Pawan Shrestha

✆ 203.809.5170 | pawanshrestha@gmail.com

Lives in Waltham, MA. Works in AWS with experience in Databases | https://pkshr.com

SKILLS/ EXPOSURE

Cloud

AWS/ GCP

Scripting

Bash shell/ Python/ Ruby/ Power shell/ PERL/ JavaScript

Database

Oracle / Vertica / MySQL/ Aurora/ MS-SQL / PostgreSQL

Tools & Misc

Terraform/ **Jenkins**/ **Bamboo**/ **Chef**/ **Ansible**/ **GIT**/ **SVN**/ **Looker**/ DataDog/ **Sensu**/ **Nagios**/ Docker/ Kubernetes/ Nginx/ Tomcat/ VMWare/ Proxmox/ **Cloudhealth**

OS

Linux / Windows

EXPERIENCE

Pillar Biosciences, Natick, MA — *Sr. DevOps*

December 2020 - Present

Mastercard, Boston, MA — *Sr. DevOps*

August 2020 - December 2020

Uplight, Cambridge, MA — *Sr. SRE/DevOps*

August 2018 - August 2020

Cotiviti, Waltham, MA — *Sr. DevOps/DBA*
(Formerly Verscend, VeriskHealth, D2Hawkeye)

Oct 2006 - August 2018

D2Hawkeye Nepal, Kathmandu, Nepal — *JSP/Database Engineer*

August 2002 - Oct 2006

HitechValley, Kathmandu, Nepal — *Software Engineer*

November 1999 - July 2002

EDUCATION

Brandeis University, Waltham, MA — *M.S in Software Engg.*

May 2015

Institute of Engineering, Nepal — *B.S in Electronics & Telecomm Engineering.*

November 1999

Summary

- 21+ years of IT experience (DevOps/Database Administration/Programming/Scripting).
- Cloud (mainly AWS) and On-Premise datacenter hands-on experience in:
    - Comprehensive knowledge in multiple industry standard software configuration management tool suites and individual point products from system/tool administration through SDLC usage.
    - Efficient in performance monitoring, diagnostic and tuning, Backup, and Recovery, managing users and security.
    - Supportive Individual capable of Managing incidents, being proactive in terms of preventive maintenance and knowledgeable in terms of risk management.
    - Expert in Support and Maintenance, Performance tuning, shell scripting, upgrades and enhancements.
    - Knowledgeable in all aspects of software development life cycle (SDLC): requirements gathering, prototyping, database design, development, support, and QA.
    - Experience in programming and Scripting in Linux and Windows platform
    - AWS: Setting up VPC, Subnets, Routes, NACL, Security Groups, Load Balancers (ALB/NLB/ELB), Launch config/templates, ASG, EC2, RDS, Route53, Cloudwatch, IAM securities, Key/Secrets management, Certificate management, Lambda, API Gateways, EKS, Cost management, Batch
    - Experience in automating CI/CD process using multiple tools (e.g Chef, Jenkins, Bamboo, GIT, Bitbucket, GITLab, Puppet, Ansible) and custom scripts wherever necessary
    - Experience in Docker and Containers, orchestration (tools used Kubernetes, KOPS, Kadm)
- Good understanding and Implementation expertise in RAC, Oracle Exadata Machines, MSSQL Servers, MSSQL AlwaysOn Cluster, Vertica Cluster, Postgres, MySQL

Major Accomplishments

Pillar Biosciences:

- Migrated locally hosted Atlassian suites (Bitbucket, Confluence and JIRA) to an AWS VPC
    - Built a VPC, site-to-site VPN allowing company LAN and SSL VPN to AWS infra
    - Built all required AWS resources (VPC, Subnet, NACL, Security Groups, EC2, RDS, Route53, ALB, classic ELB) using terraform
- Automated manual AWS deployment process using Terraform, Ansible and Atlassian Bamboo
- Automated an on-premise alerting/monitoring tool Nagios using Atlassian JIRA, Atlassian Bitbucket, Atlassian Bamboo and Ansible
- Configured and setup multiple on-premise resources/services using IAC (Infrastructure As Code) - mainly via Ansible
- Automated cleanup of server storage (used by Atlassian Bamboo) using webhook in Atlassian JIRA, Web-API, Slack and Atlassian Bamboo

MasterCard:

- Migrated multiple CI/CD pipelines from Jenkins to GitLAB
- Prepared Infra as Code for multiple environments (Prod, UAT, Stg, Dev) using Terraform

Uplight:

- **Tenant management automation** - A solo project built on python that would create or delete a tenant and the associated resources. The objective of this project was to replace manual operation that involved 20+ steps across various internal/external portals, api calls, database creation, multiple AWS service operations (RDS schema creation, Route 53 record entries, S3 bucket/key creation). I built a Jenkins job with an interactive UI to gather inputs from the user and execute python with the input supplied.

- **AWS Cost/Usage monitoring/alert in Looker** - The company was using a third party vendor "CloudHealth" for AWS cost/usage monitoring and alerting. The company already had a Looker reporting tool. I taught myself Looker and initiated/implemented a Looker project that would provide all necessary information and alerts. I was able to discontinue "CloudHealth" and reduce the Ops cost. Since it was home grown and highly customizable, I was able to present various data points on a high level and granular level. The Looker dashboard has been the ***default browser homepage of my manager*** and the first thing he checks is the looker dashboard which gives a quick snapshot of expenses and usage across multiple AWS accounts.

- **Consolidated single page Jenkins deployment** - I proposed and led this project to integrate multiple jenkins jobs into a single jenkins job, for the release & deployment. I carried out this project in multiple phases. The first phase involved calling downstream jobs to show a working concept. Gradually, I attracted and involved all of the team members to contribute and we came up with the final version. As a by-product, we had to review existing Jenkins codes/ pipelines and were able to do a lot of cleaning and make the code lean. The Jenkins job has been our one-stop shop for the release and deployment. By injecting HTML elements dynamically within Jenkins UI, we were able to create a much efficient UI in Jenkins that would allow end user to provide inputs for the deployment

- **Migration of Chef-cookbooks/recipe to Terraform** - Uplight used Chef as infrastructure provisioning tool. Recipes/cookbooks (in **Ruby**) were being used to create all of the required AWS resource mainly - VPC, Subnet, SecurityGroups, LaunchConfig, Autoscaling Group, ALB (LB and Target), ELB (LB and Target), Route53 records, S3 buckets/keys, AMI (with packer), Docker image, RDS etc. I was part of the team to convert Ruby recipes/cookbooks to Terraform.

- **On-demand building a scaled down version of the entire stack on-the-fly for developers** - The objective of this project was to provide a highly scaled down version of a product stack on the fly to a developer so that he can build and test on an isolated environment. A developer would login into a Jenkins UI for this project and after making his selection on a set of predefined parameters, he/she would get a scaled down fully functioning environment. This environment could either be hosted in a docker container or a regular ec2 instance.

Cotiviti:

- Built a number of automation scripts - one of which was to build custom nagios plugins where I found a way to check connectivity of Microsoft Analysis Server (used for OLAP databases).
- Built a custom inventory of infrastructure (using Python, Django, MySQL) to manage records of systems. This was used as a source of truth for the inhouse monitoring tool that we built on nagios.
- Converted a single node server (that used a paid version Weblogic) into a cluster of multi-nodes server (using open-source Tomcat)
- Involved in on-premise data center migration during company acquire and merger - Lead a team responsible for Database migration
- Participated in building an automation tool for batch data processing in Oracle
- Built a custom ETL application and PL/SQL scripts for
    - homogenous data transfer for processing Oracle server to production Oracle RAC server
    - heterogeneous data transfer for processing Oracle server to production Vertica server
    - Did multiple POCs and research/reviews