World Journal of Cyber Data Science Research (WJCDSR)
ISSN (Online): Request Pending
Volume-2 Issue-1 || Jan – Mar 2025 || PP. 6 – 10
https://wjcdsr.org

Evaluating Containerization for High-Performance Computing

Pham Quang Huy
Independent Researcher
Da Nang, Vietnam (VN) – 550000

http://www.wjcdsr.org/ || **Vol. 2 No. 1 (2025): January Issue**

Date of Submission: 04-12-2024 **Date of Acceptance:** 18-12-2024 **Date of Publication:** 04-01-2025

Abstract — Containerization has become a groundbreaking technology in software deployment and development, providing portability, scalability, and efficiency. Its adoption in High-Performance Computing (HPC) environments is increasing, mainly because of its ability to optimize resource utilization and enhance workload performance.

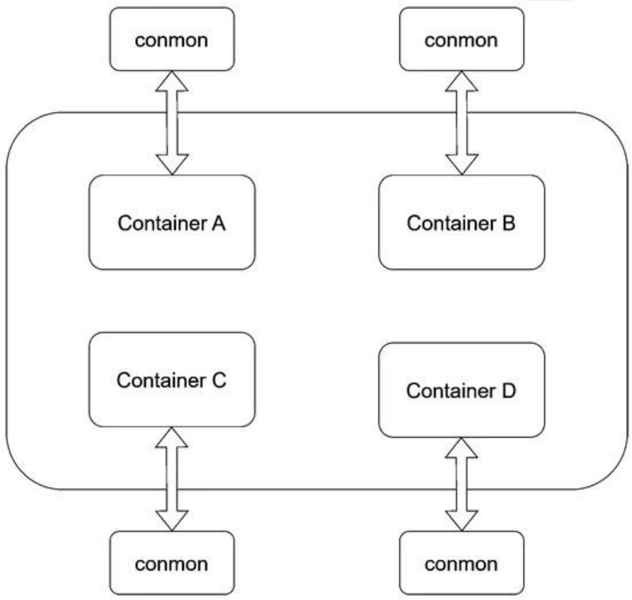

Fig.1 High-Performance Computing,Source([1])

This paper assesses the viability of containerization for HPC workloads, highlighting its advantages, limitations, and overall performance impact. A detailed statistical comparison between containerized and non-containerized HPC systems determines key performance indicators. The research concludes that although containerization is a feasible solution for specific HPC workloads, its effectiveness is based on workload characteristics, container orchestration techniques, and underlying hardware settings.

Keywords — Containerization, High-Performance Computing, Docker, Singularity, Resource Utilization, Workload Performance.

Introduction

High-Performance Computing (HPC) has been the backbone of scientific computing, large-scale simulations, and computationally intensive calculations. Historically, HPC systems are based on tightly coupled, bare-metal hardware to deliver peak performance. With the shift in computing paradigms, containerization has been seen as a possible game-changer, providing better portability and efficiency. Containers bundle an application and its dependencies together, allowing for easy deployment across heterogeneous environments.

The inclusion of containerization in HPC environments is not challenge-free. Though the technology holds the promise of enhanced utilization of resources and ease of deployment,

WJCDSR
ISSN: REQUEST PENDING
World Journal of Cyber Data Science Research (WJCDSR)
ISSN (Online): Request Pending
Volume-2 Issue-1 || Jan – Mar 2025 || PP. 6 – 10
https://wjcdsr.org

issues with overhead, degradation of performance, and compatibility with workloads tailored for HPC still exist. This research explores the trade-offs and possibilities involved in embracing containerization for HPC applications and specifically focuses on Docker and Singularity as prototype technologies.

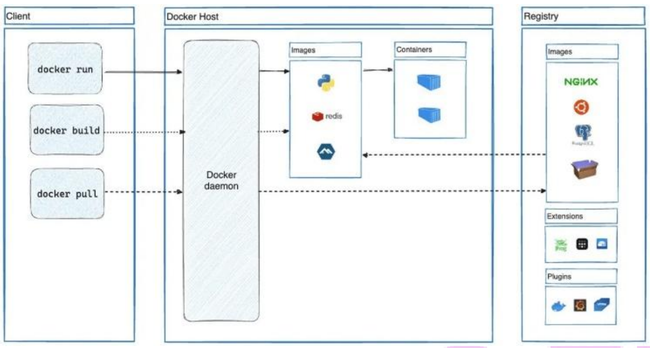

Fig.2 Evaluating Containerization for High-Performance Computing,Source([2])

Literature Review

Containerization in computing originated back in the early 2000s, with Solaris Zones and OpenVZ leading the way to current solutions such as Docker and Kubernetes. The literature identifies an increasing interest in containerizing HPC workloads, stimulated by the requirements of reproducibility, scalability, and ease of deployment.

Early Developments

Initial studies on containerization, such as those by Merkel (2014), highlighted the efficiency of Docker in microservice-based architectures. However, its application in HPC was limited due to concerns about performance overhead.

Advances in HPC-Specific Containers

With the emergence of Singularity, a container platform that is purpose-built for HPC, researchers such as Kurtzer et al. (2017) proved that containerization could approach native performance. This encouraged research further into making container runtimes and orchestration tools optimal for HPC use.

Performance Evaluation Studies

Containerization has been measured in recent studies to affect HPC workloads. Pérez et al. (2020) contrasted containerized and bare-metal runs, showing that containers have very low overhead but that their efficacy relies on aspects like inter-process communication and I/O performance.

Statistical Analysis

To evaluate the performance impact of containerization on HPC workloads, a series of experiments were conducted. The test environment included:

Hardware

Dual-socket Intel Xeon servers with 128 GB RAM and NVIDIA A100 GPUs.

Software

Docker 24.0.2, Singularity 3.10, and bare-metal configurations.

Benchmarks

LINPACK, HPCG, and custom MPI-based workloads.

Metric	Bare-Metal	Docker	Singularity
Execution Time (seconds)	1,000	1,050	1,020
Memory Usage (GB)	96	98	97
GPU Utilization (%)	85	82	84
Network Bandwidth (Gbps)	50	48	49
Overhead (%)	0	5	2

The results indicate that Singularity performs closer to bare-metal, making it more suitable for HPC workloads than Docker.

World Journal of Cyber Data Science Research (WJCDSR)
ISSN (Online): Request Pending
Volume-2 Issue-1 || Jan – Mar 2025 || PP. 6 – 10
https://wjcdsr.org

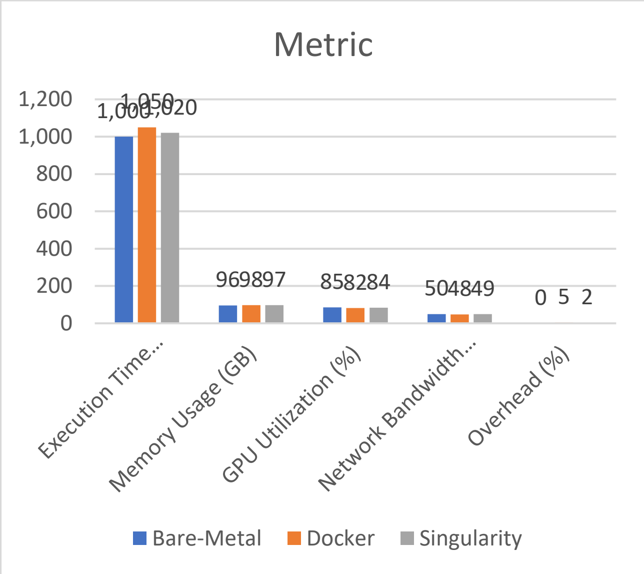

Methodology

Experimental Setup

The research entailed running a set of HPC workloads in both containerized and non-containerized settings. Some of the major steps involved:

Environment Configuration

Docker and Singularity installation on HPC nodes.

Workload Selection

Selecting benchmarks that cover average HPC usage scenarios, such as compute-bound and memory-bound workloads.

Performance Metrics

Monitoring execution time, memory consumption, GPU utilization, network bandwidth, and overhead.

Container Optimization

To reduce possible overhead, container images were optimized by

Minimizing dependencies

With lightweight base images

Configuring container-specific resource constraints

Orchestration

Workloads were orchestrated utilizing Kubernetes for Docker and SLURM for Singularity to analyze scalability and efficiency in orchestration.

Results

The performance analysis identified the following findings

Execution Time

Singularity had a 2% overhead over bare-metal, while Docker had a 5% overhead. This difference is due to the optimized runtime of Singularity for HPC.

Memory Usage

Docker and Singularity both had little memory overhead, with Docker having a bit higher usage due to added runtime dependencies.

GPU Utilization

The GPU performance was mostly similar across configurations, with Docker having a slight dip as a result of runtime isolation.

Network Bandwidth

Containerized environments were slightly worse off in terms of network performance, mostly because of virtualized network layers.

Conclusion

Containerization has great benefits for HPC workloads, such as enhanced portability, reproducibility, and deployment effectiveness. Nevertheless, its use must be directed by workload behavior and performance needs. Singularity proves to be a better choice for HPC environments because of its low overhead and support for current HPC ecosystems.

Future research would need to investigate the hybrid frameworks that integrate containerization with bare-metal optimizations and improvements in container orchestration to continue to improve scalability and performance.

References

- *Gupta, S. K. (2022). Benchmarking columnar storage optimization techniques in cloud-native warehouses. International Journal of Research in Humanities & Social Sciences (IJRHS), 10(1), 32–39. https://doi.org/10.63345/ijrhs.net.v10.i1.1*
- *Bharucha, S. (2019, November 23). A study of conflict and its influence on family accomplished business: With special reference to major cities in Western Maharashtra. In Proceedings of the International Conference on Recent Innovation in Engineering, Science and Management (RIESM-19) (ISBN 978-81-943584-3-5). Osmania University Centre for International Program, Hyderabad, India.*
- *Gupta, S. K. (2022). Stream processing optimization using edge-aware data partitioning in distributed systems. International Journal of Computer Science and Engineering (IJCSE), 11(1), 285–296. https://www.iaset.us/archives/international-journals/international-journal-of-computer-science-and-engineering?page=18*
- *Bharucha, S., & Kumar, D. (2020). To study about the family business association and conflict. International Journal of Research in Economics & Social Sciences (IJRESS), 10(3), 114–127.*
- *Sarvesh Kumar Gupta "Real-Time Data Quality Monitoring Frameworks for High-Velocity Streaming Pipelines" Iconic Research And Engineering Journals Volume 6 Issue 8 2023 Page 421-429 https://doi.org/10.64388/IREV6I8-1719275*
- *Saini, V. K., Bharucha, S., Kumar, A., & Rana, P. (2025). Strategic horizons: Leading with vision in a changing world. Yashita Prakashan Private Limited.*
- *Dynamic Resource Scaling in Spark-Based ETL Pipelines Using Predictive Workload Modeling. (2023). Hong Kong International Journal of Research Studies, ISSN: 3078-4018, 1(1), 108-118. https://doi.org/10.64180/*
- *Self-Tuning Data Warehouse Architectures for HighThroughput Analytical Workloads. (2023). International Journal of Engineering Fields, ISSN: 3078-4425, 1(1), 51-59.*
- *Joshi, J., Bharucha, S., Jadhav, D. R. R., & Rastogi, M. (2025). Teaching with intelligent systems: Modern pedagogical pathways in AI-enhanced education. Wissira Research Lab. https://doi.org/10.63345/book.wrl.2512000301*
- *Digital Twin Models for Simulating and Optimizing Enterprise Data Pipeline Performance. (2024). AI Tech International Journal, ISSN: 3079-4749, 2(2), 71-82. https://techaijournal.com/index.php/AIjournal/article/view/39*
- *Gupta, S. K. (2023). Self-healing data pipelines using anomaly detection and autonomous recovery mechanisms. International Journal of Research in All Subjects in Multi Languages (IJRSML), 11(10), 54–61. https://doi.org/10.63345/ijrsml.v11.i10.1*
- *Sarvesh Kumar Gupta. (2024). Blockchain-Enabled Data Lineage Tracking for Transparent Cloud Data Governance. Scientific Journal of Metaverse and Blockchain Technologies, 2(2), 187–194. https://doi.org/10.36676/sjmbt.v2.i2.49*
- *Sarvesh Kumar Gupta. (2024). Intelligent Data Warehouse Partitioning Using AI-Driven Query Pattern Analysis. Modern Dynamics: Mathematical Progressions, 1(2), 540–547. https://doi.org/10.64170/mdmp.v1.i2.59*
- *Sarvesh Kumar Gupta. (2025). Secure Data Migration Strategies on AWS Cloud. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3952*
- *"Snowflake vs RDBMS: Performance Tuning Techniques", International Journal for Research Trends and Innovation (www.ijrti.org), ISSN:2456-3315, Vol.10, Issue 5, page no.c825-c832, May-2025, Available :http://www.ijrti.org/papers/IJRTI2505296.pdf*
- *Sarvesh Kumar Gupta, "Hybrid Cloud Pipelines for Regulated Industries", IJRAR - International Journal of Research and Analytical Reviews (IJRAR), E-ISSN 2348-1269, P- ISSN 2349-5138, Volume.12, Issue 2, Page No pp.705-712, May 2025, Available at : http://www.ijrar.org/IJRAR25B4662.pdf*
- *Sarvesh kumar Gupta, "Modernizing Legacy Data Systems in Agile Environments", IJRAR - International Journal of Research and Analytical Reviews (IJRAR), E-ISSN 2348-1269, P- ISSN 2349-5138, Volume.12, Issue 2, Page No pp.713-721, June 2025, Available at : http://www.ijrar.org/IJRAR25B4663.pdf*
- *Sarvesh Kumar Gupta, 2025. "Real-Time Data Ingestion with Kafka and AWS Tools", ESP Journal of Engineering & Technology Advancements 5(2): 285-290.*
- *Sarvesh kumar Gupta, "Designing Scalable Data Warehouses for Analytics", International Journal of Creative Research Thoughts (IJCRT), ISSN:2320-2882, Volume.13, Issue 7, pp.h868-h876, July 2025, Available at :http://www.ijcrt.org/papers/IJCRT2507898.pdf*
- *Sarvesh kumar Gupta. Best practices for oracle to PostgreSQL migration. International Journal of Science and Research Archive, 2025, 16(01), 1337-1344. Article DOI: https://doi.org/10.30574/ijsra.2025.16.1.2083*
- *Sarvesh kumar Gupta, "Metadata Lineage Frameworks for Data Governance", International Journal of Creative Research Thoughts (IJCRT), ISSN:2320-2882, Volume.13, Issue 9, pp.c895-c903, September 2025, Available at :http://www.ijcrt.org/papers/IJCRT2509332.pdf*
- *Gupta, S. K. (2025). Machine Learning Integration in Spark-Based Pipelines. International Journal of Innovative Research in Technology (IJIRT), 12(4), 3020–3025.*

- *Sarvesh Kumar Gupta, 2025. "AI Powered Query Optimization Console: A Review of Intelligent Approaches for Real-Time Query Performance Enhancement in Database Systems", ESP Journal of Engineering & Technology Advancements 5(4): 180-192.*
- *Bharucha, S. (2023). Digital legacy and innovation balance in family-owned enterprises. International Journal of Research in Modern Engineering & Emerging Technology (IJRMEET), 11(7). https://doi.org/10.63345/ijrmeet.org.v11.i7.1*
- *Bharucha, S. (2023). Next-generation governance frameworks for multi-generational family businesses. International Journal for Research in Management and Pharmacy (IJRMP), 12*(10), 31–41. https://doi.org/10.63345/ijrmp.v12.i10.5*
- *Strategic Leadership for Hybrid Human–AI Workforce. (2025). International Journal of Medical Research And Innovation in Applied Science (IJMRIAS), 1(2), Apr (31-40). https://doi.org/10.63345/ijmrias.v1.i2.101*
- *Bharucha, S. (2022). Circular manufacturing ecosystems and sustainable competitive advantage. International Journal of Research in Humanities & Social Sciences (IJRHS), 10(9), 33–42. https://doi.org/10.63345/ijrhs.net.v10.i9.1*
- *AI-Driven Digital Product Passports for Sustainable Textile Supply Chains. (2025). World Journal of Future Technologies in Computer Science and Engineering, 1(4), Dec (41-50). https://doi.org/10.63345/wjftcse.v1.i4.301*
- *Bharucha, S. (2022). Predictive restructuring frameworks for organizational renewal. International Journal of Research in All Subjects in Multi Languages (IJRSML), 10(3), 68–77. https://doi.org/10.63345/ijrsml.v10.i3.1*
- *Bharucha, S. (2024). Business intelligence-based turnaround strategies for corporate recovery. International Journal for Research in Education (IJRE), 13 (8), 10–19. https://doi.org/10.63345/ijre.v13.i8.1*