



CASE STUDY

Hindustan Petroleum Corporation Limited Adds Turbo Power to Their OLTP and OLAP Transactions



Solution Focus

- OLTP and OLAP performance

Challenges

- Legacy storage systems could no longer support HPCL's expanding workloads
- Difficulty in adding more users to overloaded system
- Subpar performance and limited scalability had a major negative impact on business

Solution

- PCIe 3.2TB & 1.2TB as a physical solution

Key Results

- Performance increase—at least 10x greater

Summary

Hindustan Petroleum Corporation Limited (HPCL) recognized a growing need to support faster response times on their legacy storage system for OLAP/OLTP transactions. By deploying 3.2TB¹ and 1.3TB PCIe Application Accelerators, HPCL systems were able to decrease latency and achieve higher IOPS.

Background

HPCL is a Government of India Enterprise with a Navratna status, and a Forbes 2000 and Global Fortune 500 company incorporated as a company under the Indian Companies Act of 1913. HPCL is listed on the Bombay Stock Exchange (BSE) and National Stock Exchange (NSE), India.

HPCL owns and operates two major refineries producing a wide variety of petroleum fuels and specialties, one in Mumbai (west coast) of 6.5 million metric tonnes per annum (MMTPA) capacity and the other in Visakhapatnam, (east coast) with a capacity of 8.3 MMTPA. HPCL also owns and operates the largest lube refinery in the country producing lube base oils of international standards, with a capacity of 428 TMT. This lube refinery accounts for over 40% of India's total lube base oil production. Presently HPCL produces over 300+ grades of lubes, specialties and greases. HPCL in collaboration with Mittal Energy Investments Pte. Ltd. is operating a 9 MMTPA capacity refinery at Bathinda with 49% equity in Punjab and also holds an equity of about 16.95% in the 15 MMTPA Mangalore Refinery and Petrochemicals Ltd. (MRPL).

HPCL has the second largest share of product pipelines in India with a pipeline network of more than 3015 kms for transportation of petroleum products and a vast marketing network consisting of 13 zonal offices in major cities and 106 regional offices facilitated by a supply and distribution infrastructure comprising terminals, pipeline networks, aviation service stations, liquefied petroleum gas (LPG) bottling plants, inland relay depots and retail outlets, lube and LPG distributorships.

Consistent excellent performance has been made possible by a highly motivated workforce of over 11,000 employees working all over India at its various refining and marketing locations.

"Post implementation of the [Western Digital] PCIe solution the average response time for both read and write has gone down to below 1ms and the average storage performance of the application has improved from 3x to 10x depending on workload."

Hemant Gote
Chief Manager—Data Centre, HPCL

Background of Application

HPCL has been using legacy storage systems (OLTP workloads) for booking of LPG cylinders by a large network of their distributors across India. Critical requirement and key consideration from HPCL for their OLTP transactions were very high read performance, high speed I/O response along with a requirement of completing concurrent transactions at a short time to improve user experience. Speed and efficiency of batch processing workloads at a certain time of the day was also critical to the solution. Response time in legacy solutions ranged between 15ms to 300ms per transaction. The expected response time was 10ms. The average response time achieved post deployment of the PCIe solution is less than 1m.

Solution

HPCL had previously deployed an enterprise storage solution with a caching layer—SQL Server™ installed on a virtual machine (VM) and the VM was ultimately talking to the storage layer.

The Western Digital PCIe solution was deployed as a physical solution to the SQL Server environment.

Architecture consists of the following database servers:

1. Node-4: Physical (bare metal) server with PCIe cards with production SQL instance
2. Node-1 and 2: Virtual servers which make SQL failover cluster of two nodes
3. Node-3: Virtual server which hosting standalone SQL instance. It also has Read caching using PCIe cards

All 4 nodes are part of Windows® failover cluster at production site.

The business requirement was of a storage system which required handling high OLTP operations beside read operations with minimum latency (in microseconds/less than 10ms).

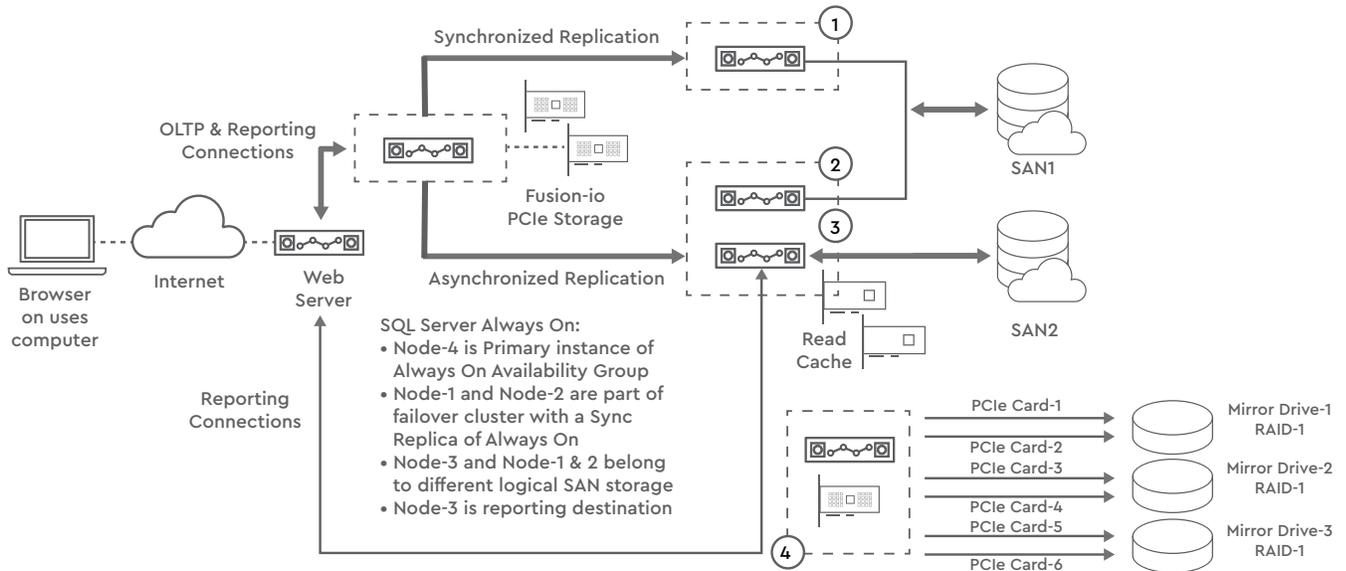
Journey

The initial plan was to create Node-4 in virtual environment using PCIe cards as pass-through. But six cards each of capacity 3.2TB were found to be incompatible under pass-through for attaching in single virtual server Node-4. Subsequently cards were introduced as VM disk with software RAID-1 in virtual server but I/O response was very poor, sometimes going to seconds.

Finally, all six cards were added in physical (bare metal) server to overcome virtual environment hurdles. To ensure redundancy, two cards were put under mirroring RAID-1. So installed capacity was close to 18TB. Effective usable capacity: 9TB. Performance of PCIe cards under physical environment was excellent. Average read and write I/O was less than 10 milliseconds (sometimes in microseconds).

At present, Node-4 is our production database server with Node-1 and 2 as synchronous replica and Node-3 as asynchronous replica.

Block Diagram of Western Digital PCIe Card Implementation



Western Digital Solution Advantage: Increased Performance, Lower Power Consumption, Smaller Footprint

According to Hemant Gote, Chief Manager – Data Centre, HPCL, “Post implementation of the [Western Digital] solution the average response time for both read and write has gone down to below 1ms and the average storage performance of the application has improved from 3x to 10x depending on workload”.

Because of the successful results from the PCIe deployment, HPCL is now considering the possibility of using a similar solution in its data centres.

Special Note

Even greater application performance enhancement is now available with the latest Ultrastar DC SN200 NVMe™ PCIe SSDs. Compared to the prior generation PCIe application accelerator, the Ultrastar DC SN200 NVMe SSDs can deliver an additional 4x performance for read-heavy workloads like reporting and 2x performance for mixed read-write workloads like OLTP. The Ultrastar DC SN200 contains its own CPU and memory for flash management, so no server overhead is needed. NVMe support eliminates the need for custom drivers and simplifies deployment and maintenance. In addition, the Ultrastar DC SN200 is offered in two form factors, a HH-HL model, and 2.5" U.2 drive, delivering capacities up to 7.68TB.

To learn more about the Ultrastar DC SN200 NVMe PCIe SSDs, visit www.westerndigital.com.

Western Digital.

5601 Great Oaks Parkway
 San Jose, CA 95119, USA
US (Toll-Free): 800.801.4618
International: 408.717.6000

www.westerndigital.com

©2018 Western Digital Corporation or its affiliates. All rights reserved. Western Digital, the Western Digital logo, and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. The NVMe word mark is a trademark of NVM Express, Inc. SQL Server and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other marks are the property of the respective owners.

¹ One megabyte (MB) is equal to one million bytes, one gigabyte (GB) is equal to 1,000MB (one billion bytes), and one terabyte (TB) is equal to 1,000GB (one trillion bytes) when referring to storage capacity. Accessible capacity will vary from the stated capacity due to formatting, system software, and other factors.