Aravindh Sampathkumar

Systems Engineer | Site Reliability Engineer(SRE) | DevOps
aravindh@fastmail.com | +45 91 67 69 54
Website: aravindh.net | LinkedIn: aravindhsampath

PROFILE

A seasoned Systems Engineer with 12 years of broad experience with building, operating and analysing IT infrastructure with a heavy focus on Performance Engineering and Automation.

EDUCATION

MASTERS DEGREE IN COMPUTER SCIENCE

CLEMSON UNIVERSITY, USA - GRADUATED 2013
MASTER THESIS : VIRTUALIZING INTELLIGENT RIVER®
A quantitative and qualitative comparison of *KVM*, *Xen*, and *Linux Containers* based on their virtualization overhead, resource entitlement and isolation facilities, operational flexibility, scalability, and security.
[Publication on the web]

BACHELORS DEGREE IN INFORMATION TECHNOLOGY

ANNA UNIVERSITY, INDIA - GRADUATED 2007

SKILLS

Programming	⋆ Python ⋆ C ⋆ Bash ⋆ Golang ⋆ SQL ⋆ HTML/CSS
Infrastructure and Ops	Operating Systems : ⋆ GNU/Linux (RHEL derivatives) ⋆ FreeBSD Network services : ⋆ DNS ⋆ DHCP ⋆ SSH ⋆ NFS ⋆ DM/FreeIPA ⋆ Routing ⋆ Firewalls Storage : ⋆ ZFS (extensive experience) ⋆ LVM ⋆ XFS ⋆ Backups Infrastructure as code : Ansible + Python + Bash + Git ⋆ Gitlab-CI Monitoring and Analysis : ⋆ Prometheus ⋆ Elasticsearch ⋆ Grafana ⋆ Kibana Containers : ⋆ Docker/Podman ⋆ Kubernetes (currently exploring) Web : ⋆ Nginx ⋆ Web services in Python and Golang Public cloud : ⋆ Amazon Web Services(AWS) ⋆ Google Cloud Platform(GCP)
Performance Engineering	⋆ Performance benchmarking ⋆ Data analysis and visualisation ⋆ Capacity planning ⋆ Queueing theory
Data Analysis	⋆ SQL ⋆ Python ⋆ Jupyter ⋆ D3/C3 (Javascript) ⋆ Unix command line tools
System Design	⋆ Asynchronism ⋆ Caching ⋆ Performance vs Scalability trade-offs ⋆ Observability

EXPERIENCE

AARHUS UNIVERSITY | SYSTEMS ADMINISTRATOR

Aug 2017 – Current | Aarhus, Denmark

- Architected, built and manage a Slurm based HPC cluster that enables over 140 researchers from all over the world to run genetics and genomics workloads.
- Responsible for capacity planning, hardware procurement, and performance analysis for our research cluster.
- Operating the cluster using idempotent code using Ansible. E.g. Adding new compute node takes < 4 minutes.
- Automated day-to-day administration of a fleet of bare-metal CentOS/RHEL Linux servers (over 1000 CPU cores) using scripts written in Bash and Python.
- Planned and built several purpose-built [ZFS storage servers].
- Built and utilise a centralised monitoring and systems intelligence solution using Prometheus and Grafana(metrics), and ElasticSearch and Kibana(logs).
- Instituted a central identity system with MFA using FreeIPA and Duo.
- Wrote extensive documentation in the form of [user guides and documentation] for users of the research cluster, and a private Admin's handbook.
- Collaborate with and advice researchers on wide variety of system design and performance problems in HPC space.

Aravindh Sampathkumar

Systems Engineer | Site Reliability Engineer(SRE) | DevOps
aravindh@fastmail.com | +45 91 67 69 54
Website: aravindh.net | LinkedIn: aravindhsampath

NETAPP INC | Performance Engineer
Jan 2014 – Jun 2017 | Raleigh, USA

- Performance Analysis: Storage workload characterization, design and guidance of performance oriented projects – Data ONTAP, release performance.
- Benchmarking: Made use of micro benchmarks and specialized in custom workload generation with [SPEC SFS 2014].
- Validation and analysis: Served as one of NetApp's representative at SPEC SFS sub committee and helped validate benchmarking harness and new workloads for SFS 2014 SP2.
- Tooling and Automation: Developed easy-to-use CLI tools and Flask based web APIs for accessing and analysing performance data. Experimented with mining performance data for behaviour search.

Achievements:

- Analysed and helped resolve performance issues reported by customers thereby aiding business continuity.
- Tools that I developed and productised are used by analysts and developers across the organization - saving time and improving efficiency.
- My work on improving parallelism and thread management resulted in improved product performance.

INTELLIGENT RIVER PROJECT, CLEMSON UNIVERSITY | Back-end and Infrastructure Engineer
Jan 2012 – Dec 2013 | Clemson, USA

- Infrastructure: Designed and deployed highly available infrastructure based on clustered RabbitMQ, replicated MongoDB and polyglot middleware.
- Logs based intelligence: Deployed an ELK stack for monitoring and added intelligence using Simple Event Correlator(SEC).

Achievements:

- [Intelligentriver.org] - Live monitoring and analysing water quality at several locations across south east USA.

BANK OF AMERICA MERRILL LYNCH, (BA CONTINUUM LTD) | Senior Software Engineer
Feb 2011 – Aug 2011 | Chennai, India

- Developed a database-backed, financial document queueing application.
- Performed impact analysis, change management, and peer code reviews.

TATA CONSULTANCY SERVICES LTD. | Asst. Systems Engineer
Jan 2008 – Jan 2011 | Chennai, India

- Experimented with and administered hundreds of Linux Virtual Servers on IBM System Z.
- Benchmarked and presented several proof-of-concept architectures to banking and financial companies.

TRAINING

GUERILLA CAPACITY AND PERFORMANCE (GCAP) | Graduate
Jan 2020

- Intensive study of Queueing models, Scalability laws, Performance analysis and their applications with Dr. Neil J. Gunther.

PERSONAL

Nationality	Indian
Work permit	India (passport), Denmark (job based work/residence permit), EU(partner based).
Languages	English (fluent), Tamil (native)
Free time interests	Woodworking, online privacy, homelab

REFERENCES

Available upon request