information Security and Risk Management (ISM 6328)

Software Developer:

- * Developed responsive web pages for college of global sustainability using Bootstrap, HTML5, CSS3 and JavaScript.

Cognizant*Software Engineer*

Chennai, TN, IN

*Aug 2008 - Nov 2013***Performance Engineering:**

- * Responsible for profiling JVM and heap dump analysis, to identify potential memory leaks and slow running code.
- * Tuned JVMs and recommended GC policies appropriate for the application, which improved the scalability and reduced memory footprint.
- * Developed RESTful web services, providing in-house performance engineering tools as SAAS services for internal development teams.

EDUCATION

University of South Florida*Master of Science in Management Information Systems; GPA: 3.94/4.0*

Tampa, FL, US

*Jan. 2014 - May. 2015***Madras Institute of Technology***Bachelor of Engineering in Computer Science; GPA: 6.8/10.0*

Chennai, TN, IN

*Aug. 2004 - July. 2008***PROGRAMMING SKILLS**

Languages: Go, Python, TypeScript, Scala**Cloud Platforms:** Kubernetes, RedHat OpenShift, IBM Cloud, AWS**Infrastructure:** Chef, Ansible**CI/CD:** ArgoCD, Travis, Tekton**Databases:** RDBMS (PostgreSQL, MySQL), NoSQL (Redis, MongoDB)**ML Frameworks:** R, Tensorflow**Others:** Containers, Git