



Cost-effective deep learning infrastructure with NVIDIA GPU

Aatiz Ghimire^{*a,b,c}, Shahnawaz Alam^b, Siman Giri^b, and Madhav Prasad Ghimire^a

^aCentral Department of Physics, Institute of Science and Technology, Tribhuvan University, Nepal.

^bHerald College Kathmandu, University of Wolverhampton, United Kingdom.

^cSchool of Mathematical Sciences, Institute of Science and Technology, Tribhuvan University, Nepal.

Abstract

The growing demand for computational power is driven by advancements in deep learning, the increasing need for big data processing, and the requirements of scientific simulations for academic and research purposes. Developing countries like Nepal often struggle with the resources needed to invest in new and better hardware for these purposes. However, optimizing and building on existing technology can still meet these computing demands effectively. To address these needs, we built a cluster using four NVIDIA GeForce GTX 1650 GPUs. The cluster consists of four nodes: one master node that controls and manages the entire cluster, and three compute nodes dedicated to processing tasks. The master node is equipped with all necessary software for package management, resource scheduling, and deployment, such as Anaconda and Slurm. In addition, a Network File Storage (NFS) system was integrated to provide the additional storage required by the cluster. Given that the cluster is accessible via ssh by a public domain address, which poses significant cybersecurity risks, we implemented fail2ban to mitigate brute force attacks and enhance security. Despite the continuous challenges encountered during the design and implementation process, this project demonstrates how powerful computational clusters can be built to handle resource-intensive tasks in various demanding fields.

Keywords: Deep learning infrastructure; Beowulf cluster; High-performance computing (HPC); GPU cluster architecture.

1. Introduction

Deep learning has become a cornerstone of modern artificial intelligence, driving advances in various fields such as natural language processing, computer vision, and autonomous systems. However, training deep learning models often requires substantial computational resources, which can be prohibitively expensive for many institutions and researchers. Cloud-based GPU solutions, while offering scalability, come with recurring costs and potential data privacy concerns, making them less feasible for long-term or resource-intensive projects [1, 2, 3].

High-performance computing (HPC) or Clusters is increasingly shifting towards heterogeneous GPU-based systems, driven by the growing demand for massively parallel computing in deep learning and scientific applications [4]. With the exponential growth of AI-driven research and simulations, traditional CPU-based architectures often struggle to keep up with computational demands. GPUs, with their high throughput and parallel processing capabilities, have emerged as a crucial component in accelerating tasks such as model training, data analysis, and large-scale simulations.

Building an in-house, cost-effective deep learning infrastructure offers a practical alternative to address these challenges. Local clusters equipped with GPUs not only reduce operational costs but also provide greater control over data and computational resources. However, achieving a balance between performance, scalability, and affordability remains a significant challenge, especially for setups that use consumer-grade GPUs and standard hardware components [5].

This study focuses on designing and implementing a cost-effective deep learning cluster using readily available hardware

and open source software. The objectives are to evaluate the feasibility of such a system, address its limitations, and highlight its potential for supporting computational research in resource-constrained environments. Using locally available resources and optimizing configurations, this study aims to contribute to the growing demand for accessible and efficient AI infrastructure.

The Beowulf Raspberry Pi cluster has gained attention as an affordable and scalable alternative for parallel computing applications. Various studies have explored its implementation, highlighting its potential for educational and experimental purposes [6, 7, 8, 9, 10, 11]. While Raspberry Pi clusters provide a cost-effective way to introduce students and researchers to distributed computing concepts, their performance is limited due to the low computational power of individual nodes. Furthermore, Raspberry Pi lacks dedicated GPU acceleration, restricting its usability for deep learning and large-scale parallel computing tasks. This limitation makes it unsuitable for workloads requiring significant matrix operations, such as AI model training.

Another notable low-cost alternative is the Odroid-XU4 board, which offers more processing capability than Raspberry Pi. Studies have demonstrated its viability in HPC education and small-scale computing clusters [12]. Odroid-XU4 features an ARM-based architecture with multi-core processing, making it more efficient for parallel computations than Raspberry Pi. However, similar to Raspberry Pi, it lacks dedicated GPU support, making it impractical for deep learning tasks that heavily rely on GPU acceleration. While software optimization can help mitigate some performance constraints, these devices are not suitable for handling large-scale AI and scientific computing workloads that demand high floating-point performance.

Recycling old computers by integrating them into a Rocks clus-

^{*}Corresponding author. Email: hello@aatizghimire.com.np

ter is another cost-efficient strategy for setting up HPC environments [13]. Many academic institutions have repurposed outdated lab machines to build functional compute clusters, extending their usability for scientific computing and parallel processing workloads [14]. However, these older systems often rely solely on CPUs and lack modern GPUs, which are essential for accelerating deep learning computations. This absence of GPU support significantly limits their effectiveness in AI research and computational tasks that require tensor operations. While CPU-based clustering can still be useful for general-purpose parallel processing, it falls short in performance compared to modern GPU-accelerated systems.

Despite the availability of low-cost solutions, pre-compiled DGX OS, designed specifically for deep learning systems, does not support commodity hardware, restricting its deployment to specialized NVIDIA DGX systems [15]. The NVIDIA DGX systems, optimized for deep learning workloads, has been widely adopted in AI research due to its integrated software stack and high-performance GPUs [16, 17, 18, 19]. Unlike the previously mentioned hardware solutions, NVIDIA DGX systems is built with dedicated NVIDIA GPUs, enabling high-speed parallel computations essential for deep learning. However, its high cost makes it inaccessible for many institutions and researchers, pushing them to explore alternative, budget-friendly solutions that attempt to bridge the gap between affordability and computational efficiency.

2. Materials and methods

On the hardware side, our aim was to maintain a minimal configuration to make use of commonly available computers in standard computing labs. Our setup consisted of four computers equipped with GPUs, connected via a switch using CAT6 Ethernet cables for network communication. We assigned the four computers as Master, C1, C2, and C3 nodes.




Figure 1: Cluster network configuration.

The basic configuration of our cluster is shown in the figure above. The computers are connected to a switch using Ethernet cables. Alternatively, a router can be used instead of a switch. One port of the switch is connected to the internet, enabling remote access to the cluster. We configured Network Address Translation (NAT) to convert the cluster’s local private IP addresses (Master node IP) into a global public IP address for internet accessibility.

Initially, we attempted to set up the cluster using Ubuntu Server 22.04.1. However, it required extensive package installations, had higher configuration complexity, and lacked enterprise-level features such as robust security, package stability, and infrequent updates. To address these limitations, we switched to Rocky Linux 9.4, which provides RHEL features. We performed a clean installation of Rocky Linux on each cluster node and updated them to the

latest version. For simplicity in installation and cluster configuration, we only set up the root user with same password and did not create additional user accounts at this stage [20].

The next step was to configure static IP addresses, which are essential for inter-server communication needed for tasks such as SLURM, MPI, and data transfer. To achieve this, we used “nmtui”, a Network Manager tool, to assign static IP addresses to the nodes, as illustrated in Fig. 1.

To streamline communication within the cluster, we configured the ‘/etc/hostname’ file on all nodes. The host file maps IP addresses to hostnames, allowing us to use simple names instead of IP addresses for communication. We assigned ‘10.80.0.100’ to the master node and sequential IP addresses (e.g., ‘10.80.0.101’, ‘10.80.0.102’, ‘10.80.0.103’) to the compute nodes ‘c1’, ‘c2’, and ‘c3’. Additionally, we updated the ‘/etc/hostname’ file on each node to define its hostname, ensuring consistent identification across the cluster. This configuration simplifies node communication and is essential for seamless cluster operations, including MPI, SLURM, and other distributed tasks.

We installed the SSH package on all nodes to enable secure communication within the cluster. To streamline operations, we configured passwordless SSH, allowing quick and seamless login between nodes without requiring repeated password entry. This setup simplifies administrative tasks and facilitates efficient execution of distributed processes across the cluster [21].

All logins to the cluster are routed through the master node, which serves as the central point for user access and management. The master node is equipped with all necessary software for package management, resource scheduling, and deployment, ensuring efficient coordination and operation of the cluster. This centralized setup simplifies user management and resource allocation across the cluster.

```

aatz@notamac:~$ ssh root@10.80.0.100
root@10.80.0.100's password:
*****
*                                     *
*           Welcome to Pluto Cluster   *
*                                     *
*           Developed by:               *
*   Aatz Ghimire, Shahnawaz Alam, & Siman Giri *
*                                     *
* Rocky Linux 9.4 (Blue Onyx) (RHEL)    *
*                                     *
* General Purpose Nodes:                *
* 3 compute nodes with:                 *
* - 12 cores Intel(R) Core(TM) i5-10400F @ 2.90GHz *
* - 16 GB RAM                           *
* - 4 GB NVIDIA RTX 1650Ti              *
*                                     *
* Internal Network: 100Mb/s Ethernet Network *
*                                     *
* SLURM 22.05                           *
*                                     *
* For a guide on Pluto Cluster:          *
* https://pluto.heraldcollege.edu.np    *
* For support: pluto@heraldcollege.edu.np *
*                                     *
* In EVIDENCE:                          *
* - Please, use STORAGE0 Folder for work! *
*                                     *
*****
  
```

Figure 2: Cluster login screen via SSH.

To streamline software deployment across multiple nodes in the cluster, we utilize pdsh, a parallel shell tool that enables simultaneous execution of commands. First, pdsh is installed on the

head node using `dnf install -y pdsh`, ensuring SSH access is configured for passwordless authentication across all compute nodes. With `pdsh`, software packages can be installed in parallel, significantly reducing deployment time. For instance, executing `pdsh -w node[1-4] 'sudo dnf install -y package-name'` installs the required software on all designated nodes simultaneously. This method enhances administrative efficiency, ensuring consistency across the cluster without manual intervention, making it well-suited for high-performance computing environments running Rocky Linux [22].

Additionally, maintaining time synchronization across the cluster is crucial for job scheduling, logging accuracy, and distributed computing. We achieve this by installing and configuring Network Time Protocol (NTP) on all nodes. Using `pdsh`, NTP can be installed in parallel with `pdsh -w node[1-4] 'sudo dnf install -y chrony'`. The NTP service is then enabled and started with `pdsh -w node[1-4] 'sudo systemctl enable --now chronyd'`. The head node is configured as the primary NTP server, while compute nodes sync their time with it. This setup ensures clock consistency, preventing synchronization issues that could disrupt HPC workloads and SLURM job execution.

Similarly, we installed MUNGE, which is a lightweight authentication service essential for secure, token-based communication between nodes in clusters. It simplifies user authentication, integrates seamlessly with job schedulers like SLURM, and provides scalability and ease of maintenance across clusters, making it ideal for efficient cluster management [23].

We added an additional 1TB of storage to the system to accommodate large-scale dataset downloads and enable long-term storage of results. The storage was mounted at `/mnt/storage0` and bind-mounted to each user's directory as `/home/username/storage0`. This setup allows users to access the additional storage directly from their home directory, and we recommend utilizing this location for data management [24].

We set up a Network File System (NFS) on `/mnt/storage0` to enable efficient file sharing between the server and client nodes in our cluster. The configuration involved installing the `'nfs-utils'` package, creating shared directories with appropriate permissions, and defining NFS exports in the `'/etc/exports'` file. We configured the firewall to allow NFS services and created mount points on the client nodes. The NFS shares were mounted manually for testing and configured in `'/etc/fstab'` for persistence after reboots. This setup ensures secure and reliable file sharing across the cluster [25].

We utilized FreeIPA for centralized user management and synchronization across all cluster nodes. FreeIPA provided seamless integration of user authentication and access control, ensuring consistent user accounts and credentials throughout the system. This streamlined administrative tasks, enhanced security, and simplified user management in the cluster environment.

User management in HPC clusters requires a balance between security and ease of administration. While NIS is widely used in many HPC environments due to its simplicity, it is considered less secure than modern alternatives [26, 27]. FreeIPA, which integrates LDAP and Kerberos, offers enhanced security and centralized authentication but is more complex to install and configure [28]. Our cluster utilizes FreeIPA with LDAP to provide a robust identity management system, ensuring secure user authentication and streamlined access control. However, for clusters prioritizing ease of use over security, NIS remains a simpler alternative despite its vulnerabilities.

Slurm is an open-source workload manager designed for efficient resource allocation, job scheduling, and queue management in high-performance computing (HPC) clusters. To set up Slurm in our cluster, we installed the necessary packages on all nodes

```
[root@login ~]# sinfo
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
normal*    up       infinite    3     idle c[1-3]
```

Figure 3: Slurm command `sinfo` for listing cluster.

and configured the `'slurm.conf'` file to match our cluster's specifications. The configuration file was distributed to all nodes, and appropriate permissions were set for Slurm directories. On the controller node, the Slurm controller daemon (`'slurmd'`) was enabled and started, while the Slurm node daemon (`'slurmmd'`) was activated on all worker nodes. The installation was verified using commands like `'sinfo'` for cluster information and `'srun'` for test jobs. This setup ensures efficient job scheduling and resource management, optimizing the performance of the cluster [23]. We configured the master node to support computations but chose not to use it for this purpose. Utilizing the master node for processing tasks would add significant load to it, as it is already responsible for managing multiple tasks, including job scheduling, resource allocation, and acting as the login node for multiple users. Prioritizing these critical functions ensures the stability and efficiency of the cluster.

At the time of setup, NVIDIA driver installation was not supported directly through the Rocky Linux package manager (`dnf`). Therefore, we manually installed the driver using the `'run'` file obtained from the NVIDIA website. The GPU driver was installed first, followed by the installation of CUDA to enable GPU-accelerated computations in the cluster.

To enable GPU integration in SLURM, we configured the `'slurm.conf'` file to recognize GPU resources by adding the `'Gres'` (Generic Resources) parameter. Each node was configured with its specific GPU details, including type and count. SLURM's `'gres.conf'` file was updated accordingly, ensuring seamless allocation of GPU resources to jobs. This setup allowed efficient scheduling and utilization of GPUs across the cluster [29].

We installed Lmod, a Lua-based environment modules system, to efficiently manage software environments in the cluster. Lmod enables users to dynamically load, unload, and switch between different versions of software and their dependencies. Using Lmod, we created module files for various versions of Python, PyTorch, etc., allowing users to easily switch between them based on their project requirements. After installation, we configured the module files directory and integrated Lmod with the shell environment by updating the system's profile files. This setup, powered by Lua scripting for custom module configurations, ensures flexibility and simplicity in managing diverse software stacks, optimizing the cluster for various computational workflows [30].

We configured Lmod and Lua on all cluster nodes to ensure consistent environment management and seamless job execution across the cluster. Lmod and Lua were installed on each node, and module files were synchronized using a shared directory accessible to all nodes. The shell environment on each node was updated to load Lmod and set the appropriate `'MODULEPATH'`. This setup ensures that all nodes can access the same software environments, allowing SLURM to execute jobs reliably and efficiently with the required configurations.

To further streamline software management, we integrated EasyBuild, an automated framework for building and managing software installations in HPC environments. EasyBuild simplifies the process of compiling and installing complex scientific software by handling dependencies and environment configurations automatically. By using EasyBuild alongside Lmod, we generated module files dynamically, ensuring that all installed software could be easily managed within the cluster's module system. This approach reduces manual configuration efforts while maintaining re-

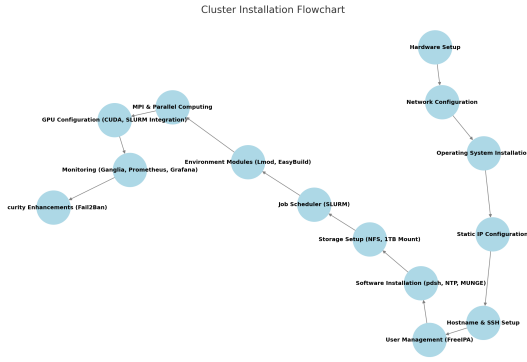


Figure 4: Deep learning cluster installation flowchart.

producibility and consistency across all nodes. Additionally, EasyBuild allows us to install optimized versions of libraries and applications, improving performance for computational workloads [31].

While our cluster primarily utilizes EasyBuild, Spack [32] is also a powerful and flexible alternative for software deployment. Spack enables users to install multiple software versions and manage dependencies efficiently, offering a more modular approach than traditional package managers. Its extensive package repository and customizable builds make it a preferred choice for many HPC environments. Although we opted for EasyBuild due to its structured installation and seamless module integration, Spack remains an excellent choice for clusters requiring greater flexibility in software management.

We installed Anaconda in a shared directory accessible across the cluster and integrated it with Lmod for dynamic environment management. A custom Lmod module file was created to configure Anaconda, setting environment variables like 'PATH' for seamless activation. Using the shared directory, all cluster nodes can access the Anaconda installation, enabling users to load and unload Anaconda dynamically through Lmod. This centralized setup allows users to create isolated Conda environments, which can be loaded as part of SLURM job scripts, ensuring consistent software environments for diverse computational tasks across the cluster.

We installed multiple versions of Python alongside Anaconda in a shared directory to provide flexibility for various computational tasks. Each Python version was configured with its own Lmod module file, allowing users to dynamically load the desired version as needed. These Python modules, like Anaconda, are accessible across all cluster nodes via the shared directory. Users can load specific Python versions using Lmod commands and utilize them in SLURM job scripts, ensuring compatibility with diverse project requirements while maintaining a consistent and centralized environment management system.

We set up the Message Passing Interface (MPI) using MPICH on the cluster to enable efficient parallel computing. Open MPI was installed on all nodes, ensuring the availability of runtime libraries, compiler wrappers, and development tools. Environment variables such as 'PATH' and 'LD_LIBRARY_PATH' were configured globally or managed dynamically using Lmod. We verified the installation by compiling and running test MPI programs using 'mpicc' and 'mpirun'. Static IPs and passwordless SSH were configured already before to facilitate seamless communication between nodes. Additionally, MPI was integrated with SLURM, utilizing 'srun' for job execution to leverage SLURM's resource management capabilities. This setup provides a robust framework for scalable parallel computing in the cluster.

The NVIDIA GTX 1650 GPUs in our cluster support CUDA for GPU-accelerated parallel computations but lack the features required for clustering GPUs across nodes, such as GPUDirect RDMA. This

limitation prevents the GPUs from being used together in a unified multi-GPU setup. [33] Instead, each GPU operates independently, and tasks must be run on individual GPUs within their respective nodes. While this restricts the scalability of GPU workloads, it remains suitable for single-node GPU computations and smaller parallel tasks.

Monitoring is a critical aspect of HPC cluster management, ensuring system health, performance optimization, and fault detection. We installed Ganglia due to its ease of setup and widespread use in HPC environments, allowing efficient tracking of system metrics across all compute nodes [13, 10]. Ganglia provides real-time resource utilization insights, helping to maintain workload balance and optimize job scheduling.

Additionally, we tested Prometheus (on <http://10.80.0.100:9090>) and Grafana (on <http://10.80.0.100:3000>) as alternative monitoring solutions. Prometheus, with its time-series monitoring and alerting capabilities, allows detailed system metric collection via exporters [17]. Grafana, when integrated with Prometheus, provides an intuitive visualization interface for tracking CPU, memory, and GPU usage in real time. While both Prometheus and Grafana are viable options under the given IP configurations, we preferred Ganglia for its simpler deployment and efficient monitoring of our SLURM-managed cluster.

```

aattiz@notamac: ~$ ssh root@10.80.0.100
root@10.80.0.100's password:
Activate the web console with: systemctl enable --now cockpit.socket

Last failed login: Tue Oct 1 11:06:29 +0545 2024 from 36.138.99.175 on ssh:notty
There were 1880 failed login attempts since the last successful login.
Last login: Mon Sep 30 17:59:29 2024 from 10.80.0.254
[root@login ~]#
  
```

Figure 5: Failed login attempts in cluster.

Our cluster faced a significant cybersecurity threat daily due to internet connectivity with nearly 2000-6000 failed login attempts to the root account within 24 hours, which we suspect were primarily bot-driven attacks. To mitigate this risk, we disable login with root account in SSH and we installed the Fail2Ban package to enhance SSH security. Fail2Ban monitors login attempts and, after six consecutive failed password entries, automatically blacklists the offending IP address for an hour, blocking further SSH communication. This proactive approach effectively reduced unauthorized access attempts and strengthened the cluster's overall security.

3. Results and discussion

The implementation of our cluster setup demonstrates a cost-effective and practical solution for high-performance computing tasks. The master node efficiently managed resource allocation and user authentication, ensuring seamless operation across the cluster. Security was significantly enhanced with the integration of Fail2Ban, which successfully reduced unauthorized login attempts by blacklisting suspicious IP addresses.

The individual GPU performance was validated for CUDA computations, proving effective for single-node tasks despite the hardware limitation of not being able to cluster GPUs. MPI communication over Ethernet was reliable, providing acceptable latency and bandwidth for medium-scale workloads. Additionally, the integration of MPI with SLURM ensured efficient resource scheduling for distributed computations.

The adoption of Lmod and Lua simplified software environment management, enabling users to switch dynamically between various versions of Python, and other required software. The setup of Anaconda further streamlined Python-based workflows, offering flexibility and efficiency for diverse computational tasks. These features significantly improved the user experience in managing software environments within the cluster.



Figure 6: Cluster setup.

The addition of 1TB of shared storage, bind-mounted to user directories, facilitated the handling of large-scale datasets and ensured easy access to results. This enhanced the system's ability to support data-intensive applications and long-term storage needs.

Containerization plays a crucial role in deep learning workflows by providing isolated and reproducible environments. In our cluster, we evaluated both Singularity and Docker as potential solutions for containerized deep learning environments. While Docker is widely used for its flexibility and ease of deployment, it requires root privileges, making it less suitable for shared HPC environments. Conversely, Singularity enables users to run containers securely without requiring elevated permissions, making it a more HPC-friendly alternative. Both containerization tools facilitate seamless execution of deep learning frameworks like TensorFlow and PyTorch, ensuring dependency consistency across different nodes. Based on our evaluation, Singularity is preferred for our cluster due to its security model and compatibility with HPC resource managers like SLURM, but Docker remains a viable option for standalone deep learning projects and development environments [16] [19] [34].

In summary, the results demonstrate the feasibility of building a robust HPC cluster using commonly available hardware and open-source software. While the system effectively balances cost, scalability, and performance, limitations such as the inability to cluster GPUs and reliance on Ethernet for MPI communication introduce constraints for high-demand, large-scale applications. Future improvements could focus on integrating high-speed interconnects and upgrading to HPC-oriented GPUs to enhance the cluster's performance and scalability further.

Table 1: Cluster configuration summary.

Component	Specification
Master node	10.80.0.100, Intel i5 CPU, 16 GB RAM
Compute nodes	c1, c2, c3 (10.80.0.101-103), Intel i5 CPU, 16 GB RAM
GPUs	NVIDIA GTX 1650 (on all nodes)
Storage	1TB (shared across all nodes)
Network	Ethernet (1 Gbps)

Table 2: Software stack.

Software	Version	Purpose
Rocky linux	9.4	Operating system
SLURM	22.05	Job scheduling
Lmod	Installed	Environment management
Anaconda	2023.09	Python-based workflows

4. Conclusion

Setting up a local GPU cluster within a small department or university presents a highly cost-effective alternative to renting cloud GPUs. Our cost analysis demonstrates that operating a locally installed cluster with four NVIDIA GTX 1650 GPUs (4 x 4GB = 16GB total) incurs a monthly electricity cost of only NPR 5,760 (based on an 800W power consumption per node). In contrast, renting an NVIDIA T4 GPU (16GB) in the cloud costs approximately NPR 123.50 per hour, amounting to NPR 29,640 per month for similar usage [35, 36].

This stark difference underscores the long-term financial benefits of local GPU clusters, particularly in regions like Nepal, where electricity costs are relatively low. Beyond cost savings, a locally managed cluster offers greater control over data privacy, independence from internet connectivity, and uninterrupted access to computational resources, making it an ideal solution for research and education. While the initial hardware investment may seem substantial, many institutions can leverage existing computing labs to set up these clusters efficiently. Ultimately, the autonomy, sustainability, and reduced recurring costs of an in-house GPU cluster make it a far more viable and strategic investment for academic and research institutions compared to reliance on cloud-based alternatives.

Acknowledgments

We express our sincere gratitude to Herald College Kathmandu for providing the research space, hardware, and infrastructure essential for this project.

For detailed instructions on installing, and configuring, we have made our setup process and related scripts available at <https://github.com/aatizghimire/pluto-cluster>. For running this cluster, We have also made slurm with gpu sbatch file and instruction at <https://github.com/aatizghimire/deep-learning-gpu-slurm-template>.

References

- [1] Chahal D, Mishra M, Palepu S & Singhal R. Performance and cost comparison of cloud services for deep learning workload. In: *Companion of the ACM/SPEC International Conference on Performance Engineering*, ICPE '21. Association for Computing Machinery, New York, NY, USA (2021). ISBN 9781450383318, p. 49–55. <https://doi.org/10.1145/3447545.3451184>.
- [2] Lawrence J, Malmsten J, Rybka A, Sabol D A & Triplin K, Comparing tensorflow deep learning performance using CPUs, GPUs, local PCs and cloud. URL https://academicworks.s.cuny.edu/cgi/viewcontent.cgi?article=1054&context=bx_pubs.
- [3] Munhoz V, Bonfils A, Castro M & Mendizabal O. A performance comparison of HPC workloads on traditional and cloud-based HPC clusters. In: *2023 International Symposium on Computer Architecture and High Performance Computing Workshops (SBAC-PADW)*. IEEE. ISBN 979-8-3503-8160-3, pp. 108–114. <https://doi.org/10.1109/SBAC-PADW60351.2023.00026>.
- [4] Han J, Xu L, Rafique M M, Butt A R & Lim S H. A quantitative study of deep learning training on heterogeneous supercomputers. In: *2019 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE. ISBN 978-1-7281-4734-5, pp. 1–12. <https://doi.org/10.1109/CLUSTER.2019.8890993>.
- [5] Vaihela E. The performance and optimization of HPC clusters (2022). URL <https://urn.fi/URN:NBN:fi:amk-2022061017370>.

- [6] Penyala H, Ibrahim S & El Mesalami A. The raspberry pi education mine: For teaching engineering and computer science students concepts like, computer clusters, parallel computing, and distributed computing. In: *2020 IEEE International Conference on Electro Information Technology (EIT)*. IEEE. ISBN 978-1-7281-5317-9, pp. 624–628. <https://doi.org/10.1109/EIT48999.2020.9208242>.
- [7] Cicirello V A. Design, configuration, implementation, and performance of a simple 32 core raspberry pi cluster. <https://doi.org/10.48550/arXiv.1708.05264>.
- [8] Vargas-Pérez S. Designing an independent study to create HPC learning experiences for undergraduates. In: *2022 IEEE 29th International Conference on High Performance Computing, Data and Analytics Workshop (HiPCW)*. IEEE. ISBN 979-8-3503-3388-6, pp. 6–11. <https://doi.org/10.1109/hipcw57629.2022.00006>.
- [9] Moses Mwasaga N & Joy M. Implementing micro high performance computing (μ HPC) artifact: Affordable HPC facilities for academia. In: *2020 IEEE Frontiers in Education Conference (FIE)*. IEEE. ISBN 978-1-7281-8961-1, pp. 1–9. <https://doi.org/10.1109/FIE44824.2020.9273986>.
- [10] Mollova S, Zhekov M, Kostadinov A & Georgieva P. Laboratory model for research on computer cluster systems. In: *2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE. ISBN 978-953-233-095-3, pp. 1388–1393. <https://doi.org/10.23919/MIPRO.2018.8400250>.
- [11] Penyala H, Ibrahim S & El Mesalami A. The raspberry pi education mine: For teaching engineering and computer science students concepts like, computer clusters, parallel computing, and distributed computing. In: *2020 IEEE International Conference on Electro Information Technology (EIT)*. IEEE. ISBN 978-1-7281-5317-9, pp. 624–628. <https://doi.org/10.1109/EIT48999.2020.9208242>. URL <https://ieeexplore.ieee.org/document/9208242/>.
- [12] Alvarez L, Ayguade E & Mantovani F. Teaching HPC systems and parallel programming with small-scale clusters. In: *2018 IEEE/ACM Workshop on Education for High-Performance Computing (EduHPC)*. IEEE. ISBN 978-1-7281-0190-3, pp. 1–10. <https://doi.org/10.1109/EduHPC.2018.00004>.
- [13] Rao C M & Shyamala K. Analysis and implementation of a parallel computing cluster for solving computational problems in data analytics. In: *2020 5th International Conference on Computing, Communication and Security (ICCCS)*. IEEE. ISBN 978-1-7281-9180-5, pp. 1–5. <https://doi.org/10.1109/ICCCS49678.2020.9277362>.
- [14] Kumar D, Memon S & Thebo L A. Design, implementation & performance analysis of low cost high performance computing (HPC) clusters. In: *2018 12th International Conference on Signal Processing and Communication Systems (ICSPCS)*. IEEE. ISBN 978-1-5386-5602-0, pp. 1–6. <https://doi.org/10.1109/ICSPCS.2018.8631769>.
- [15] NVIDIA. *NVIDIA DGX OS 6 User Guide – Release Notes* (2024). URL https://docs.nvidia.com/dgx/dgx-os-6-user-guide/release_notes.html, accessed: 2025-03-25.
- [16] Emil V. Creation and optimization of HPC clusters. URL http://www.theseus.fi/bitstream/handle/10024/754650/Vaihela_Emil.pdf?sequence=2.
- [17] Majee A. DeepOps & SLURM: Your GPU cluster guide. <https://doi.org/10.48550/arXiv.2405.00030>.
- [18] NVIDIA, Deploying a scalable GPU-as-a-service platform and building a deep learning project in under 80 minutes | GTC digital march 2020 | NVIDIA on-demand. URL <https://www.nvidia.com/en-us/on-demand/session/gtcsj20-s22086/>.
- [19] Volkov A & Palmer B, Deploying rich cluster API on DGX for multi-user sharing. URL <https://developer.nvidia.com/blog/deploying-rich-cluster-api-on-dgx-for-multi-user-sharing/>, accessed: 2025-01-30.
- [20] Soyinka W. *Installing Rocky Linux 9 - Documentation*. Rocky Linux Project (2024). URL <https://docs.rockylinux.org/guides/installation/>, accessed: 2024-11-29.
- [21] Saive R, How to setup SSH passwordless login in Linux, Tecmint. URL <https://www.tecmint.com/ssh-passwordless-login-using-ssh-keygen-in-5-easy-steps/>, accessed: 2024-11-29.
- [22] Moffatt R, Linux cluster sysadmin – Parallel command execution with PDSH, Rittman Mead Blog. URL <https://www.ritmanmead.com/blog/2014/12/linux-cluster-sysadmin-parallel-command-execution-with-pdsh/>, accessed: 2025-01-30.
- [23] SergioMEV. Slurm for dummies: A dummy's guide to setting up (and using) HPC clusters on Ubuntu 22.04 LTS using slurm and munge. GitHub repository (2024). URL <https://github.com/SergioMEV/slurm-for-dummies>, accessed: 2024-11-29.
- [24] Rackspace. *Bind Mounts in Linux*. URL <https://docs.rackspace.com/docs/bind-mounts-in-linux>, accessed: 2024-11-29.
- [25] HowtoForge. *How to Set Up an NFS Mount on Rocky Linux 8*. URL <https://www.howtoforge.com/how-to-set-up-an-nfs-mount-on-rocky-linux-8/>, accessed: 2024-11-29.
- [26] Barone G B, Bottalico D, Carracciulo L, Doria A, Michelino D, Pardi S, Russo G, Sabella G & Spisso B. Designing and implementing a high-performance computing heterogeneous cluster. In: *2022 International Conference on Electrical, Computer and Energy Technologies (ICECET)*. IEEE. ISBN 978-1-6654-7087-2, pp. 1–6. <https://doi.org/10.1109/ICECET55527.2022.9872709>.
- [27] Kumar M. Distributed execution of dask on HPC: A case study. In: *2023 World Conference on Communication & Computing (WCONF)*. IEEE. ISBN 979-8-3503-1120-4, pp. 1–4. <https://doi.org/10.1109/WCONF58270.2023.10234994>.
- [28] Arisal A, Iryanto S B & Akbar Z. Managing multi-services for multi-users in heterogeneous cluster computing system. In: *2019 International Conference on Computer, Control, Informatics and its Applications (IC3INA)*. IEEE. ISBN 978-1-7281-5540-1, pp. 103–107. <https://doi.org/10.1109/IC3INA48034.2019.8949602>.
- [29] Manager S W. *Generic Resource (GRES) Scheduling*. Slurm SchedMD. URL <https://slurm.schedmd.com/gres.html>, accessed: 2024-11-29.

- [30] Documentation L. *How to Transition to Lmod (or How to Test Lmod Without Installing It for All)*. Lmod. URL https://lmod.readthedocs.io/en/latest/045_transition.html, accessed: 2024-11-29.
- [31] Analytics B D. *Modules, LMod, and EasyBuild*. URL <https://news-tutorials-bd-dl.readthedocs.io/en/latest/hands-on/easybuild/>, accessed: 2025-01-30.
- [32] Gamblin T, LeGendre M, Collette M R, Lee G L, Moody A, de Supinski B R & Futral S. The spack package manager: bringing order to HPC software chaos. In: *SC '15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC15. ACM, pp. 1–12. <https://doi.org/10.1145/2807591.2807623>.
- [33] Olmos J V, Liss L, Oved T, Binshtock Z & Goldenberg D. Advanced software architectures and technologies in high performance computing and data centers. In: *Optical Fiber Communication Conference (OFC) 2020*. Optica Publishing Group. ISBN 978-1-943580-71-2, p. T3K.3. <https://doi.org/10.1364/OFC.2020.T3K.3>.
- [34] McMillan S, How to run NGC deep learning containers with singularity, *NVIDIA Technical Blog*. URL <https://developer.nvidia.com/blog/how-to-run-ngc-deep-learning-containers-with-singularity/>.
- [35] GlobalPetrolPrices. Nepal electricity prices, March 2024. URL https://www.globalpetrolprices.com/Nepal/electricity_prices/, accessed: 2024-11-29.
- [36] Kleban C, NVIDIA Tesla T4 GPUs now available in beta. URL <https://cloud.google.com/blog/products/ai-machine-learning/nvidia-tesla-t4-gpus-now-available-in-beta>, accessed: 2024-11-29.