

# Northern California Oracle User Group FALL CONFERENCE – 2017

## Oracle Hybrid Transaction/Analytical Processing

Ganesh Balabharathi  
Principal Database Architect  
[Ganesh@vexata.com](mailto:Ganesh@vexata.com)

# Agenda

- ✓ **Quick Review: Oracle OLTP application** - What it does or What it needs (from the base solution)?
- ✓ **Quick Review: Oracle OLAP application**
- ✓ **Quick Review: HTAP - Hybrid Transactional/Analytical application**
- ✓ **Deep Dive: HTAP 'key' requirements**
- ✓ **Demo on Oracle HTAP application performance with Vexata Storage array**
- ✓ **Q&A**

# On-Line Transactional Processing (OLTP)

What it does or What it needs???

# OLTP (On-Line Transaction Processing)

- ✓ **Short transactions by large number of concurrent users**
  - ✓ In industry terms, it require huge IOPS
- ✓ **Frequent data loads, data maintenance operations**
  - ✓ In industry terms, it require huge THROUGHPUT
- ✓ **Quick response time**
  - ✓ In industry terms, it require low LATENCY

# OLTP (On-Line Transaction Processing)

## ✓ **Short transactions by large number of concurrent users**

Short transactions (inserts/updates/deletes) involves small amount of data and time.

Large number of concurrent users can access the same data at the same time.

Example applications:



Online / Retail stores

# OLTP (On-Line Transaction Processing)

How do we measure transactions, concurrent users?

AWR: Transactions per second, us

Load Profile		Per Second	Per Tr
DB Time(s):		209.9	
DB CPU(s):		42.5	
Background CPU(s):		11.7	
Redo size (bytes):		141,575,164.6	
Logical read (blocks):		1,043,457.4	
Block changes:		908,905.7	
Physical read (blocks):		637,538.6	
Physical write (blocks):		389,177.4	
Read IO requests:		637,535.8	
Write IO requests:		302,144.6	
Read IO (MB):		4,980.8	
Write IO (MB):		3,040.5	
IM scan rows:		0.0	
Session Logical Read IM:		0.0	
RAC GC blocks received:		11.2	
RAC GC blocks served:		4.8	
User calls:		17.2	
Parses (SQL):		6.4	
Hard parses (SQL):		0.2	
SQL Work Area (MB):		1.9	
Logons:		2.8	
Executes (SQL):		29,432.3	
Rollbacks:		0.0	
Transactions:		14,124.6	

Key Instance Activity Stats				DB/Inst: VXDB/vxdb1	Snap: 61-62
-> Ordered by statistic name					
Statistic	Total	per Second	per Trans		
user commits	12,755,997	14,124.6	1.0		
user rollbacks	0	0.0	0.0		

Instance Activity Stats				DB/Inst: ORCL/orcl	Snap: 22-23
-> Ordered by statistic name					
Statistic	Total	per Second	per Trans		
user logons cumulative	1	0.0	0.1		
user logouts cumulative	1	0.0	0.1		

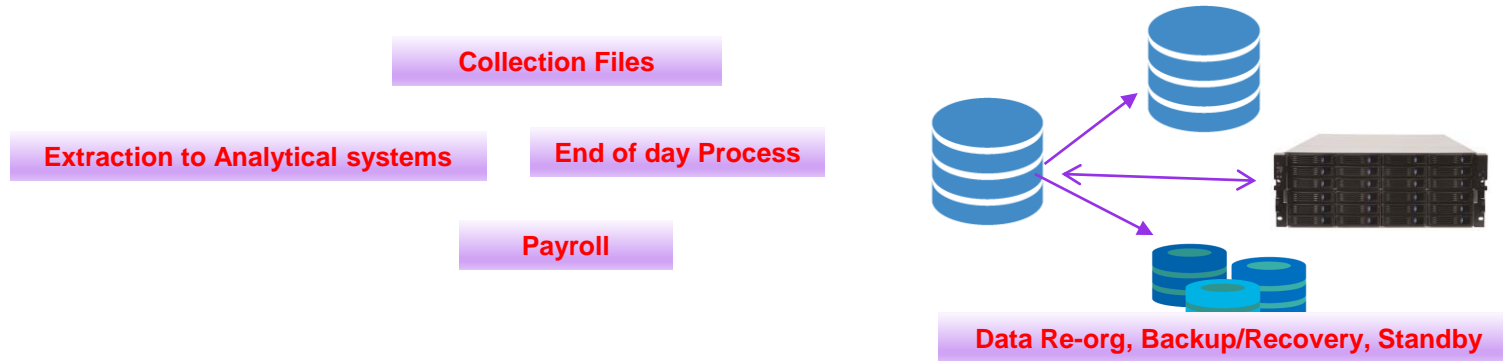
Instance Activity Stats - Absolute Values				DB/Inst: ORCL/orcl	Snap: 22-23
-> Statistics with absolute values (should not be diffed)					
Statistic	Begin Value	End Value			
logons current	320	323			

# OLTP (On-Line Transaction Processing)

## ✓ Frequent data loads, data maintenance operations

OLTP systems runs many small/medial batch/ETL processing (Daily loads, Data Extractions).

OLTP systems transmit data to standby systems, runs backups/recovery, data reorganization.



# OLTP (On-Line Transaction Processing)

## ✓ **Faster Response time**

OLTP systems require instant response in order for the application and users to remain productive

Example applications:



Check-ins



Deposits/Withdraw



Survey

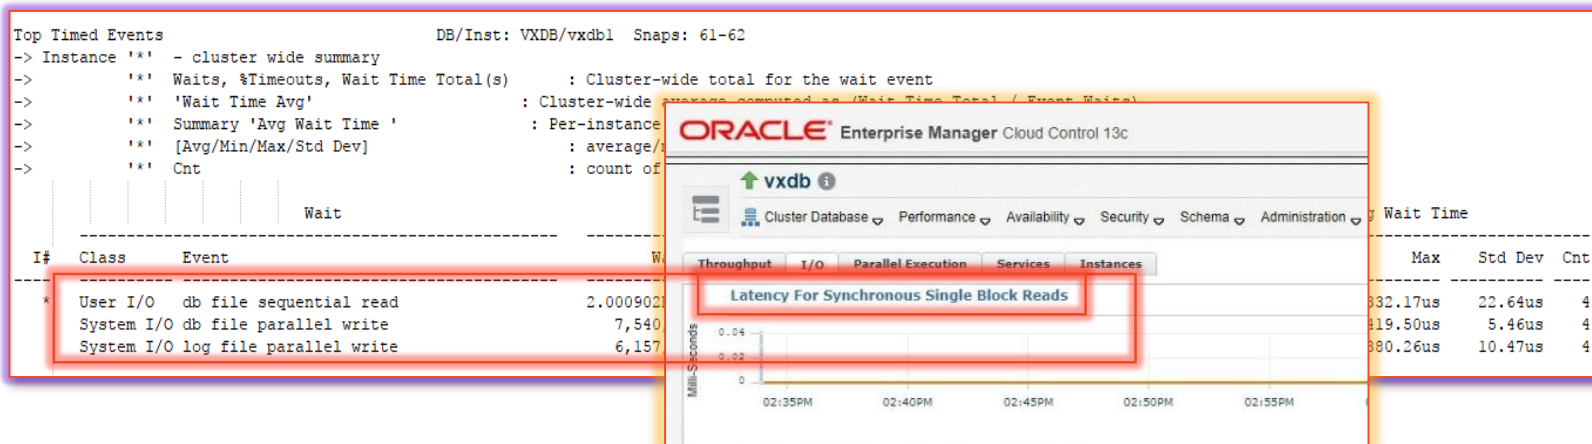
# OLTP (On-Line Transaction Processing)

- How do we measure response time?

AWR and EM Cloud control:

Read Latencies: DB File Sequential Reads, Latency For Synchronous Single Block Read

Write Latencies: DB File Parallel Write, LOG File Parallel Write,



# OLTP (On-Line Transaction Processing)

## Summary:

- ✓ **Huge IOPS** - Short transactions by large number of concurrent users
- ✓ **Huge Throughput** - Frequent data loads, data maintenance operations
- ✓ **Low Latency** - Quick response time

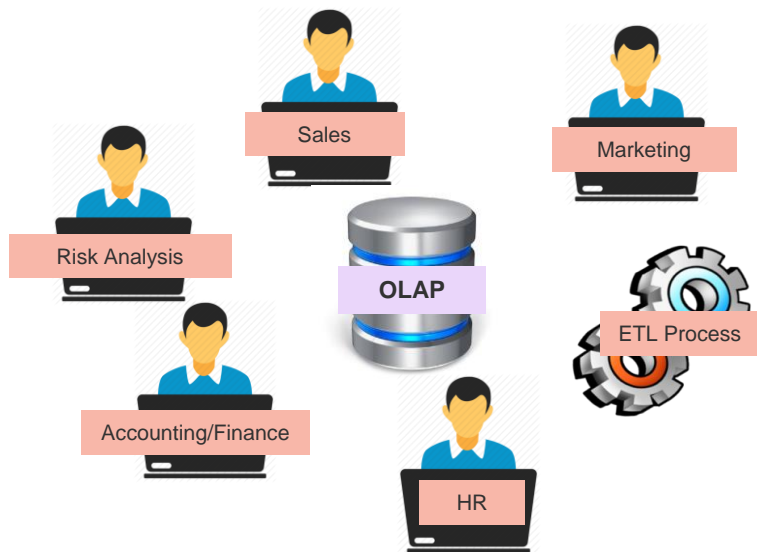
# On-Line Analytical Processing (OLAP)

What it does or What it needs???

# OLAP (On-Line Analytical Processing)

✓ **High Throughput** with preferably low response time

OLAP systems require massive scan capability on large data with low response time



# OLAP (On-Line Analytical Processing)

## How do we measure Throughput and Latencies?

AWR Latencies: Direct Path Read/Write, DB File Sequential/Scattered Read

Throughput: IO Profile section

Top Timed Events			DB/Inst: VXDB/vxdbl Sns: 80-81			
			IO Profile (Global)			
			DB/Inst: VXDB/vxdbl Sns: 80-81			
I#	Class	Event	Statistic	Read+Write/s	Reads/s	Writes/s
*	User I/O	direct pa	Total Requests	49,097.11	44,303.50	4,793.61
	User I/O	direct pa	Database Requests	49,068.43	44,280.28	4,788.16
	User I/O	db file s	Optimized Requests	0.00	0.00	0.00
	User I/O	db file s	Redo Requests	1.50	N/A	1.50
			Total (MB)	41,252.17	40,133.51	1,118.66
			Database (MB)	41,251.74	40,133.14	1,118.61
			Optimized Total (MB)	0.00	0.00	0.00
			Redo (MB)	0.01	N/A	0.01
			Database (blocks)	5,280,223.06	5,137,041.54	143,181.52
			Via Buffer Cache (blocks)	1,860.28	1,689.54	170.74
			Direct (blocks)	5,278,362.68	5,135,351.89	143,010.78

# OLAP (On-Line Analytical Processing)

## Summary:

### ✓ **Throughput**

(Small number of users but long transactions, Large volume of data scans)

**So far so good...but**

Lets say..

You got an email from your management team to explore  
HTAP option (aka Real Time Analytics)!

Because..

# The conventional ways are not good enough!

---

## **Marketing/Analytical team:**

OLAP data is not current

Analytical decisions are being made on old data

Too many versions of same data

## **Infrastructure/Storage team:**

Redundant storage & servers

Duplicate data is stored at many places

## **Operational/IT team:**

Too many jobs to run

Too many tools to support

Too many licenses & support costs

They are right!  
Lets explore ..

Hybrid Transaction Analytical Processing (HTAP)

# OLTP + OLAP = Hybrid Transaction/Analytical Processing

What is a Well Tuned Solution means?



OLTP

## Well Tuned OLTP Solution

Database  
Storage

## Database Parameters/Options

Block Size  
Caching Options  
Parallelism  
Indexes / Full Table Scans

## Storage Parameters/Options

Volume Block Size  
IO Types (Seq / Random)  
IO Mix



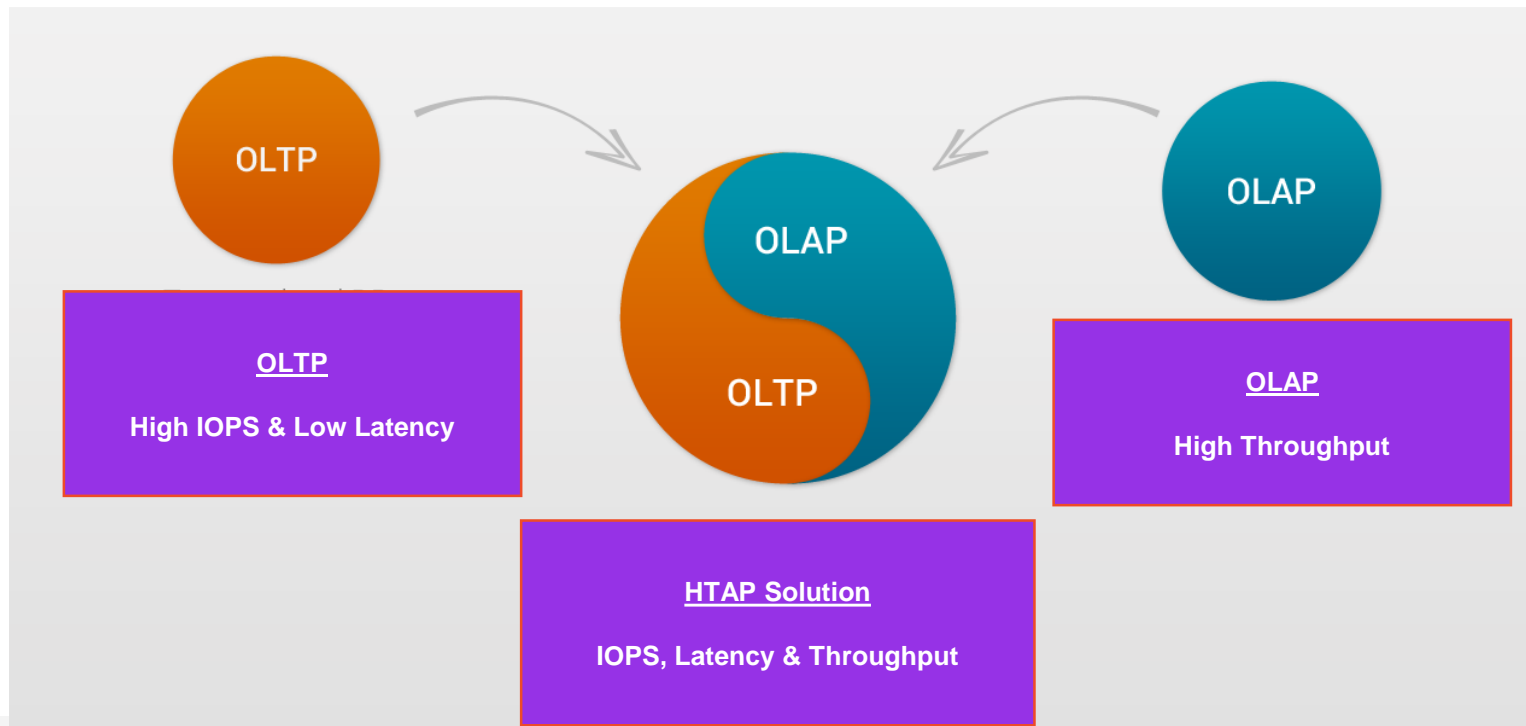
OLAP

## Well Tuned OLAP Solution

Database  
Storage

# OLTP + OLAP = Hybrid Transaction/Analytical Processing

In summary, HTAP is a **superset** of OLTP and OLAP..

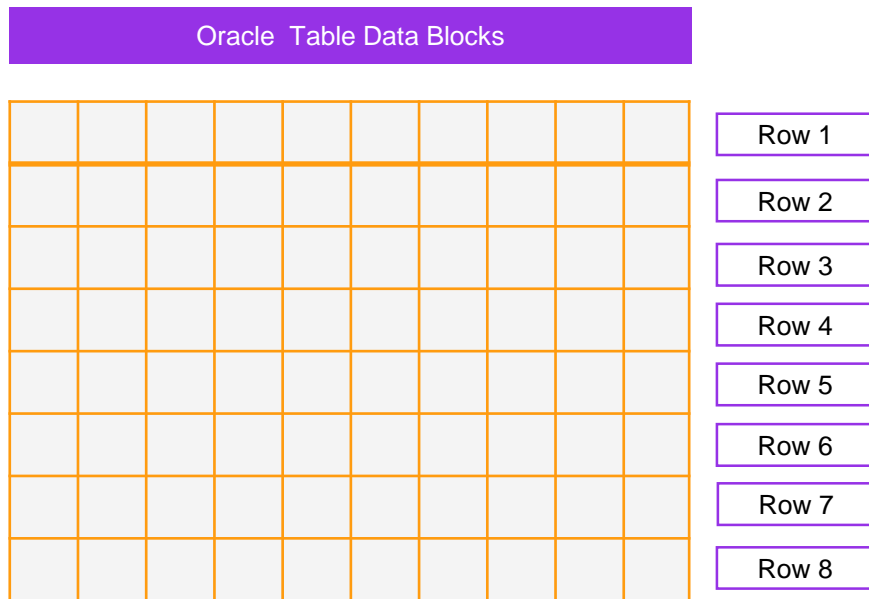


OLTP + OLAP = Hybrid Transaction/Analytical Processing

HTAP key requirements (deep dive)

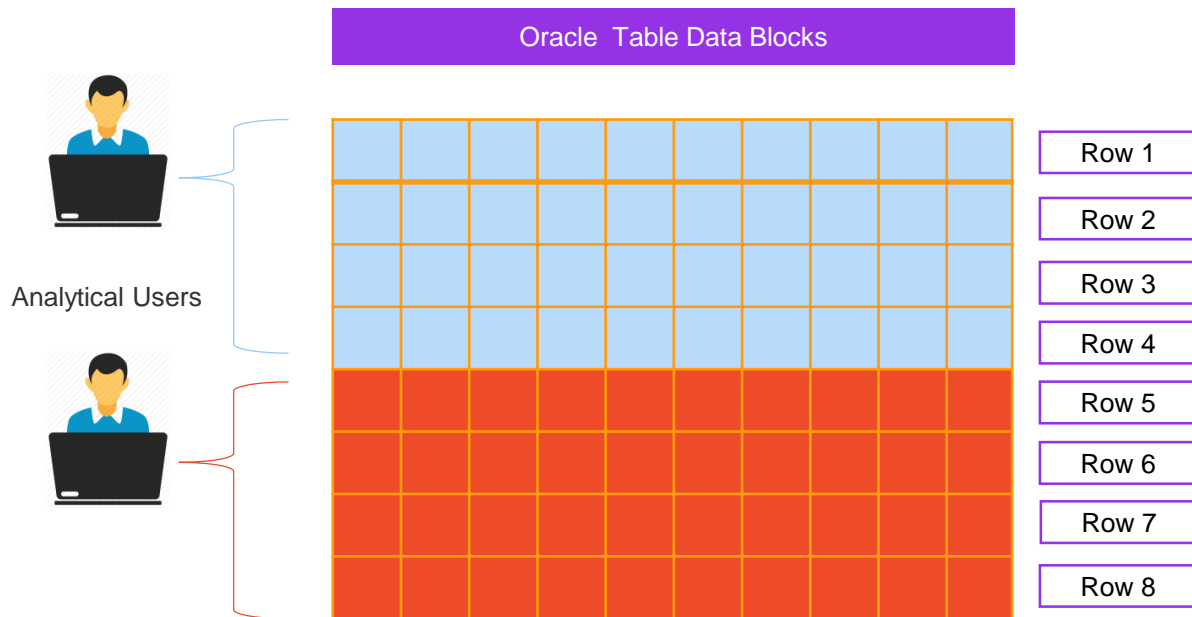
# HTAP key requirements

- Solution should support effective data consistency, concurrency and integrity:



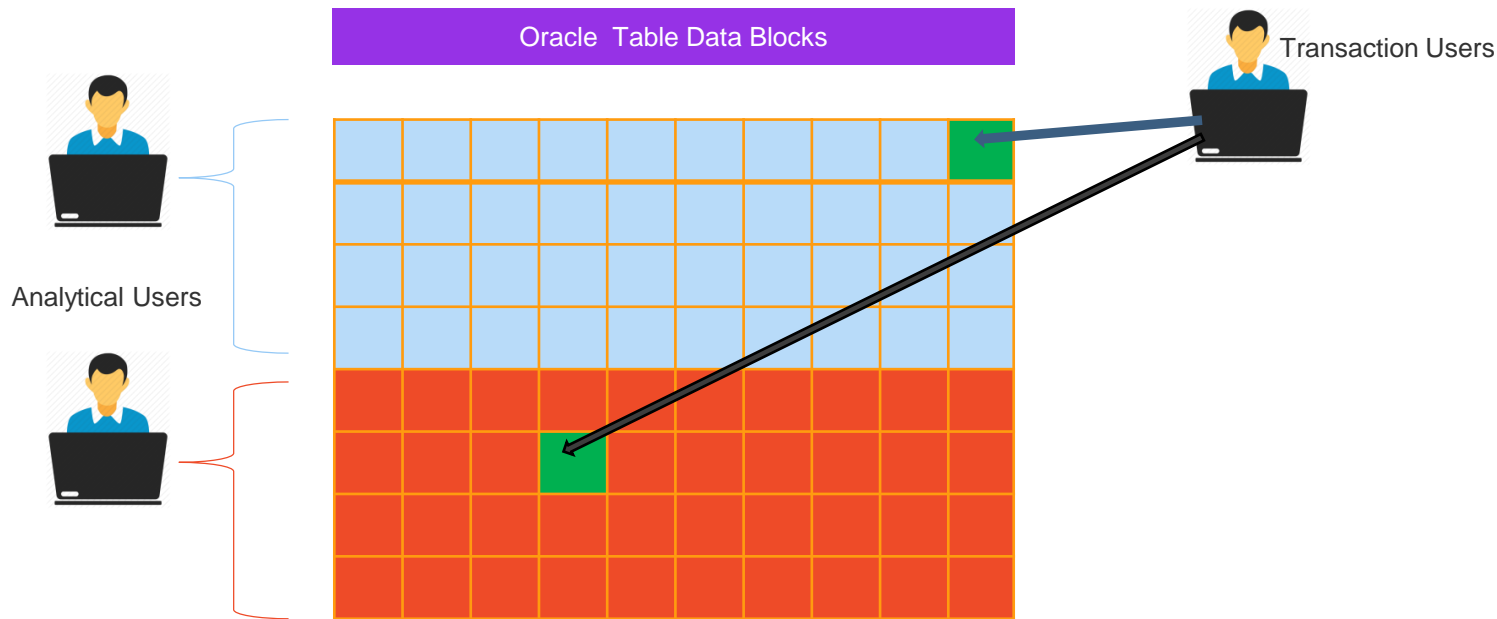
# HTAP key requirements

- Solution should support effective data consistency, concurrency and integrity:



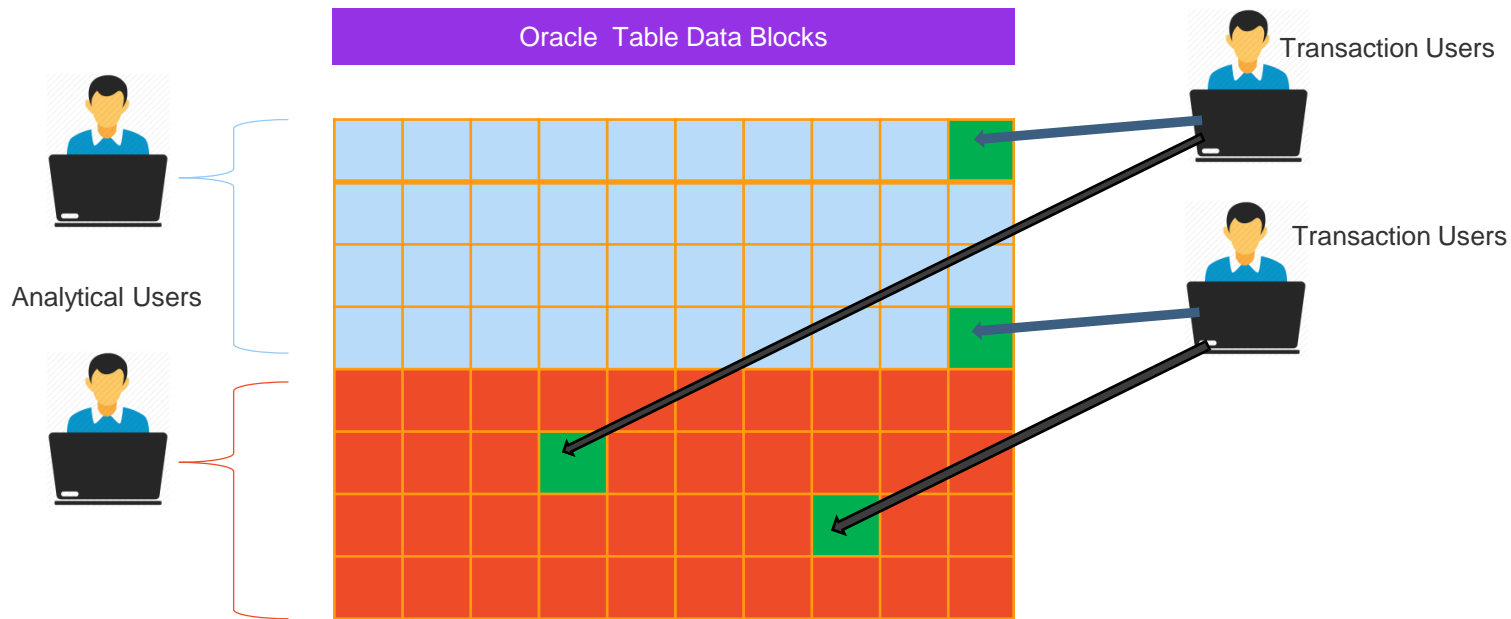
# HTAP key requirements

- Solution should support effective data **consistency, concurrency** and integrity:



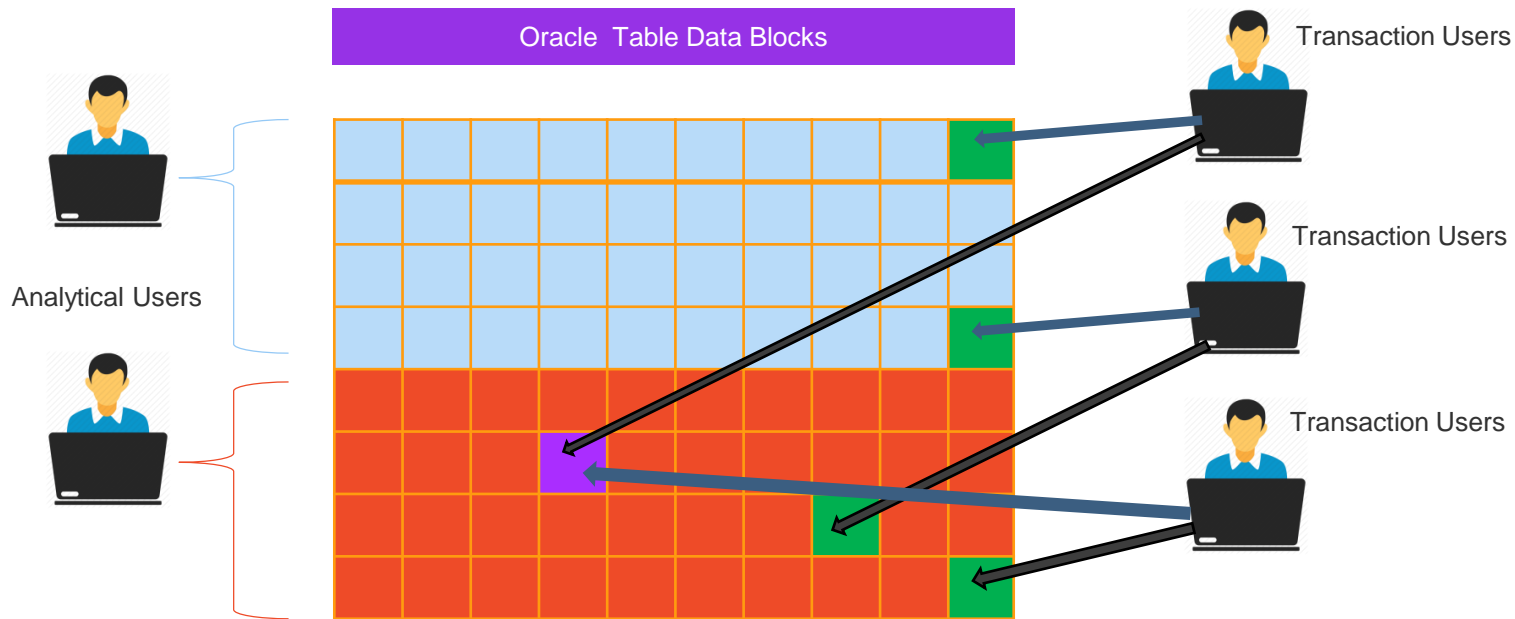
# HTAP key requirements

- Solution should support effective data **consistency, concurrency** and integrity:



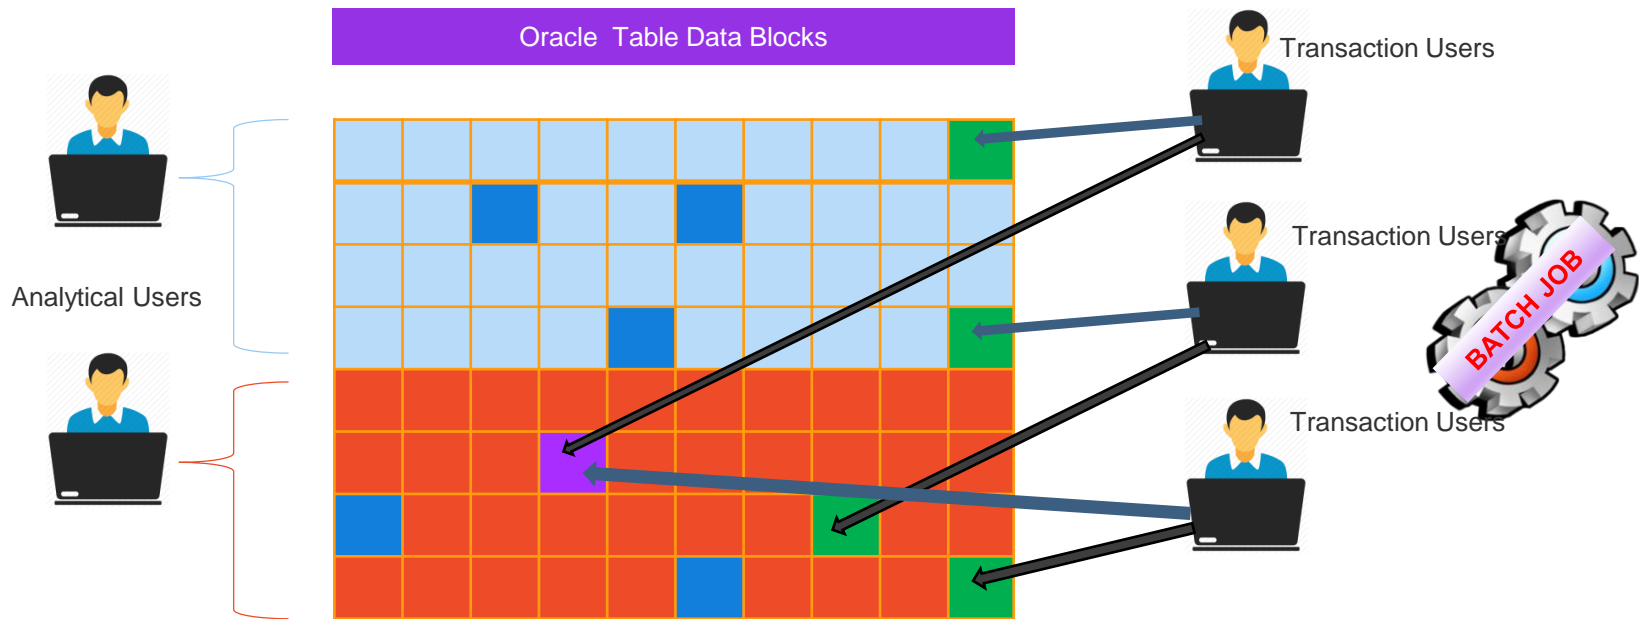
# HTAP key requirements

- Solution should support effective data **consistency, concurrency** and integrity:



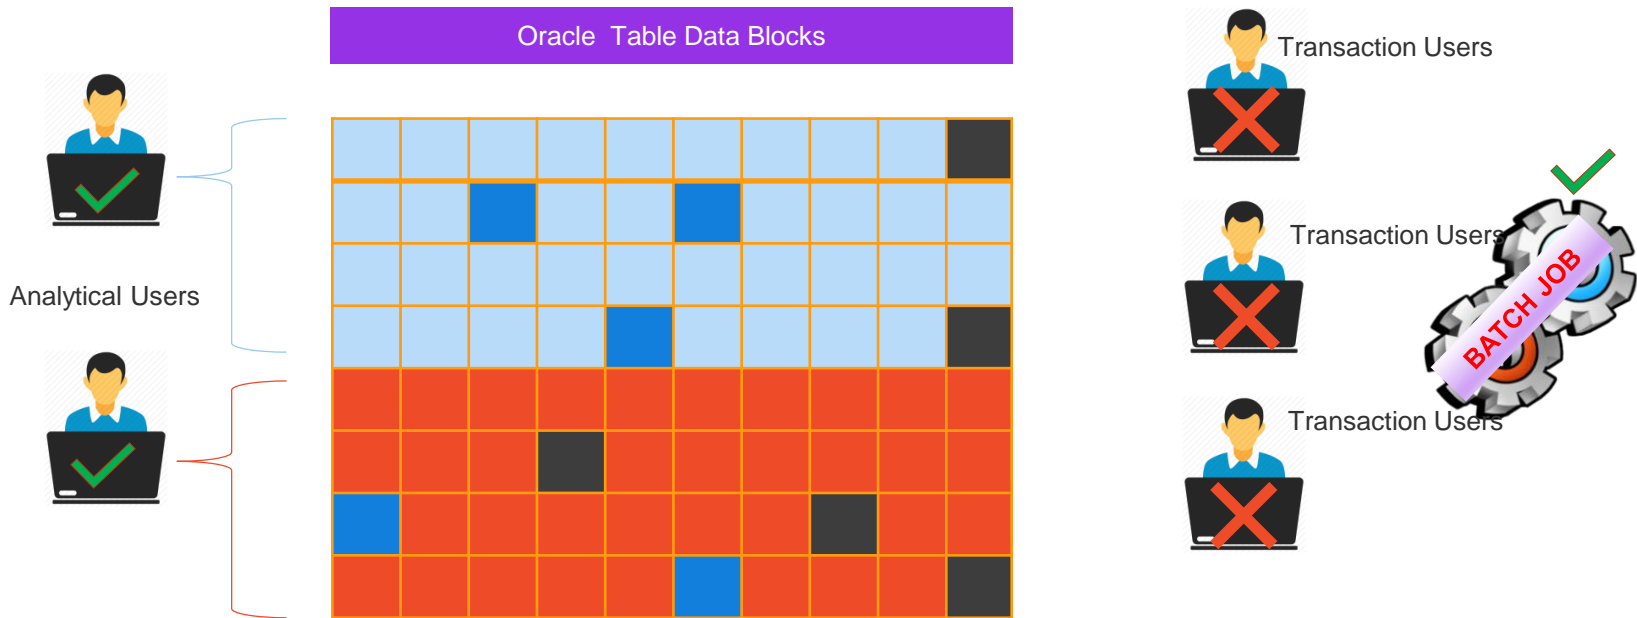
# HTAP key requirements

- Solution should support effective data **consistency, concurrency** and integrity:



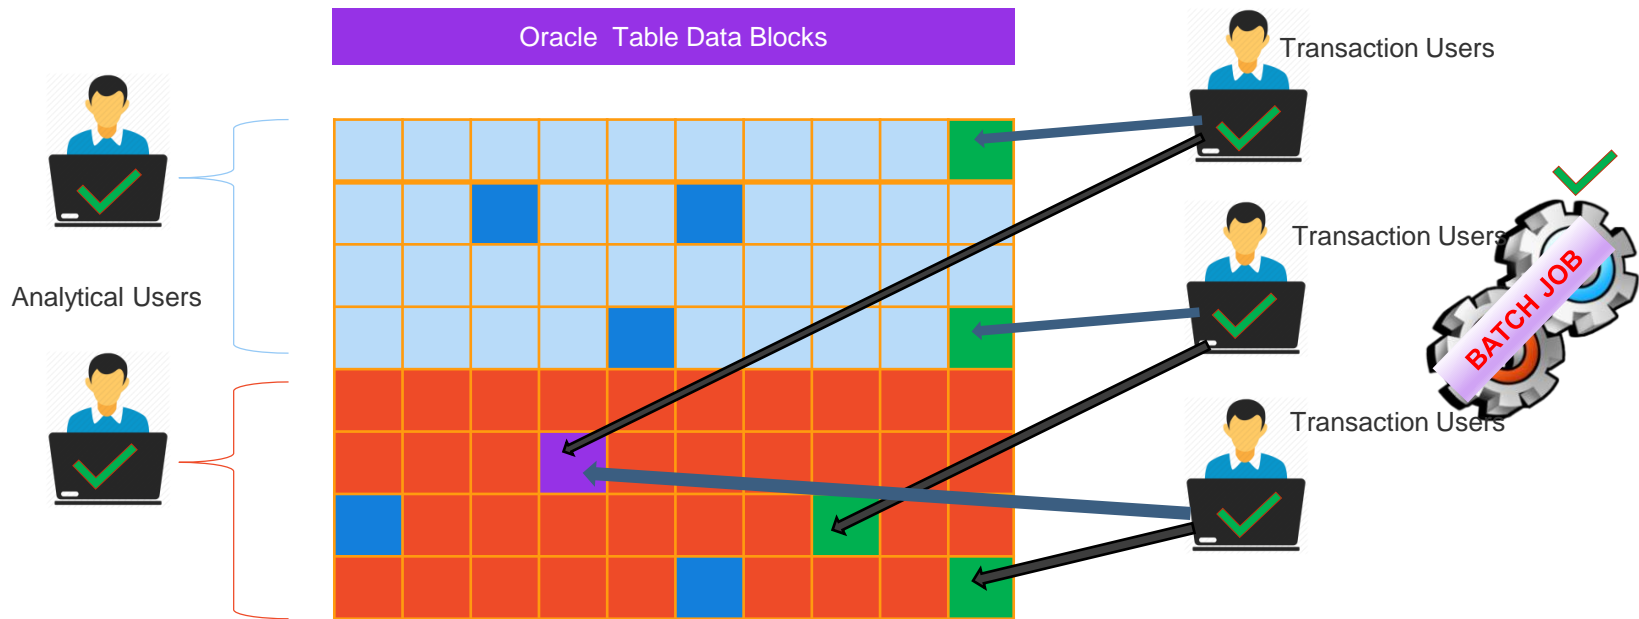
# HTAP key requirements

- Solution should support effective data **consistency, concurrency and integrity**:



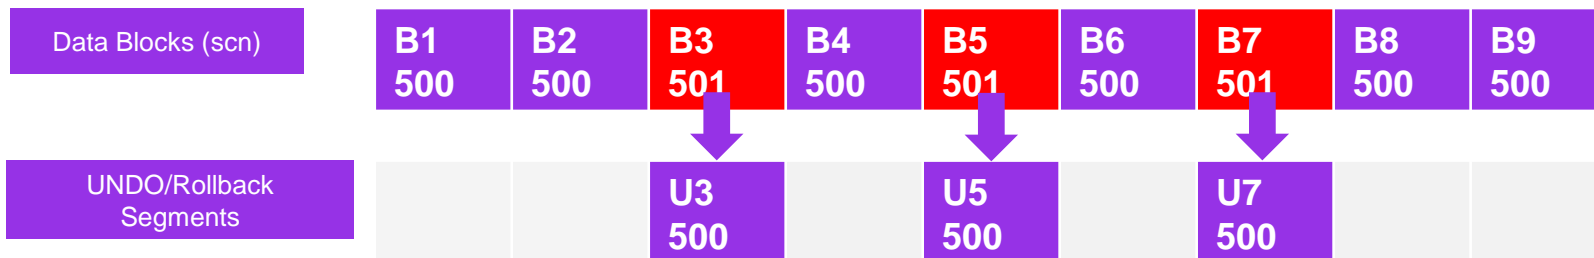
# HTAP key requirements

- Solution should support effective data **consistency, concurrency and integrity**:



# HTAP key requirements

- ✓ **Solution should support effective data consistency, concurrency and integrity:**
  - Oracle uses **multi-version consistency model** and various types of locks and transactions.
  - These methods ensures consistency, concurrency and integrity and ensures no data is blocked for read only queries.
  - Example: **Select \* from customer where <some conditions>**
  - When the above SQL enters execution stage, lets assume Oracle uses '**Statement-Level Read Consistency**' and **SCN 500** is determined..



# HTAP key requirements

- **Solution should handle dynamic workload changes without configuration changes.**
  - Oracle database with Vexata storage solution is capable of supporting dynamically varying workloads without having to adjust the configurations.



**ORACLE®**  
DATABASE **12<sup>c</sup>**

# HTAP key requirements

- **Solution should have better cache management:**
  - Oracle caching options helps to **KEEP** data required for OLTP.
  - Vexata array can be used for FLASH CACHE options.



OLTP applications benefit from CACHE  
(Buffers cache, SMART DB Flash cache)



OLAP applications may not benefit from CACHE  
But can bring unnecessary data to buffers



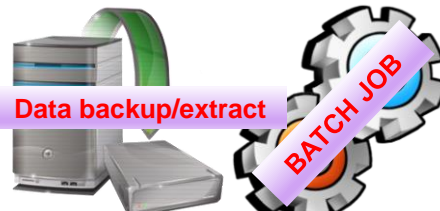
**ORACLE®**  
DATABASE **12<sup>c</sup>**

# HTAP key requirements

- **Solution should have both OLTP and OLAP characteristics:**
  - Oracle 12c + Vexata solution supports both OLTP and OLAP needs.



High IOPS, Ultra Low Latency, High concurrent users



High Throughput, High volume data scanning, Long running Queries



**+** **ORACLE®** **12<sup>c</sup>**  
DATABASE

# HTAP Benefits

- **Operational Efficiency:**
  - Data **doesn't have to move** from OLTP to OLAP (No More ETL!)
- **Real time Analytics:**
  - Transactional Data is **instantly available** for analytics
- **Cost Savings:**
  - Eliminates the **duplicate data and related infrastructure** investments
- **Accurate Business Decision**
  - Business decisions are based on **real time data** (No more Guestimates!)

# Vexata + Oracle Hybrid database performance

## 4 Node Oracle RAC + Vexata VX-100M Performance Highlights...

### DBAs + Developers

HammerDB TPC-C + HammerDB TPC-H

OLTP – TPM , OLAP - Throughput

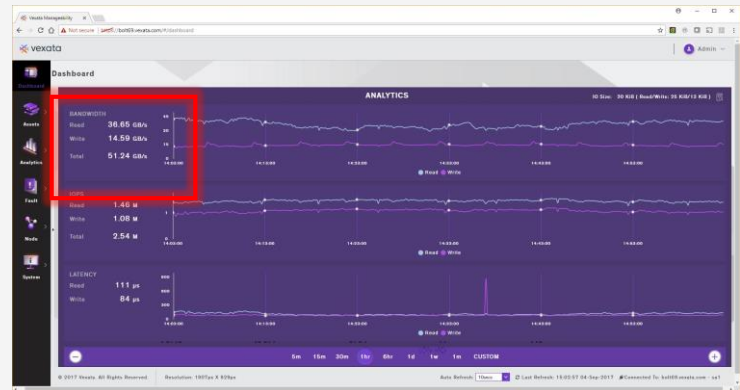
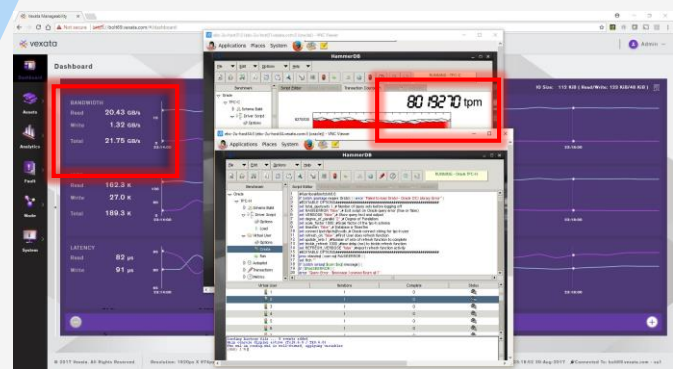
8M Transactions Per Minute + 20GB/s Throughput

### DBAs + Storage Admins + Architects

SLOB + HammerDB TPC-H

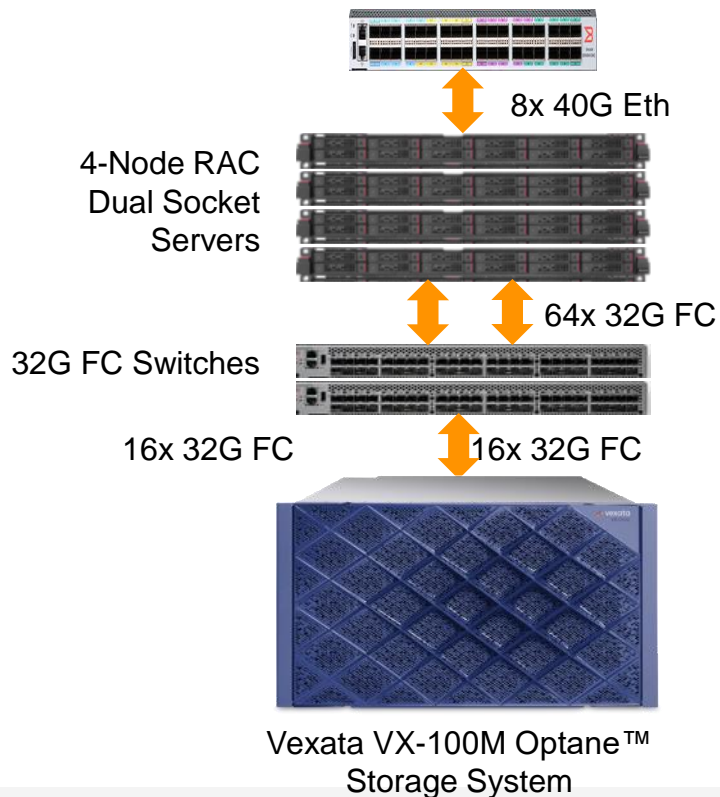
OLTP + OLAP – IOPS/Throughput

37 GB/s Reads and 15 GB/s Write Throughput



# Hybrid Transactional and Analytical Processing - Demo

# Oracle on Vexata VX 100-M Optane Storage System



## ■ Host Setup:

- 2x Intel Xeon (Skylake) CPUs with 28 cores/socket
- 512GB RAM, 2x Dual Port 32Gbps HBAs
- Oracle Linux 6.8 64bit
- Oracle Grid Control 12.2.0.1
- Oracle Database 12.2.01
- Oracle Enterprise Manager Cloud Control 13c

## ■ Network Setup

- 1x Brocade VDX 6940 Switch w/ (36) 40GbE ports
- 2x Brocade G620 w/(48) 32Gbps FC ports

## ■ Storage Setup

- Vexata VX-100M Optane Storage System
- Vexata VX-OS 3.0.0 (v3.0.1-5)
- 16 x 32Gbps FC Ports
- 16 x ESM with 4x375GB PCIe Optane™ (3D Xpoint) Drives
- 15.5 TB Usable Capacity (RAID 6)

# HTAP - Demo

C:\NoCOUG\nocoug-htap-11092017.mp4