Sraddhanand C

SRE, Full Stack Developer, Automation enthusiast, Cloud & DevOps Architect (Exp: 10+).

Bengaluru,Karnataka
(+91) 805-608-4026
sradhu.anand@gmail.com
https://www.linkedin.com/in/sraddhanand-c/

EXPERIENCE

Pega Systems, Bengaluru, IN — Principal System Engineer - Cloud

October 2018 - PRESENT

Roles & Responsibilities:

- Define and implement SLI/SLO/SLA.
- Design, Develop, deploy and maintain the high available, scalable and reliable microservice architecture
- Identifying the pain points in the operations and automating the solutions to **reduce the toil.**
- Develop automations to decrease the Mean Time To Detect and Mean Time to Recovery.
- Design and implement the **Event Driven Auto Remediation** for critical alerts triggered by **datadog** and Pega APM.
- Design and implement observability practices using open source softwares such as **prometheus, grafana, Jaeger, Elasticsearch - Beats** for microservice architecture.
- Design solutions to measure **uptime, availability, cost-optimization, capacity monitoring.**
- **Proactive monitoring** for a large fleet of cloud resources and reporting them and designing the solutions to fix the addressed issues.

Automations & *bilities:

Forensic Analysis: To reduce **Mean Time to Detect** the system outage, a python package is created to scan a fleet of instance's cloudwatch logs for the given time window and display html output consisting of exceptions, timestamp, number of occurrences per instance along with the whole summary of cluster for the affected environment.

Stack Health Validation: A consolidated **status page** of the particular client environment is designed using AWS SSM run command, python script to gather the indicators and health of the customer environment consist of multiple tiers of instances.

Database Health: SQL statements based on SOP along with postgresql's psql utility to get the health of postgres engine running in AWS RDS. It lists out

CERTIFICATIONS

AWS - Certified Solution Architect - Associate

PRIMARY SKILLS

AWS, Kubernetes, Docker, Prometheus, Terraform, Mongodb, Ruby on Rails, FastAPI, Linux, Python, AWS-SDKs, DevOps, Grafana, CICD.

SECONDARY SKILLS

React.js, Jenkins, NGINX, postgreSQL, Shell, Ruby, Chef, Ansible, git, Redis, Elastic Stack, Datadog

SUMMARY

I'm SRE as defined by Google who can do operations work and automation. I'm a Cloud/ software / language agnostic engineer who enjoys working independently, solving complex challenges and architecting solutions.

long running queries, top tables by size, number of connections by clients, Vacuum and bloat info, list of schemas, etc..

JVM dump collection & Analysis: A simple utility to get the thread dumps and heap dumps from the tomcat container and then push it to the analyzing tool.

SLA and SLI Metrics: With the help of different AWS and datadog metrics, measuring the Service Level Availability and Indicators for the month for all clients.

Event Driven Auto-Remediation: Handling of 3 types of alerts are automated using EDAR principles. The solution uses event source (Datadog/APM), SNS, Lambda, Python REST API and AWS SSM run commands to accomplish the tasks.

1. **Disk cleanup**: It cleans the disk space usage in the ec2 instances.
2. **JVM OutOfMemory & Memory Analysis**: It collects the heap dump, remediated the JVM, analyzes the heap dump and notifies over the Webex Chat.
3. **Kafka Offline Partitions or URP**: It identifies the kafka broker which has offline partitions or under replicated partitions and recovers the service.

Time Series metrics: Created kubernetes cronjob to run the python script (uses boto3 + Pandas) to collect the metrics using DB query from ~2000 RDS instances, feeds the output csv into InfluxDB. Grafana dashboards are created to display these metrics.

Data Persistence: Deployed MongoDB replica set and Redis cluster with kubernetes persistent volumes, claims and statefulset to achieve HA, persistence of data.

High-availability: Used AWS ASGs, kubernetes' affinity and anti affinity to schedule the distribution of pods across the AZs.

Scalability: Kubernetes' Horizontal Pod Autoscaling (HPA) for ruby on rails and fastapi pods level scaling and ASG are used for instance level scaling.

Reliability: Configured monitoring and alerting systems using prometheus, grafana, elasticsearch, kibana and beats. Implemented kubernetes' probes on pods.

Security and compliances: Developed solutions using kubernetes secrets, service accounts, private network topology, AWS's Secret manager, OKTA openID oAuth for authentication and authorization, JWT for communication between microservices. All the solutions honors GDPR policy.

Persistent Systems, Bengaluru, IN — Senior Team Lead - Devops

January 2017 - October 2018

- Designed and Developed one button blue-green deployment using Chef and AWS resources.
- Written infrastructure automation using Chef and AWS CloudFormation.
- Developed and managed multiple CICD pipelines.
- Troubleshoot the issues using AppDynamics, Splunk

OPEN SOURCE

https://github.com/RLOpenCatalyst/automationlibrary

INTERESTED AREAS

Architecturing the solutions, Distributed System, Chaos Engineering, Reliability, Artificial Intelligence and Machine Learning, Contribution to open source.

LANGUAGES

English, Hindi, Telugu

- Service reliability activities on INTUIT's e-commerce application.
- War room support during the US Tax season.
- Participate in the week based on-call rotation.

Relevance Lab, Bengaluru, IN — Senior Software Engineer - Devops

November 2015 - December 2016

- Designed and wrote Chef Automation library.
- Mastered in Configuration Management and Infrastructure as code skills using Chef, AWS, Docker, Terraform.
- One click automated dependent, multi-tier application deployment for different tech stacks

Cognizant, Chennai, IN — Senior Systems Executive

December 2014 - September 2015

- Learnt Docker, AWS, Chef, Ruby, Ruby On Rails.
- Developed Chef cookbooks and recipes for the client.
- Explored container technology using docker and wrote dockerfiles as a beginner for various sample applications.

Tech Mahindra, Chennai, IN — Software Engineer

August 2010 - December 2014

- Mastered in CICD, Version Management, Source Control management, Build and deployments.
- Implementation of Warehouse application for New Parts distribution centre.
- Active participation as a team member in application migration from Oracle 10g to 11G along with major enhancements.

EDUCATION

Indian Institute of Management, Calcutta,IN — Executive program in Supply chain Management

2015 - 2016

MVSR Engineering College, Hyderabad,IN — Bachelor of Engineering - Electronics and Communications

2006 - 2010