SOLUTION BRIEF

ORACLE | Partner

PLiOPS EXTREME DATA PROCESSOR

High Density Oracle Database Consolidation Solution

EXECUTIVE SUMMARY

Oracle Database is known for its exceptional performance, scalability, availability, and security. With Oracle Database 21c release, this converged database offers customers best of breed support for data types including relational, JSON, XML, spatial, graph, OLAP, etc. This means that customers can efficiently handle different data formats within a single database. It simplifies database consolidation, allowing organizations to bring together diverse workloads while maintaining isolation and helping to ensure that each workload receives the necessary resources and performance. Oracle Database 21c aims to provide a comprehensive and efficient solution as multi-model, multi-tenant, multi-workload database platform.

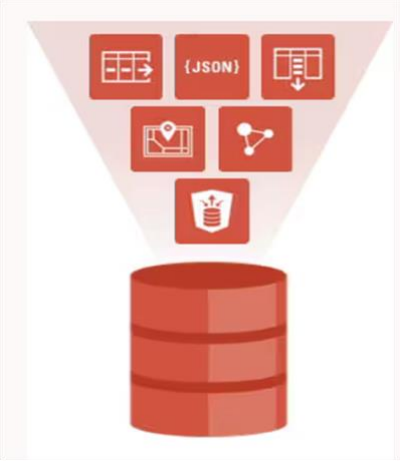

Figure1: Oracle Database 21c converged database.

Pliops Extreme Data Processor (XDP) is a breakthrough data accelerator that delivers performance acceleration, capacity expansion and reliability benefits that can increase the density of pluggable databases with Oracle Database 21c. It also simplifies the set-up, configuration, and administration of infrastructure for Oracle Database, enabling a faster refresh of earlier versions of Oracle Database infrastructure. This solution brief focuses on the combined benefits of Pliops XDP and Oracle Database 21c for performance acceleration, multi-tenancy, database backup offloading, and restore operations.

Performance Scaling with Pliops

As shown in Figure 2, HammerDB4.7 was employed to evaluate the combined benefits of Pliops setup with Oracle Database 21c for multi-tenant deployment. Oracle Database 21c with container and four pluggable database testing environment is configured on Dell PowerEdge R750 server using the Oracle Linux 8.6 operating system with the Unbreakable Enterprise Kernel (UEK) 5.4.17-2136.307.3.1.el8uek.x86_64 version. The servers are set up with Intel 8368 38-Core processor and 1Tera Byte of Memory and storage is configured with 4 Gen4 3.84TB NVMe SSDs (Samsung PM9A3) that are directly managed by Pliops XDP.

Pliops–Oracle Solution Benefits

- Increase the consolidation density of Pluggable Databases by 50%.
- Easy to provision and move the Oracle Database 12c,18c workloads to a new Oracle database 21c solution.
- 66% higher aggregate performance with PDBs
- Increased redundancy with external storage array benefits at significantly lower costs
- Close to 3X additional database capacity gains with Oracle ASM plus XDP external redundancy for ASM.
- Faster database backup and recovery benefits to meet service level objectives using HW offload benefits.
- Higher reliability and resiliency of Oracle database to tolerate multiple SSD failures.
- Enables to reduce the SSD $/TB cost with capacity and Endurance benefits.

Oracle Database 21c: HammerDB Performance

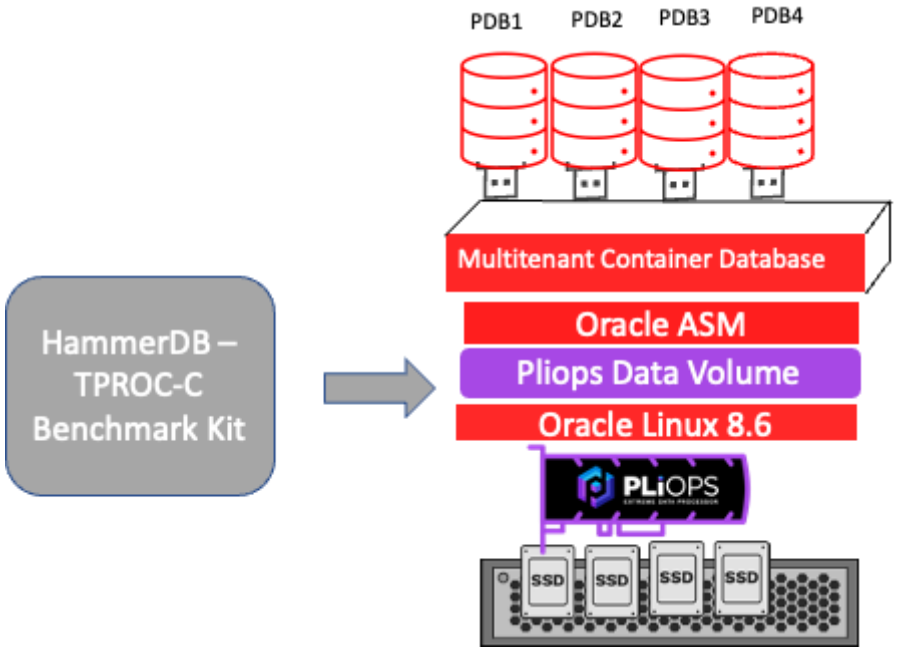

Figure2: Pliops XDP with Oracle Database 21c

To identify multi-tenancy database capacity and performance of Pliops XDP, Oracle Automatic Storage Management (ASM) is initially configured with internal redundancy with NVMe SSDs. Later, Oracle ASM is configured with external redundancy with Pliops providing the data protection for NVMe SSDs. Pliops XDP exposes a high-performance, highly reliable data storage volumes with built-in data protection to tolerate SSD failures. To execute an OLTP workload, each PDB is loaded with 1TB of dataset that is HammerDB TPROC-C benchmark-configured with 10,000 warehouses. Our benchmark testing included both Use All-Warehouses option enabled and disabled to evaluate benefits in both customer scenarios. Enabling the Use All-Warehouses option will increase the I/O load on the storage system significantly by accessing all 10K warehouses and simulating highly random access to the dataset.

Oracle Database Capacity Savings:

Query Result

SQL | All Rows Fetched: 2 in 0.022 seconds

	NAME	TOTAL_MB	USABLE_TOTAL_MB	FREE_MB	FREE_PCT	USED_MB	USED_PCT	USED (%)
1	ASM_DISKGROUP	14651312	5494242	615404	11	4878838	89	[#########_] 89%
2	ASM_PLIOPSXDP	14656820	14656820	9778180	67	4878640	33	[###_____] 33%

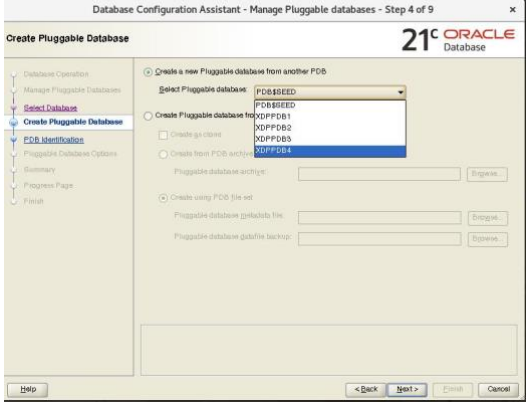

Figure3: Oracle ASM Diskgroup and PDB Setup

As shown in Figure 3, during Oracle ASM with internal redundancy setup, the total usable capacity was 5.4TB. However, after provisioning 4 Pluggable Databases (PDBs) with 1TB each (index files and log files total size increases to 1.35TB), the free capacity reduces to 650MB. When utilizing a Pliops-accelerated and protected volume with the same SSD setup, it exposed 14TB of space. Provisioning 4 PDBs with 1TB each resulted in 9.8TB of free space. This significant increase of close to 10TB of additional space provides a notable benefit for Oracle customers. It allows them to provision additional PDBs or accommodate database growth without incurring additional storage costs. Growing enterprise data continues to drive the need for higher capacity, higher performance, and easier manageability. Oracle Database 21c multi-tenancy simplifies the management of pluggable databases to greater extent and when combined with Pliops XDP, it enables higher capacity to provision additional databases or extend the database capacity, providing further flexibility and scalability to Oracle customers.

50% Higher PDB Consolidation with Pliops

With the database capacity expansion benefits realized using Pliops XDP, the next objective was to assess the performance benefits. The Transactions per Minute (TPM) metric was captured after executing HammerDB TPROC benchmark for a 60-minute duration with four pluggable databases initially without Pliops XDP and compared with Pliops XDP setup. Subsequently, two pluggable databases were added to evaluate the aggregate performance gains with Pliops XDP. The automatic workload repository and active session history reports were also captured to assess the read and write latency and CPU, IO information. In addition to this, write amplification was also captured to assess SSD durability and performance challenges.

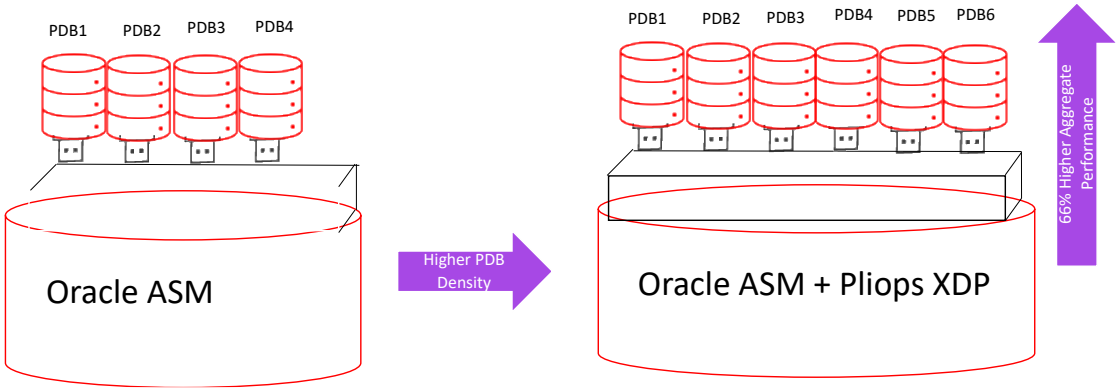

Figure 4: 4 Pluggable Databases to 6 Pluggable Database with 60% Aggregate Performance Gains

As depicted in Figure 4, setting up Pliops XDP enables the provisioning of two additional Pluggable Databases (PDBs), and accomplish 66% higher aggregate performance gain.increases increases. It enables oracle customers to increase pluggable database consoldation (PDB) density and increased performance gains.

AWR Top10 Wait events Without XDP

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
log file sync	2.3E+08	358.4K	1.56ms	47.1	Commit
DB CPU		225.1K		29.6	
log file switch (checkpoint incomplete)	50,457	148.2K	2937.73ms	19.5	Configuration
db file sequential read	45,763,119	22.8K	497.29us	3.0	User I/O
library cache: mutex X	1,485,659	4356.7	2.93ms	.6	Concurrency
latch: redo writing	1,293,935	601.4	464.75us	.1	Configuration
buffer busy waits	255,358	542.8	2.13ms	.1	Concurrency
cursor: pin S	313,110	470.9	1.50ms	.1	Concurrency
latch: enqueue hash chains	160,774	468.1	2.91ms	.1	Other
SQL*Net message to client	2.5E+08	301.4	1.19us	.0	Network

AWR Top10 Wait events With XDP

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
DB CPU		484K		46.4	
log file sync	3.8E+08	412.3K	1.07ms	39.6	Commit
db file sequential read	83,826,157	128.1K	1.53ms	12.3	User I/O
library cache: mutex X	3,854,548	11.3K	2.92ms	1.1	Concurrency
latch: enqueue hash chains	571,502	2423	4.24ms	.2	Other
cursor: pin S	650,233	1128.1	1.73ms	.1	Concurrency
SQL*Net message to client	4.2E+08	703.1	1.66us	.1	Network
latch: cache buffers lru chain	694,815	474.4	682.76us	.0	Other
latch: redo allocation	189,692	409.3	2.16ms	.0	Other
log file switch completion	64,082	405.8	6.33ms	.0	Configuration

Figure 5: AWR Report without XDP vs AWR Report with XDP

The AWR report (Automatic Workload Repository) as shown in figure 5 provides insights into the reasons behind the higher performance achieved using Pliops XDP with six PDBs compared to the setup without Pliops XDP with four PDBs. Firstly, the average latency associated with redo log write events, specifically the logfile sync, is reduced from 1.56 milliseconds to 1.07 milliseconds. This indicates improved efficiency in writing transaction log data. Secondly, the CPU utilization for executing transactions increases from 29.6% to 46.4%. This demonstrates that the system can handle a higher workload and process transactions more efficiently. Additionally, the logfile checkpoint incomplete event, which had an average wait time of 2937 milliseconds, is eliminated with the Pliops XDP setup. This suggests improved checkpointing operations, leading to better database performance. Finally due to a higher number of pluggable databases and higher number of concurrent user requests with Pliops XDP setup, the database file sequential read that is associated with single block read events is at 1.53 milli seconds compared to 497 microseconds without Pliops XDP setup.

Database Backup Offloading with Pliops

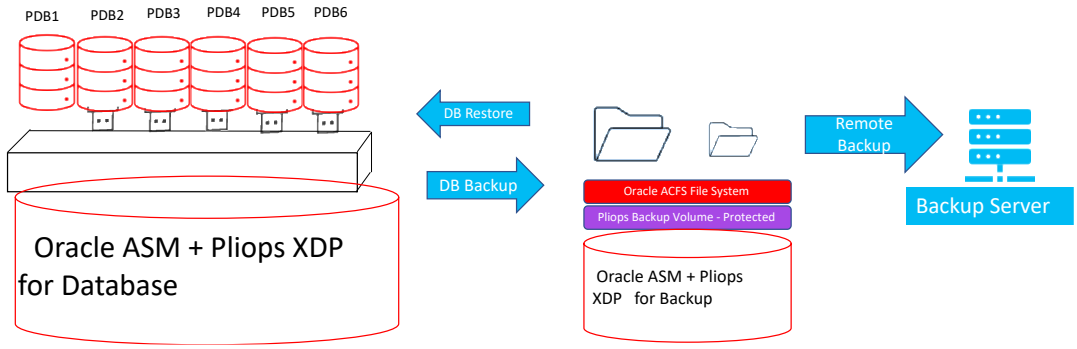

Figure 6: Oracle Database Backup & Recover Testing with Pliops XDP

```
RMAN> SHOW ALL;

RMAN configuration parameters for database with db_unique_name XDPCDB are:
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE DISK PARALLELISM 20 BACKUP TYPE TO BACKUPSET;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT   '/u01/bckup_xdp_r/cold_bckp_r/%U';
CONFIGURE MAXSETSIZE TO UNLIMITED; # default
CONFIGURE ENCRYPTION FOR DATABASE OFF; # default
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE COMPRESSION ALGORITHM 'BASIC' AS OF RELEASE 'DEFAULT' OPTIMIZE FOR LOAD TRUE ; # default
CONFIGURE RMAN OUTPUT TO KEEP FOR 7 DAYS; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/u01/app/oracle/dbs/snapcf_xdpcdb.f'; # default
```

Database backup and restore is one of the main responsibilities of Oracle Database administrators to help ensure the databases are protected against storage media corruption or accidental drop of business-critical tables. We wanted to explore how Pliops XDP can be utilized not only as a database accelerator but also serve as database backup acceleration and offloading as shown in figure 6.

Figure 7: Oracle PDB Full Database Backup & restore testing with Pliops XDP

The additional 9TB of disk space exposed by Pliops XDP is leveraged to carve out ASM backup disks, and an ACFS file system is mounted to serve as the backup location for these pluggable databases.

Backup Operation	Database Size (GB)	Duration (Minutes)	Backup Size (GB)	CPU Utilization
Full Database Backup for 5 PDBs	5500	1.07:51	4100	6%
Validate Backup Set	2500	0.12.35		7.5%
Restore Backup for PDBs	5500	0.23.47	4100	7.3%

The test metrics of duration, storage IO, capacity and CPU utilization were captured during full database backup, differential backup, and restore operations. As shown in the table below 5 pluggable databases with around 5.5TB database is backed up with compression in 1 hour and 7 minutes with just 6% CPU utilization.

Table 1: Database Backup & Recovery Testing

The database restoration for these 5 pluggable databases is under 25 minutes with 7% CPU utilization. Based on additional experiments, the backup duration can be considerably reduced with Oracle Recovery Manager's (RMAN) parallel operations and multi-channel implementation. The goal of the testing is to validate and obtain the out-of-the-box Oracle Database backup and recovery operations metrics and this can be further improved with additional optimization.

SSD Write Amplification Challenges

Chart 1: Mixed workload performance benefits of Pliops XDP

The primary challenges with traditional NVMe SSD storage configurations with RAID10 or RAID 5 are the write and read amplification overhead that not only limits the database performance but also reduces the longevity of SSD usage. Write amplification increases the amount of data that's written to SSDs compared to database application workloads. Since SSDs have a limited number of write cycles, write amplification results in shortening SSDs' durability. Pliops XDP mitigates the SSDs durability issue significantly by providing the in-line compression of Oracle Database data and then data shaping activities that include packing and merging the variable length data and finally converting random writes to sequential writes. This process minimizes the write amplification to a significant level and thereby increasing database performance and SSD durability. Chart 1 demonstrates that during Oracle pluggable database workload testing IO throughput, the Pliops XDP delivers a robust write throughput of 4.7 GB/s to the Oracle Database. Yet, the per-SSD write rate is just 620 MB/s a 7.5X reduction.

Collectively, all four SSDs incur a 2.4 GB/s total disk write of, nearly halved. In contrast, without the Pliops XDP setup, each SSD experiences 2.3 GB/s high amount of disk writes—equivalent to the combined write for all SSDs within the XDP configuration. Based on these test results we can draw on conclusion that with Pliops enabled setup drives higher performance and improves the longevity (useful life) of the SSDs by minimizing the disk writes from database applications

Oracle Linux

Oracle Linux is a robust, secure, and optimized operating system for application development and deployment. It supports X86 and Arm for on-premises, MultiCloud and Edge environments. based applications. It's tuned for business-critical demanding workloads like Oracle database at cloud scale and on-premises deployments. It Increase scalability, security, and productivity by automating manual processes in DevOps with open source, ansible-based automation tools.

Conclusion

The combination of Oracle Database 21c on Oracle Linux with Pliops XDP offers several benefits to enterprises. Oracle Database 21c with Pliops XDP provides higher pluggable database density, higher aggregate transaction performance, database capacity expansion, and backup offload benefits to consolidate databases and simplify overall IT operations. The built-in performance acceleration, capacity expansion, and data protection features enable enterprises to efficiently manage data growth challenges without performance and reliability impacts. The built-in data protection enables enterprises to increase database uptime and resiliency. This solution also significantly improves the system and storage utilization by offloading database backups and provides a faster recovery to meet or enhance availability service level objectives for critical applications.