

Mayank Raj

SaaS Founder
& Associate Director of
Engineering

aws
certified
Solutions
Architect
Associate

aws
certified
Solutions
Architect
Professional

I'm a solutions architect with extensive experience in not only building high performant products powered by AI/ML & Big Data but also designing resilient and fault-tolerant cloud infrastructure for it. I have architected products that can scale to handle 100x load within minutes and yet remain cost-effective. I have also built a graph database engine to work with more than billions of nodes and edges.

JavaScript | Python | AI | NLP | BigData | AWS | Graph Database | System Architecture | Distributed Systems
Serverless Systems | CI/CD | Cloud Automation | Tech Author | Automated Drones

JavaScript, NodeJS, Python

AWS, GCP, EC2, ECS, S3, Lambda,
RedShift, Kinesis, API Gateway,
CloudFormation etc

GrayLog, Sentry, CloudWatch, X-Ray

RDS, Postgresql, RedShift,
DynamoDB

Neo4j, AWS Neptune, DGraph

Cost effective, Secure and Reliable
Infrastructure Design, Serverless
Architecture

Reproducible Infrastructure -
CloudFormation, SAM, CDK
Docker, Git

AI, ML, NLP, PyTorch, TensorFlow,
Pandas, Keras, TensorFlow.js, ML5.js,
WebVR, Ar.js, Three.js

Automated Drones, Ardupilot, PX4,
Custom Keyboard, IoT
Core Contributor to Firefox

Authoring book titled "Hands on AI
Projects" with O'Reilly

Cactus Communications

Associate Director of Engineering (Principal Solutions Architect) July '20 - Present

Cactus Communications is a technology company accelerating scientific advancement. We offer AI-powered scholarly publishing solutions for journals and researchers.

With offices in London, Princeton, Singapore, Beijing, Shanghai, Tokyo, Seoul, Aarhus, and Mumbai; a global workforce of over 3,000 experts; and customers from over 190 countries; Cactus Communications is a truly global brand.

- Bootstrapping multiple products in the domain of AI/ML and BigData
- Architected BigData platform that **ingests over 1.5Tb/week** and **generates over 4.5Tb/week**. It **manages over ~250Tb** of data in the datalake.
- Designed and developed a pipeline to process data for a Machine Learning use case with **cumulative 24K vCPU, 48Tb Memory**. Generated over ~4.5Tb of data in under 2.5 hours at **1/5th the proposed cost from AWS Big Data Team**.
- Architected a truly serverless highly available API that orchestrates **~100K of long running requests with sub-second SLA**
- Gap the bridge between the business and tech side; bootstrap projects from the idea stage to enterprise scale, oversee hiring across verticals
- Lead a team of 15+ across India and Denmark
- **Lead a team of Solutions Architects** who consult other teams on their workloads

Apptale.io

Dec '20 - Present

Sideproject grown out of personal need, now used to monitor services by many

- Designed architecture to offered as a SaaS at **1/15th the cost of competitors**
- Built to have monitoring **service orchestrated across all 51 AWS availability zones** for a truly highly available monitoring system
- API layer built with support for GraphQL & REST, global database via DynamoDB and payment & subscriptions via Stripe
- **Built with Amplify, CloudFormation, Appsync, API Gateway, DynamoDB, Lambda** and more

✉ hello@mayankraj.com

☎ +91 83693 39460

Mayank Raj

SaaS Founder
& Associate Director of
Engineering

aws
certified
Solutions
Architect
Associate

aws
certified
Solutions
Architect
Professional

JavaScript, NodeJS, Python

AWS, GCP, EC2, ECS, S3, Lambda,
RedShift, Kinesis, API Gateway,
CloudFormation etc

GrayLog, Sentry, CloudWatch, X-Ray

RDS, Postgresql, RedShift,
DynamoDB

Neo4j, AWS Neptune, DGraph

Cost effective, Secure and Reliable
Infrastructure Design, Serverless
Architecture

Reproducible Infrastructure -
CloudFormation, SAM, CDK
Docker, Git

AI, ML, NLP, PyTorch, TensorFlow,
Pandas, Keras, TensorFlow.js, ML5.js,
WebVR, Ar.js, Three.js

Automated Drones, Ardupilot, PX4,
Custom Keyboard, IoT
Core Contributor to Firefox

Authoring book titled "Hands on AI
Projects" with O'Reilly

Cactus Communications

Solutions Architect

May '18 - Jun '20

- Report to the CTO
- Designed and **developed a graph database** engine that **queries ~8Tb of raw data in under \$30/query** with avg query time of 8 minutes.
 - Benchmarked to **outperform Neo4j and AWS Neptune**, in cost and speed
 - Collaborated with the AWS Neptune team on benchmarking
- Designed and maintained the infrastructure of projects that achieved state-of-the-art scores in Natural language Detection & Correction, NLP and Vision
- Setup the best practices and implementation guidelines for the team to operate on Cloud across AWS, Azure and GCP. Templatised infrastructure
- Managed team that grew from 1 to 16. Overall grew from 1 to >60
- Reported and **worked with AWS S3 team** on fixing bug on prefix throughput
- Developed autoscaled microservices spread across Windows and Linux.

Senior Software Engineer

Jul '16 - Apr '18

- Built clickstream application with AWS Kinesis and Redshift. Designed to **ingest ~5M records/day**, operating at **under \$9/day**.
- Developed daemons to **monitor** activities on **45+ web properties and 12+ servers**
- Developed Data Lake with Hot and Cold Storage layer to optimize for cost and data availability. Manages over **4Tb+ data, powering 120+ Dashboard** for BI.
- Designed and built a service that connects any user on the same page, in real-time using web sockets. Service **scaled from 47 users to 2.7K in 36 minutes**.

Directi

Software Developer - Consultant

May '16 - Jul '16

- Designed and developed a full stack application to handle recruitments and track the candidate from the time they enter to the time they are hired
- Also handle walk-in candidates by generating their profiles on site
- Generate automated reports and email trigger

WorkAmp, AIESEC, Frapp.in, TCS, HamaraEvent...

Software Developer - Consultant

Feb '14 - Apr '16

Education

Bachelor in Technology, Computer - NMIMS University, Mumbai, India

✉ hello@mayankraj.com

☎ +91 83693 39460