Mustafa Arif

HPC Professional

I am a committed professional possessing robust leadership skills and extensive expertise in High-Performance Computing (HPC), Artificial Intelligence (AI), and Cloud infrastructure operations, with a specific focus on establishing scalable computing platforms for research and development organizations. My career spans over a decade in the field of HPC infrastructure operations, during which I have consistently excelled in devising and executing technical strategies to achieve organizational objectives and drive success.

emarif12@gmail.com

+447767730003

Cambridge, United Kingdom

mustafaarif.com

linkedin.com/in/mustafaarif

github.com/mustafaarif

TECHNICAL SKILLS

HPC Platforms	Cray, DELL, HPE, Nvidia DGX
Cloud	AWS, Azure, OpenStack
Network	InfiniBand, Ethernet, Firewall
Automation Tools	Ansible, Bright Cluster Manager, CRAY Cluster Manager
Parallel Frameworks	MPI/OPENMP, CUDA, Matlab Parallel, Python Dask
Infrastructure Monitoring	Checkmk, Prometheus, Graffana
HPC WLM	SLURM, PBS
File System	Lustre, Panasas, NFS
Containers	Docker, Singularity, Kubernetes, Nvidia GPU Cloud.
Data Analytics	TensorFlow, scikit-learn, RAPIDS
Programming	R, Python, C/C++, Bash
Virtualization	KVM, OpenStack

WORK EXPERIENCE

HPC Team Lead

European Bioinformatics Institute

06/2023 - Present *Cambridge, UK*

Achievements/Tasks

- Leading HPC Team to deliver HPC Compute services for over 800 research users.
- Managing multiple HPC (DELL) and GPU (A100s) Clusters with SLURM Job scheduler, IB network and Vast storage for scratch space.
- Implementing in-band and out-of-band infrastructure monitoring solutions (Checkmk) to ensure the reliability, performance, and security of EMBL-EBI HPC systems.
- Deploying reporting interfaces (XDMoD) for SLURM Job scheduler to get visibility into infrastructure usage and resource utilisation.
- Optimising provisioning methods by introducing stateless provisioning using Warewulf.
- Enabling remote visualization to minimize data transfers and enable users to view results in close proximity to the computing and storage resources.

WORK EXPERIENCE

Systems Manager (HPC)
KTH Royal University

01/2022 - 06/2023

Parallelldatorcentrum (PDC), Stockholm, Sweden

Achievements/Tasks

- Responsible for HPC (Compute, GPU, Storage and Network) Infrastructure planning, deployment and operations.
- Advising leadership on budgeting, planning and procuring compute infrastructure for research.
- Overseeing Cray EX HPC system with 70k+ AMD Epyc cores, 12PB Lustre storage, and substantial GPU compute for AI and ML workloads.
- Designing & implementing private OpenStack cloud for HPC users to support their HPC computational needs.
- Implementing Interactive HPC infrastructure for remote visualization and interactive use of SLURM job scheduler.
- Managing routine infrastructure tasks, including upgrades, monitoring, and documentation. This encompasses applying security patches, automating backups, streamlining system administration, and maintaining current documentation.
- Collaborating closely with stakeholders to meet end users' computational needs, which involves liaising with application teams and occasionally offering direct technical guidance on HPC infrastructure utilization.
- Training and consulting with users on scientific computing and AI projects.

Sr. IT Consultant (HPC)
Texas A&M University

03/2014 - 01/2022

Research Computing group at TAMUQ

Achievements/Tasks

- Involved in multiple projects related to HPC systems acquisition, operations and data platforms development.
- Working with team in planning and defining strategy for future scientific computing growth in the organization.
- Successfully administered Linux based servers which includes; Cray XC 40 HPC system, Bull HPC Cluster, GPU cluster, Hadoop Cluster and Bare metal servers. Installation and compilation of scientific packages on the system.
- Consulting scientific computing users in code development from different domains e.g. Data science, Image Processing, Astronomy, Fluid Codes, Multi-Engineering domains and Bioinformatics.
- Promoted use of HPC and GPU computing among user community by organizing training and identifying use cases. Consulting users on their research projects and assisting them in parallel programming to port their workloads on HPC system.
- Containerized multiple HPC applications for research users so to allow them to run computation with more control on software stack. Thus allowing users to quickly deploy their scientific applications on the system and perform cross system migrations with much ease.
- Delivering training on scientific software and packages. Also, organizing external training where required to make sure that scientific community training requirements are fulfilled. This has a great positive impact in boarding new users on the system and optimizing workloads of existing users.
- Preparing technical guides of HPC system and making them available on internal wiki page for easy accessibility. https://rc-docs.qatar.tamu.edu/.

CERTIFICATIONS

Red Hat Certified Engineer (RHCE)
(02/2022 - 02/2025)
Certification ID: 150-041-309

Red Hat Certified System Administrator (RHCSA)
(02/2022 - 02/2025)
Certification ID: 150-041-309

AI in Data Center - NVIDIA (03/2020 - Present)

EDUCATION

MS Computer Engineering
National University of Science and Technology, Pakistan

2012 - 2016

Thesis

- Parallel Gene Ranking framework for early prediction of cancer

BS Computer Engineering
Comsats University, Pakistan

2007 - 2011

CONFERENCES / COURSES

Project Management Professional Training (02/2022)

TIA Education, Udemy

- Learned various project leadership skills and processes to build connection between projects and organizational strategy.

Supercomputing Conference (2019)

Denver, Colorado, US

- This year focus was on use of HPC infrastructure in AI. Learned new advancements in HPC technologies for AI.

OpenStack Administration (2017)

LFS452-10467 - The Linux Foundation

- OpenStack is an open source cloud computing platform which is deployed as infrastructure as a service. The administration training was essential in deploying and maintaining a cloud platform for research users.

Docker Administration & Operations (2017)

RMS-KS-DC-10038-2609/4807 - Koenig Solution Pvt. Ltd.

- The docker administration training helped in understanding micro services concept and how this framework can be used in research computing environment.

Architecting on AWS Accelerator (2020)

ExitCertified, Virtual Training

- A comprehensive course on making architectural decisions based on AWS architectural principles and best practices.

Advanced Python: Best Practices and Design Patterns (06/2018)

Learning Tree International, UK

- Explored Object-oriented programming features of python. Learned Unit testing functions of python which was very beneficial in regression testing of HPC system after the upgrades.

DDN ExaScaler Storage Administration (2017)

DataDirect Networks

- DDN ExaScaler systems are Lustre based file system. The new HPC system at TAMUQ was integrated with DDN based storage. This training was helpful in learning core concepts of DDN Lustre storage systems.

Data Science & Data Engineering Boot camp (2016)

Data Science Dojo, Chicago, USA

- Data science boot camp was 5 day extensive training session. It was quite helpful in understanding how data analytics is being used in industry and the potential use cases.

PROJECTS

Interactive High Performance Computing (09/2022 - 03/2023)

- Interactive use of HPC resources using ThinLinc and gfxlauncher tools. Allowing user to remotely connect to a desktop environment where they can interactively launch SLURM Jobs, Run GUI applications with hardware rendering.

Private Cloud for HPC users (10/2022 - 04/2023)

- At PDC, KTH there was in immense need of a private cloud to facilitate users for Data and HPC workflows. Consulted with various stakeholders and designed an efficient solution which will be deployed end of this year. Once in production this system will facilitate HPC and AI users for running Data simulations and Implementing Machine learning Pipelines.

Hadoop Cluster Deployment (04/2020 - 08/2020)

- Planned and deployed Hadoop cluster on-perm using Singularity Containers. YARN and SPARK was setup for users to run their big data applications. The cluster could scale and shrink based upon user requirements.

HPC Data Visualization (11/2019 - 01/2020)

- Compute nodes in HPC system are meant for fast processing and doesn't contain powerful graphic processing unit. Users had to move their datasets from HPC system to GPU capable workstations for data visualization. Planned and deployed HPC Visualization infrastructure which enabled users to visualize scientific data sets on powerful Viz nodes without offloading data from HPC storage.

PROJECTS

Containers for HPC and AI (09/2019 - 12/2019)

- Introduced Singularity containers for the community which enabled users to bring their applications to supercomputer with least effort and in minimal time. Worked with researchers to deploy containerized scientific codes on HPC system which otherwise would not have been possible on native operating system.

AI/ML Cluster Deployment (01/2019 - 07/2019)

- Planned procurement and installation of GPU Cluster for research needs. Successfully completed physical installation and integration with the current infrastructure. The configuration of the software stack on the system was planned in accordance with emerging technologies and research users needs. Prepared training for users to help them get started on the new system.

HPC Analytics & Accounting (2018 - 2020)

- Implemented completely open-source HPC compute resources accounting and analytics framework. Resource accounting to identify top users in community and observing usage trends. Whereas, workload logs analytics gave insights on improving resource scheduling and optimizing resource distribution in queues.

HPC System deployment (2016 - 2017)

- In 2016 we acquired new HPC system. Worked closely with facilities to prepare data center to host new system. Compiled required software and packages on new system for users. Made sure that required packages and compilers are available. Migrated data from old storage system to new. Provided consultation and training to users to make them comfortable in using new system.

INTERESTS

Coaching | Traveling | Networking | Internet of things | 3D Printing

SOFT SKILLS

Self-motivation | Teamwork | Decision Making | Ability to Work Under Pressure | Adaptability