ISSN: 2754-6659

Journal of Artificial Intelligence & Cloud Computing

SCIENTIFIC Research and Community

Review Article

Open Access

Optimizing Oracle DB Workloads for Enterprise SAN Environments using NVMe/FC on Cisco UCS and NetApp Storage Converged Infrastructure

Sriramaraju Sagi

NetApp, USA

ABSTRACT

In this research paper the authors delve into the utilization of NVMe/FC, on Cisco UCS and NetApp Storage converged infrastructure to enhance the performance of Oracle database workloads in enterprise SAN systems. They highlight the advantages of employing NVMe oF, a rapid storage standard that facilitates communication with cloud-based solutions. By leveraging existing Ethernet or InfiniBand infrastructure NVMe oF offers throughput and Improved IOPS. The authors also discuss the significance of hardware calibration utilizing FIO, IOPS tests on the x410c M7 Server, and SLOB (silly little Oracle Benchmark) user scalability tests for total IOPS. Furthermore, they emphasize the cost effectiveness and scalability benefits of transitioning from hardware to commodity hardware running Linux. The article concludes by exploring how localized infrastructure provides organizations, with enhanced adaptability to meet their demands and requirements thereby boosting production efficiency.

***Corresponding author**

Sriramaraju Sagi, NetApp, USA.

Received: September 07, 2023; **Accepted:** September 16, 2023; **Published:** September 25, 2023

Keywords: NVMe-oF, Oracle, NetApp, Cisco UCS, Enterprise SAN, Scale-Out Architecture

Introduction

Optimizing technology, in enterprise workloads plays a role in managing resources and achieving satisfactory business outcomes. It also ensures service delivery to clients especially when facing increases in workload from business applications. By reconfiguring the existing infrastructure and streamlining workflow processes we can reduce core traffic leading to performance and efficient utilization of available resources. In the case of enterprise workloads like Oracle DB optimizing technology becomes essential for enhancing performance, resource allocation and service delivery. To handle spikes in traffic caused by business applications it is vital to optimize the buffer size in Oracle DB by increasing the buffer cache. This helps reduce disk I/O operations and enhances database performance while maintaining data locality and minimizing disk accesses.

In times companies have been shifting away from proprietary hardware, towards commodity hardware running on Linux. This transition offers cost savings and scalability benefits. By leveraging 18 core Intel chips that're readily available companies can operate large Linux compute nodes at a fraction of the cost compared to traditional enterprise UNIX servers.

Moreover Linux has demonstrated its credibility as an operating system that offers reliability, to UNIX based systems like IBM AIX, HP/UX and Sun Solaris. Through optimizations Linux can efficiently handle Oracle databases. Deliver enhanced performance making it an excellent option, for enterprise workloads.

Ultimately when firms optimize technology and adopt commodity hardware running on Linux they can experience cost savings and improved scalability. At the time they can maintain the performance of UNIX servers. By making these changes organizations can effectively manage resources and ensure service delivery, to clients ultimately leading to business performance.

In a context localized infrastructure refers to the use of technology and systems specifically designed to meet the requirements of a particular business or organization. This approach contrasts with an one. Implementing localized infrastructure allows businesses to adapt effectively to their needs and requirements resulting in increased efficiency and productivity. One of the advantages of localized infrastructure is its ability to scale easily as businesses grow or change their requirements. For example if a sudden increase in demand occurs a business can quickly ramp up its infrastructure to meet that demand. Overall localized infrastructure offers benefits for businesses such as enhanced scalability and the ability to adapt rapidly to technologies or market demands. By having an customized infrastructure, in place businesses can meet their needs more effectively and remain competitive in todays fast paced business environment.

One way to achieve localized integration is, by utilizing NVMe/FC (Non Volatile Memory Express over Fibre Channel) a storage protocol that facilitates communication between traditional systems and modern cloud based solutions. By incorporating NVMe/FC organizations can take advantage of the benefits offered by cloud computing while ensuring compatibility with their existing infrastructure. This technology enhances performance reduces latency and improves scalability resulting in data transfer and overall system efficiency. Furthermore NVMe/FC

incorporates security measures like end, to end data protection and encryption to guarantee the transmission of sensitive information. For instance a financial institution relying on legacy mainframe systems can integrate NVMe/FC to connect with a cloud based analytics platform. This integration allows them to harness the scalability and processing power of the cloud while seamlessly integrating with their infrastructure. The improved performance and decreased latency provided by NVMe/FC enable real time data analysis enabling the institution to make more decisions regarding investments, risk management and customer services. Moreover the advanced security features offered by NVMe/FC provide a layer of safeguarding for financial data.

With the implementation of end, to end encryption and robust protocols for data transfer NVMe/FC ensures that information remains protected during transit well as when it is stored. This provides reassurance, to both institutions and their customers guaranteeing the handling of their financial data. Moreover the incorporation of NVMe/FC technology enables storage and retrieval of data thereby reducing the likelihood of data loss or corruption. Ultimately embracing NVMe/FC technology empowers financial institutions to modernize their operations and remain competitive amidst the evolving landscape.

In businesses Oracle DB (database) is commonly used for purposes. It helps manage and store amounts of data, process transactions generate reports, support business intelligence and analytics well as facilitate e commerce activities. Oracle DB is a solution that caters to diverse business needs. To ensure performance, in enterprise workloads it's important to focus on areas such as hardware infrastructure, storage systems, network configuration and database tuning. By optimizing these aspects enterprises can ensure that Oracle DB operates efficiently and delivers the possible performance.

One crucial aspect of implementing Oracle DB in enterprises is the consolidation process. This process has the potential to simplify and streamline deployment and maintenance procedures significantly. It reduces complexity while enhancing efficiency. By utilizing an infrastructure IT teams can easily scale resources up or down as needed without experiencing downtime. This approach minimizes disruptions. Maximizes productivity levels. Additionally Cisco and NetApp provide integrated management tools within the converged infrastructure that simplify monitoring and troubleshooting processes. This integration enables issue resolution while relieving some of the burdens, on IT staff members.

Converged infrastructure is a platform that brings together compute, storage, networking and virtualization resources to consolidate and simplify management. This results, in improved efficiency and cost reduction for institutions. In contrast traditional IT infrastructure involves components that are managed independently leading to complexity and inefficiency. By opting for Cisco UCS and NetApp as the hardware and software solutions businesses can leverage converged infrastructures advantages to enhance security measures streamline operations and create an agile and scalable IT environment. Converged infrastructure addresses scalability challenges by offering a platform that can easily accommodate growing database needs. It also improves performance by eliminating bottlenecks caused by components while optimizing resource allocation. Additionally it enhances security measures by centralizing management and implementing security protocols across all components to ensure data protection and compliance with requirements.

The utilization of Ciscos converged infrastructure from Cisco along with NetApp greatly simplifies database management by integrating servers, storage, networking and virtualization into one platform. This consolidation process has the potential to significantly reduce complexity while enhancing efficiency through deployment and maintenance procedures. With converged infrastructure, in place IT teams can seamlessly scale resources up or down as needed minimizing downtime while maximizing productivity.

Moreover the converged infrastructure provides a combination of management tools, from Cisco and NetApp. This integration allows for monitoring and troubleshooting processes resulting in issue resolution and reducing the workload, on IT personnel.

Literature Review

Large corporations face obstacles and restrictions when it comes to optimizing Oracle DB workloads. These challenges encompass the necessity, for data storage capabilities scalability to handle amounts of data concerns about security and cost effectiveness. One potential solution that companies can explore to enhance their data storage capabilities is utilizing Cisco UCS and NetApp solutions. These solutions offer advantages over competing options. Firstly Cisco UCS provides a computing platform that combines compute, storage and networking resources enabling simplified management and improved efficiency. NetApps data storage solutions provide features like data deduplication and compression which optimize storage utilization and reduce costs.

Looking ahead to the future there are emerging trends in the industry that are anticipated to drive the adoption of NVMe/FC technology powered by Cisco UCS and NetApp solutions further. Real time analytics and decision making are becoming increasingly in demand, as one of these trends. NVMe/FC technology offers low latency and high speed data transfer capabilities that can support real time analytics applications effectively. Furthermore considering the growing amount of data generated in the world scalable storage solutions are essential. The NVMe/FC technology powered by NetApp can deliver the scalability required to handle volumes of data.

Several studies have explored the implementation of NVMe/FC, which combines the NVMe protocol with fiber channel technology. Sanyal emphasized the importance of validating the functionality of NVMe hosts to ensure they meet all criteria [5]. Numerous research projects have investigated the performance of all flash array systems based on NVMe SSDs. Jin highlighted that inefficient kernel policies and settings can contribute to fluctuating I/O latency in these systems suggesting reconfiguration, as a solution for achieving more stable latency. The key findings from this study include examining and improving input/output latency in a petabyte scale AFA system identifying kernel policies as the cause of significant I/O variations and emphasizing the need for enhanced CPU schedulers and storage maintenance protocols to minimize prolonged tail latency. Son and Son both conducted assessments of NVMe SSDs with one focusing on performance measurements and workloads while the other concentrated on I/O setups [1]. Both studies concluded that NVMe SSDs outperformed SATA SSDs in terms of data transfer rate and response time. Xu presented an analysis of NVMe drive performance, including a breakdown of access times and a comparison between their capabilities, in cloud databases versus those using SATA based SSDs [2].

These investigations highlight the potential, for performance in AFA systems by using NVMe SSDs and optimizing system software and configurations. The NVMe specification was developed by industry and academic experts to achieve performance on NVM devices. While this technology is widely used in applications it's important to evaluate the performance of NVM devices with NVMe specifications under I/O configurations and analyze workloads on storage systems to maximize their performance potential. This paper showcases the architecture that utilizes Cisco UCS and NetApp technologies demonstrating the benefits of operating an Oracle RAC Database 21c environment with enhanced performance, expandability and robust availability through the use of NVMe, over Fibre Channel (NVMe/FC).

Methodology

In this study we utilized an architecture that combined Cisco UCS and NetApp technologies to evaluate the performance of an Oracle RAC Database 21c environment. Our successful installation of the Oracle RAC database 21c, on a converged infrastructure with Cisco UCS and NetApp storage enabled us to conduct performance tests using NVMe over Fiber Channel (NVMe/FC) as the storage protocol. The objective was to assess the advantages of using NVMe over Fiber Channel (NVMe/FC) in an Oracle RAC Database 21c environment in terms of performance, expandability and availability. The outcomes revealed enhancements in these aspects when implementing NVMe over Fibre Channel (NVMe/FC) as the storage protocol.

The research paper details our setup for testing the deployment and scalability of an Oracle database on a FlexPod system. This involved integrating technologies such as NVMe/FC, Cisco UCS and NetApp storage solutions. By employing NVMe over Fibre Channel (NVMe/FC) we aimed to optimize storage processes, for speed and efficiency.

NVMe, a storage protocol designed to take advantage of the speeds provided by solid state drives in a storage network is used over fiber channel networks to enable efficient data transfer, between computing resources and storage arrays. This implementation greatly reduces delays. Improves the performance of applications that handle amounts of data like Oracle databases.

Cisco UCS, also known as the Unified Computing System is a platform, for data centers that brings together computing, networking, storage access and virtualization. Its main purpose is to minimize ownership costs and enhance scalability. Within this framework Cisco UCS acts as the infrastructure for computing. It offers an architecture through blade servers. Rack servers that can be managed from a centralized location. These servers are specifically optimized for performance to effectively handle enterprise application workloads such as hosting Oracle databases.

NetApp storage systems like the AFF A900 array mentioned in the document are used to provide the storage framework for Oracle databases. These arrays are well known for their performance in environments where demanding IOPS (input/output operations per second) and low latency requirements are crucial for database operations. This storage setup optimizes the speed and effectiveness of SSDs while prioritizing reliability and data protection features to ensure fast and secure database storage.

The Oracle databases are set up on the configured UCS servers using Oracle Real Application Clusters to achieve both availability and scalability. NetApp arrays serve as the storage solution, for these databases with data files distributed across SSDs to maximize performance. The document describes evaluations conducted to assess the effectiveness of this configuration.

These evaluations involve calibrating the hardware using FIO conducting IOPS tests on the x410c M7 Server and performing user scalability tests using SLOB (Oracle Benchmark) to measure total IOPS. We carefully analyze the performance measurements to ensure that the setup meets all the requirements, for a level corporate database system.

In summary this document presents an setup and configuration aimed at creating an efficient, flexible and reliable infrastructure, for Oracle databases through the use of NVMe/FC technology, Cisco UCS and NetApp storage solutions. The system underwent a series of tests to validate its performance and scalability in order to meet the demanding needs of enterprise applications.

Results

Flexible IO plays a role, in designing and optimizing enterprise workloads for Oracle databases. In this scenario FIO is used to assess the performance of a storage device within a time frame. To conduct the FIO tests we created eight subsystems with a total of 32 namespaces (each containing four namespaces). Each subsystem had a size of 500 GB, evenly distributed across both aggregates. These 32 namespaces were utilized across all four nodes for read/write IO operations. We conducted FIO tests to measure IOPS, latency and throughput performance by adjusting the block size parameter within each test. Additionally we modified the read/write ratio for each test including scenarios with read/write ratios of 0/100%, 50/50%, 70/30% and 90/10%. The purpose was to evaluate system performance at scales. To ensure the sustainability of this configuration under loads, over a period each test ran for at least four hours.

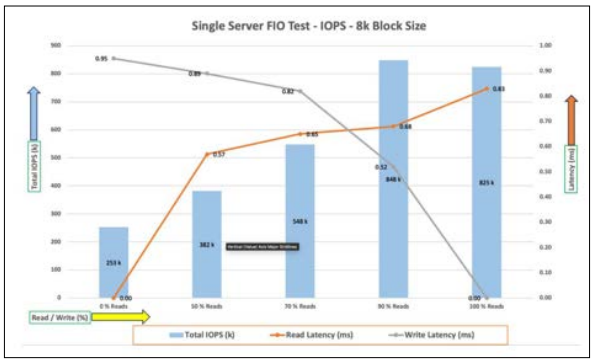

Figure 1: Single Server FIO Test, IOPS, 8K Block Size

In our tests, on a single server node we found that the average IOPS for 100% read/write was 825k, with a latency of than 1 millisecond. Similarly in the 90/10% read/write test we achieved 848k IOPS with both read and write latencies under 1 millisecond. When it came to the 70/30% read/write test we achieved around 548k IOPS with latencies. The results for the 50/50% and 0/100% read/write tests were 382k IOPS and 253k IOPS respectively both with write and read latencies under the threshold.

Figure 2: 4 Server FIO Test, IOPS, 8K Block Size

Across all four server nodes in our tests using reads/writes of size 8k we observed an average of IOPS for the test where reads constituted100% of operations and writes constituted none. The latency, for this test remained millisecond well. In the case of the 90/10%,70/30%,50/50% and0/100%read/write tests we achieved approximately 1645k IOPS,792k IOPS , 513 k IOPS, and 364 k IOPs respectively with similar latency results.

The Silly Little Oracle Benchmark (SLOB) is a toolkit designed to generate and test input/output (I/O) operations in an Oracle database. It effectively assesses the performance of the I/O subsystem using Oracle SGA buffered physical I/O. We configured SLOB to run on all four nodes of the Oracle Real Application Clusters (RAC) with a distribution of users, across each node. During our testing we gradually increased the number of Oracle users in the database starting from a minimum of 128 users and reaching a maximum of 512 users across all nodes. At each level of user load we thoroughly examined whether both the storage system and server nodes could maintain performance without encountering any issues or bottlenecks. The graphs below depict the number of input/output operations per second (IOPs) observed while executing SLOB workloads for different numbers of concurrent users. These graphs demonstrate scalability as we increase the number of users showcasing IOPS results, for scenarios where read/write ratios are set at 100%, 90%, 70% and 50%.

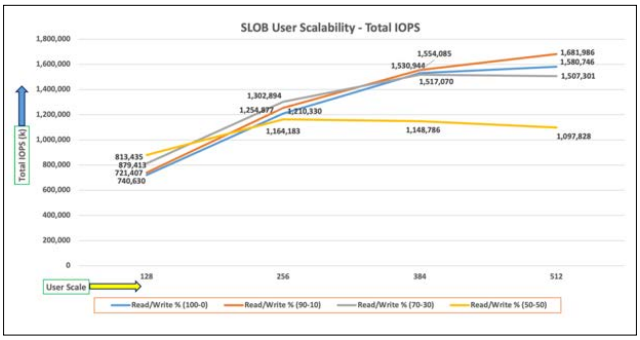

Figure 3: SLOB User Scalabiity: Total IOPS

Conclusion

In our research we delved into the optimization of Oracle Database (DB) workloads, with a focus, on implementing NVMe/FC (Non Volatile Memory Express over Fibre Channel) technology on Cisco UCS (Unified Computing System) and NetApp infrastructure. The article highlights how customized infrastructure brings advantages to businesses by allowing them to tailor their technology and systems according to their needs. This enables scalability, adaptation to technologies and responsiveness to market demands.

We introduce NVMe/FC technology as a high speed storage protocol that bridges the gap between legacy systems and modern cloud based solutions. It offers performance, latency improved scalability and advanced security features that are well suited for real time analytics and data intensive applications. To improve Oracle DB workloads it is crucial to focus on aspects such as hardware infrastructure, storage systems, network configuration and database tuning strategies like converged infrastructure solutions, from Cisco UCS and NetApp. These solutions streamline deployment processes enhance resource scalability and simplify management tasks.

Implementing NVMe/FC technology in an Oracle RAC Database 21c environment, on a converged infrastructure data center with Cisco UCS and NetApp storage leads to performance improvements, expandability and availability as per the findings of the study. Evaluations using FIO and SLOB demonstrate that Flexible IO tests exhibit IOPS, latency and throughput across different read/write ratios and block sizes. The Silly Little Oracle Benchmark tests show that IOPS scale linearly with increased users indicating performance in various scenarios. To efficiently manage resources and ensure service delivery in enterprise environments it is crucial to optimize Oracle DB workloads. By adopting NVMe/FC technology, on Cisco UCS and NetApp infrastructure organizations can achieve performance, scalability and security. To enhance efficiency and competitiveness in the changing landscape organizations should consider implementing converged infrastructure solutions for optimizing their Oracle DB workloads [3-9].

Acknowledgments

This article was supported by Tushar, Cisco, Hardik, Cisco and Bobby Oommen, NetApp – we sincerely thank them for their contribution and valuable input.

References

1. Yongseok S, Hara K, Hyuck H, Yeom H (2016) An Empirical Evaluation of NVM Express SSD. Cluster Computing 19.
2. Qiumin X, Huzefa S, Mrinmoy G, Tameesh S, Awasthi M, et al. (2015) Performance analysis of NVMe SSDs and their implications for real-world databases. Experimental Systems Conference https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=1a0af07c26d30548f2bd40c769f3961547a78179.
3. Dino Q, Thomas A, Graeme C,Amr D, Guilherme F, et al. (2013) Workload-optimized systems tuning POWER7 for analytics. IBM Redbooks https://www.redbooks.ibm.com/abstracts/sg248057.html.
4. Mir M, Ali MN, Sami J, Ansari U (2014) Review of mechanics and applications of auxetic structures. Adv Mater Sci Eng https://www.hindawi.com/journals/amse/2014/753496/.
5. Baoning N, Xue Y, Patrick M, Powley W (2010) Managing Workload Importance in Enterprise DBMSs. Sematic Scholar https://www.semanticscholar.org/paper/Managing-Workload-Importance-in-Enterprise-DBMSs-Niu-Xue/61af2a2f3dd50275f6a214e44863056bcb3ecdf0.
6. Keke G, Zihua D, Meikang Q, Hui Z (2015) Efficiency-aware Workload Optimizations of Heterogeneous Cloud Computing for Capacity Planning in the Financial Industry. IEEE 2nd International Conference on Cyber Security and Cloud ComputingAt: New York City, NY, USA https://www.researchgate.net/publication/283661641_Efficiency-Aware_Workload_Optimizations_of_Heterogeneous_Cloud_Computing_for_Capacity_Planning_in_Financial_Industry.

7. Bishwapa S, Namita P (2021) Verification Of Firmware-Controlled NVME Host. International Research Journal of Engineering and Technology (IRJET) https://www.semanticscholar.org/paper/VERIFICATION-OF-FIRMWARE-CONTROLLED-NVME-HOST-Sanyal-Palecha/5d37552d964d27c46f7fece435cdea30a8984914.
8. Youngbin J, Ahn S, Sungjin L (2018) Performance Analysis of NVMe SSD-Based All-Flash Array Systems. IEEE International Symposium https://www.semanticscholar.org/paper/Performance-Analysis-of-NVMe-SSD-Based-All-flash-Jin-Ahn/083146ec3f66cf21ee692d475f3b8f0baf77d445.
9. Yongseok S, Hara K, Hyuck H, Heon YY (2016) An empirical evaluation and analysis of the performance of the NVM Express solid state drive. Cluster Computing https://www.researchgate.net/publication/305268917_An_empirical_evaluation_and_analysis_of_the_performance_of_NVM_express_solid_state_drive.