Deepu Mohan Puthrote

Senior Software Engineer
London, United Kingdom
Email: mail@deepumohan.com
Phone: +44 7400025813
Website: https://deepumohan.com/
GitHub: https://github.com/WarFox
LinkedIn: https://linkedin.com/in/deepumohanp

Summary

A generalist software engineer, building software since 2008, across backend and event-driven systems, data platforms, infrastructure and frontend. Research-oriented continuous learner who is always looking for ways to expand my knowledge and skills in computer science. Proven track record in leading teams to solve complex technical challenges, implementing DevOps and CI/CD practices, and reducing cloud infrastructure costs.

Passionate about continuous learning and innovation, aiming to leverage expertise to drive impactful solutions in software engineering. I have consistently demonstrated my ability to propose ideas and validate their feasibility through examples. I favour solid design with fast iteration over quick-and-dirty solutions that impede future modifications. I can work independently and as part of a team and am always willing to go the extra mile to get the job done.

I enjoy collaboratively building end-to-end solutions and can do functional programming, object-oriented programming, system design, and architecture. I am driven by a desire to solve complex problems and innovate continuously. Continuous improvement is better than delayed perfection.

Experience

CAIS, London, UK, Sep 2024 – Present

VP | Senior Software Engineer, Sep 2024 – Present

- Saved ~1000USD/month on S3 by identifying and deleting millions of unnecessary objects
- Designed and implemented feature to prevent duplicates in historical records spanning multiple services
- Reduced time for adding new external integration from months to 1 week by automating the process
- Reduced time taken for configuring various data feeds by building a self service tool for the ops team. Built backend APIs in Spring/Kotlin and frontend in React
- Optimised code base by deleting ~10,000 lines of code, but still delivering the same business requirements with zero downtime. This improved the maintenance time of the system.
- Improved developer workflow by identifying and automating a task that took 45mins to 1hour per developer
- Designed and implemented prevention of historic duplicates in outgoing data files
- Implemented a paginated search api over postgres and jOOQ
- Identified and fixed several bugs and security vulnerabilities

Technologies: Kotlin, Spring Boot, Test Containers, Kubernetes, AWS, Airflow, jOOQ, React

Funding Circle, London, UK, Jul 2018 – Jul 2024

Gave several suggestions on technology choices and improvements throughout my career in tech review meetings. Taught Python to both internal and external candidates as part of DiversiTech Initiative in FundingCircle in addition to mentoring junior engineers.

Technologies: Clojure, Python, Scala, Terraform, Avro, Kafka, Mesos, Chronos, Docker, Jenkins, Circle CI, PostgreSQL, Athena, HugSQL, AWS Glue, Kafka S3 sinker

Principal Engineer, Jan 2021 – Jul 2024

Building a self-service data platform for batch analytics, machine learning development, and near-realtime processing, improving data accessibility and processing efficiency. Led technical initiatives, collaborating with the VP of Data, Director of Data, Head of Data Engineering, Head of Data Governance, Head of ML/AI, Head of Self Service Analytics, and Senior Manager of Data Platform, along with internal partners like Product, InfoSec, and SRE, to ensure alignment and technical excellence

Various POCs, RFCs, standardisations, data reliability engineering and support the data platform engineering teams. Working an next generation machine learning piplines using Sagemaker. Data access protection using Lakeformation.

RFC for AWS account segregation within organisation for data, POC for metadata platform using DataHub, metadata tagging, migration tools and projects etc. was delivered while being Principal Engineer at Funding Circle.

Wrote down Data Platfrom Principles that makes it clear for stakeholders and teams to meet the expectations. Wrote Standards for Monitoring and Alerting.

Senior Engineer II, Jan 2020 – Jan 2021

Introduced AWS Fargate as an option to run tasks instead of Mesos within Data Platform. Brought down cost of AWS Glue and S3 by redesiging pipeline with a compaction step. Introduced the idea of "decentralised" projects to isolate related resources in their on repositories. Lots of interviews.

Senior Engineer, Jul 2018 – Jan 2020

Funding Circle is a leading small business loans platform, where investors lend to small businesses.

Built software and AWS infrastructure for the data platform: ETL tooling on Kafka, Athena and Glue, and streaming, batch and event-driven jobs.

- Introduced the concept of "inversion of control" in the way data related products were setup at Funding Circle, that fundamentally changed how we built tools. This concept allowed us to deliver solutions faster for a wide range of data products and data science projects.
- Designed and developed data validation tool using Clojure that can validate data on multiple datasources like Postgres and AWS Athena
- Proposed a better way for deploying Airflow DAG within our infrastructure which my team implemented later.
- Built ETL's based on Kafka streaming, Athena SQL and Glue
- Streamlined Glue and Lambda deployments using automation tools.
- Proposed solution for a generic data loader for documents in multiple formats like JSON, XML and CSV, that supports schema validation and evolution.

Valtech, London, UK, Jun 2017 – Jul 2018

I joined Valtech as a Senior Software Engineer. At Valtech I got training on presentation skills, which I used to present a talk on Anomaly Detection in time series data.

Website customer behaviour prediction, Software Engineer/DevOps, London, UK, Jan 2018 – Jul 2018

- Research on AWS Sagemaker
- Prepare JSON data from Bson using PyMongo library
- Design and implement data pipeline using EMR
- Design and implement machine learning pipeline using EMR
- Streaming data with Kinesis
- Streaming KMeans Clustering algorithm using Apache Spark Streaming context
- Persistance in S3 and DynamoDB
- Visualising streaming clusters with LightningServer and Dash
- Infrastructure automation with Terraform and Ansible

Technologies: AWS, Spark, Hive, Python, Jupyter, EMR, Terraform, Zeppelin, AWS Glue, AWS Athena, S3, DynamoDB

Big-data pipeline and analytics, Software Engineer/DevOps, Norway, London, Oct 2017 – Dec 2017

Data project to analyse large volumes of data on Hadoop echo system on Microsoft Azure platform. This project has a data pipeline part, some data scrapping and data analysis part. Data came from various motors and sensors.

- Get vessel and weather information from web using Microdata and API's in Python
- Data pipeline built on Azure HDInsights platform
- Apply transformations to Hive tables using PySpark based on algorithms by data scientists
- Created a Workload report of HDInsights using Azure Runbooks, using Yarn and Livy APIs
- Build a data analytics platform on Azure using HDInsights, Spark and Jupyter
- Setup Ansible playbooks for configuring JupyterHub and users
- Explored Apache Ranger using Hortonworks Sandbox on Ubuntu

Technologies: Jupyter, Apache Spark, HDInsights, Azure, Ansible, Plotly, Apache Livy, Horton-works, Python, Azure Storage accounts

HMRC - Senior Scala Consultant, Senior Scala Consultant, London, UK, Client: CSHR (https://www.gov.uk/government/organisations/civil-service-fast-stream), Jun 2017 – Oct 2017

Project Type: Website, Customer responsive platforms

This is a UK government project delivered following the standards set by Government Digital Services. The UK government uses Fast Track and Fast Stream to recruit talents for the Civil Service.

- Re-design the two services, with new GOV.UK services, so that candidates can apply and attend online tests
- Back office system for candidate support
- This is a responsive website service that is easy to use on any device.

This project involved creating API, candidate frontend and admin frontend for UK Government Civil Service HR Faststream candidate application.

Part of this project is open source and is available at hmrc/fset-faststream (https://github.com/hmrc/fset-faststream), hmrc/fset-faststream-frontend (https://github.com/hmrc/fset-faststream-frontend)

Technologies: Scala (Play framework), HTML5, CSS3, JS, Progressive enhancement, Respon-sive web design

Expedia Affiliate Network, London, UK, Aug 2015 – Jan 2017

Joined Expedia on contract basis for building the next generation reporting tool for internal users and then extended the contract on to anomaly detection platform. Responsibilities included software development in various languages, design & architecture discussions, DevOps, deployment pipeline, and interviewing candidates.

Reporting Platform

The project involved building a reporting platform using micro-services architecture in a Kanban environment. RESTful API's were built using Node.js and front end was built with React.

- Developed user identity management service API and implemented authorisation algorithm using JSON Web Tokens(JWT) and oAuth.
- Designed and implement various REST API endpoints in Node.js and Restify for getting the data. HATEOAS (Hypermedia As The Engine Of Application State) pattern was used for linking between resources.
- Wrote a framework Chef Cookbook (Library cookbook) for deploying Node.js based services.
- Wrote Chef cookbooks (Application Cookbook) for various services, including rspec tests and integration tests with kitchen vagrant. Acted as pseudo DevOps member within the team.
- Implemented API Gateway using nginx and Chef.
- Implemented Service Discovery using Hashicorp Consul. This radically improved the time required for the service to be available in the gateway after deployment.
- Implemented next iteration of API Gateway with nginx, Consul and chef.
- Deployment in AWS Cloud. Setup nodes using Scalr. Fix various deployment issues.
- Did a brown bag session on Hashicorp Consul and service discovery, which was well-received.
- Worked closely with Operations team to deploy services to Production environment.

Technologies: Node.js, Chef, cookbooks, rspec, serverspec, chefspec, AWS Cloud, Teradata, Node.js, Ruby, React, redux, rubocop, Jenkinss, Scalr, bamboo, stash, gitlab, DevOps, restify, Bamboo, upstart

Anomaly Detection platform (Machine Learning)

Big data analytics project for detecting anomalies in various time series data using machine learning techniques. Worked closely with Data scientists to develop models and the platform.

- Worked closely with Data scientists to create models in Scala to run on Apache Spark. Original models were written in Python and R.
- Used Spark to parallelise forecasting based on splits of data. Wrote Scala based tool for triggering AWS Simple Notification Service (SNS).
- Automated Airflow DAG creation using Python based on configuration files in AWS S3.
- Wrote AWS Lambda functions in Python to trigger data pipeline.
- Setup Elastic Search cluster on AWS cloud. Wrote a script to automate the process.
- Optimised data ingestion to Elasticsearch from Hive on AWS Elastic Map Reduce (EMR), brought down the time taken from 3.5 hours to 20 minutes.
- Used Spark to move data from Teradata to S3 in Parquet format.

Technologies: Apache Spark, Hive, AWS Cloud, AWS Lambdas, Python, AWS SNS, AWS Elastic Map Reduce, Scala, Python, Airflow, Bamboo

Kainos, Burnley, UK, Feb 2015 – Jun 2015

In DWP location on a project related to Government Digital Service (GDS) and Gov.UK Verify using micro-services architecture as a Tech Lead in an Agile (Scrum) environment.

Actively upgraded existing source code to use Java 8 features. Helped new team members practice TDD. This project uses Elastic search to find possible matches from a matching data set. Wrote scripts to ease development process. Brought in consistency and standardisation to source code.

Responsibilities include technical discussions with architects, sprint planning, sprint elaboration, delivering sprint vision, knowledge sharing among team members, pair programming and code review.

Reported and fixed issue #953 (https://github.com/dropwizard/dropwizard/issues/953) in DropWizard project on GitHub during the stint. dropwizard/dropwizard#1001 (https://github.com/dropwizard/dropwizard/pull/1001)

Technologies: Java 8, DropWizard, JDBI, PostgreSQL, Node.js, Ruby, Git, Jenkins, JIRA, Nginx, Gitlab, JavaScript, Mongodb, Python for scripting, Elastic search, Docker, Kibana, Logstash, JUnit and Fest Assertions

Corelogic, Essex Road, LONDON, UK, Nov 2014 – Jan 2015

The project was to build new module into existing social care product developed by Corelogic. Worked on Java (Spring/ibatis), Backbone/Marionette, Microsoft SQL Server and Oracle stack. Suggested positive changes to standards and practices.

Introduced JSR-303 server side validation. Designed RESTful API's and structured controllers. Responsibilities included attending sprint meetings, code review and merging changes from other developers.

Technologies: Java, Spring, ibatis, Backbone.js, Marionette, Handlebars, Underscore.js, MSSQL, Oracle, Weblogic server, Gitlab, JIRA, TestNG, Mockito

Downtown Consulting, Harrow, UK, Oct 2014 – Present

Available to hire for contract jobs immediately. Happy to relocate to anywhere in UK.

Collective: The Audience Engine, Bangalore, India, Mar 2013 – Sep 2014

Collective is a leader in multi-screen advertising and dynamic optimized creatives and programmatic advertising that connects brands to audience with personalized ads. Joined the team at Collective to build and evolve the best dynamic creative optimized ad platform in an Agile environment.

Software Engineer

- In-house BigData Analytics product, which produces reports from the ad delivery/engagement logs generated by the ad server.
- Designed and developed secure REST API for reporting dashboard using Play Framework and Redis for caching.
- Play framework and NodeJS/ExpressJS based secure RESTful APIs
- Rich UI using AngularJS and RequireJS

RESTful Reporting API (Dashboard)

- Setup and lead Play (Scala) project from scratch (Proposed the use of this framework).
- Slick (Scala) for database access.
- Implemented a state machine based Authentication header parser and developed algorithm.
- Single sign-on integration with Authentication API
- Redis for caching
- Jackson for JSON generation.
- JAXB annotations for converting authentication message to object.
- Spring Batch for batch processing and pre populating the Cache.
- Vagrant boxes for database and caching during development.
- Pair Programming
- Test Driven Development
- In-memory H2 database with YAML data files for unit testing.
- Knowledge transfer to team members.
- Build with Jenkins, deployment with Chef

Reporting UI (Dashboard)

- Setup and lead Node/Express project from scratch.
- Visualization using HighCharts and D3JS
- Build with Jenkins, deployment with Chef

Proof Of Concept

- POC for the very ambitious dashboard project
- Created using Play (Java/Scala)
- UI with Angular JS
- Data visualization with D3JS and NVD3

Tableau Reporting

- Backbone.js and Bootstrap based app for Tableau views built on top of Tableau JavaScript API
- Yeoman based project setup.
- UAT in Heroku
- Build with Grunt and Jenkins, deployment with Chef

In-house Analytics product
- Maintained this already matured product.
- Modernised UI using Bootstrap, Struts, jspx and tagx
- Added new reports
- Scheduling with Quartz
- IBatis for SQL queries
- Python scripts for chores.
- Build and deployment with in-house tool

University of Newcastle upon Tyne, UK, Nov 2012 – Jan 2013
This was a paid project given by the University as part of Newcastle Work Experience program to build an academic assignment feedback generation application which, makes it possible to prepare constructive and comprehensive feedback in an efficient and effective manner.
- Built using rapid application development tool Spring Roo, Jasper Reports and Dojo
- User authentication and authorisation features using Spring Security.
- Java Mail API to send registration email.
- UAT deployment in Heroku and PostgreSQL
- Production deployment in Tomcat server and MySQL database located at remote University Server using ssh key authentication.

Happiest, UK, Mar 2012 – Sep 2012
Happiest is a start-up in Newcastle upon Tyne, that focused on Rewarding customers based on engagements. To be a part of an exciting start-up in the UK was really thrilling. Most exciting part was the thought of building something that is going to delight people. The internship turned into an offer for doing Master's dissertation.

MSc Dissertation Researcher, Jun 2012 – Sep 2012
- Dissertation project on Social Network Analysis and Centrality metrics
- This project was considered the initial phase of bigger data analysis project.
- Created a friend recommendation system using Neo4J graph database.
- Analysis of social network data using R, IGraph
- Happiest's network data was analysed and centrality points were given to each node in the network.
- Graph data visualization using R and Gephi.

API Developer, Mar 2012 – Jun 2012
- Developed various API end points in PHP
- Database programming in PostgreSQL
- Responsible for GitHub commits and merges.
- Scrum methodology
- Participate in daily stand-up sessions and discussions.

JMR Infotech, India, Aug 2008 – Aug 2011
JMRInfotech is a technology consulting company Banking and Financial Services Industry domain in India, where I served for 3 years including offshore support and client site development projects.

Software Engineer/Technical Consultant
- Was deployed in client site and acted as single point of contact.
- Lead the team in design and development and also gave training to team members and knowledge transfer to new members.
- Designed and developed Java based web framework for development.
- Training in Flexcube (OFSS), PL/SQL and Oracle Forms.
- Offshore support for Flexcube.
- Responsible for code maintenance and version control.
- Responsible for UAT deployment.

User and session management module
Created the complete design document for the project.

Customised account reports from Flexcube database
Lead developer in Java, JQuery, Hibernate, Jasper Reports, production deployment in Websphere server

Bank cheque truncation automation project
Lead developer in Java Swing interface between Flexcube and cheque truncation automation vendor. The Swing application generates XML files as specified by the client in a given location.

Credit card tab for Internet Banking
Developed in Java, production deployment in Websphere server. The credit card tab worked by consuming Webservice (SOAP) provided by Credit card vendor.

Bulk Signature Uploader
Upload scanned signature images from a given location. Single-handedly developed. Used XML for representing data and XSLT to translate to html.

Data double entry
Web application to verify same data entered by two employees. Single-handedly developed. Used XML for representing data and XSLT to translate to html.

Company Website
Built on DotNetNuke. Team member.

Search Engine Optimization
Webmaster

Education

M.Sc. E-Business Information Systems (Merit), University of Newcastle upon Tyne (http://ncl.ac.uk/), UK, Sep 2011 – Sep 2012
Taught postgraduate course with modules from both computing school and business school, which gave great exposure and outlook on both technology and business aspects. Skills for project management and business planning were enhanced by the consultancy project and business plan group work.

- Dissertation on Social Network Analysis and centrality metrics.
- Prepared a complete Business Plan and prototype as part of group work.
- Individual assignment and research on BigData and its uses in EBusiness, based on Gartner's hype cycle.
- Developed a simple web server in Java.
- Modified a ticket booking application using Java, Hibernate and Memcache in IntelliJ.
- Developed an Android and PHP application for Newcastle University Alumni Association, which uses Facebook SDK and Twitter API.
- Class representative - Participated in meetings and voiced issues faced by classmates.
- Actively participated in events happening in the school like Start-up Weekend (http://newcastle.startupweekend.org/), Thinking Digital Meeting.

Modules: CSC8004 Network and Web Technologies, CSC3103 Internet Technology and Electronic Commerce, NBS8207 E-Business

B.Tech Computer Science & Engineering (Merit), AWH Engineering College (http://awhengg.org), University of Calicut (http://universityofcalicut.info), India, Jun 2004 – Jun 2008
Bachelors degree course that covered in-depth concepts of computer science, mathematics, algorithms and software development. Programming lab in C, Java and Lisp was included in the course to enrich the understanding of subject.

- IEEE and CSI (Computer Society of India) associations
- Secretary of Computer Science Society
- Organised events for Computing Society
- Won programming competitions in C
- Mini project on face detection and pattern recognition
- Main project was a robot that rolled on wheels with voice recognition and response, face detection, obstacle detection and path learning capabilities.

Skills

Programming Languages: Clojure, Java, Kotlin, Scala, Python, Typescript/JavaScript
Tools & Technologies: AWS, Terraform, AWS CDK, Spark, Kafka, Kubernetes, Docker
Methodologies: DevOps, CI/CD, Agile, TDD, DDD
Disciplines: Backend Engineering, Full Stack Engineering, Data Platform, System Design, Cost Optimisation, Technical Leadership, Problem Solving

Technical Skills
Languages: Java, PL/SQL, JavaScript, shell scripting, Python scripting, Scala
Web: HTML, CSS, JavaScript, JQuery, AngularJS, Backbone.js, RequireJS, AJAX, JSON, Web Service, Bootstrap
APIs: JEE (JSP, Servlets, Filters), JPA, JAXB, JUnit, Jasper Reports, Log4J

Frameworks: Play Framework, Spring Core, Spring MVC, Spring Roo, Spring Security, Hibernate, Spring Data Neo4J, Spring JUnit
Methodologies: Agile (Scrum), OOP, Design Patterns
Middleware: Apache Tomcat, Websphere, nginx
Cloud Platform: Amazon Web Services (AWS), Microsoft Azure
Automation: Chef, Ansible, Terraform
Service Discovery: Consul
Big Data: Apache Spark, Hive, EMR
Database: Oracle, PostgreSQL, MySQL, Neo4J (NoSQL graph database)
Tools: IntelliJ IDEA, Spring Tool Suite, Eclipse, Toad, PL/SQL Developer, pgAdmin, Sequel Pro
Source control: git, svn, perforce
Cache: Redis, EhCache
Acquainted with: C, C++, R, PHP, Ruby
Unofficial: MongoDB, Travis CI

Personal Projects

Automate Email Reply
- Use Gmail API to automate reply of emails from recruiters, with the latest copy of my CV as attachment
- Written in Python
- Features include greeting with First name extraction, apology note incase I'm replying the next day, dynamic greeting based on time of day (in-progress)

Strong Password Generator, https://deepumohan.com/projects/strong-password-generator/
- Small Reactjs application, with Material Design
- CI/CD with bitbucket and Wercker

Bangalore Hangouts, http://bangalorehangouts.com/
- Ruby on Rails app that lists cool places in Bangalore to hangout.
- RESTful API based.
- People can rate and comment on places.
- Places can put up events and other marketing stuff in their pages.
- Work in progress

501 Expenses
- Webapp to store and split expenses between group of people
- Using Play framework
- REST API based
- Work in progress

Score App
- Small JavaScript based web app to enter scores while playing games.
- Using AngularJS.
- Stores score in browser
- Next step: Use Cordova to port this app into mobile.

Interests
- Movies
- Wall Climbing
- Travelling
- Sky Diving
- Blogging - I write about technology in my blog at deepumohan.com/tech (https://deepumohan.com/tech/)