


---

# Performance Baseline of Hitachi Data Systems HUS VM All Flash Array for Oracle

Storage and Database Performance

Benchmark Performance Suite Release 8.5 (Build 131015)

December 2013



## **1 System Configuration**

2 Introduction to Storage Performance Tests

3 Storage Benchmark Results – Sequential I/O

4 Storage Benchmark Results – Random I/O


5 Introduction to Database Performance Tests

6 Database Benchmark Results – Database Load

7 Database Benchmark Results – OLTP Transactions

8 Reviewing Storage and Database Benchmark Results


# System Configuration



## Server and Cluster Configuration

Server	HDS UCP Large SMP
CPU type	Intel Xeon E7-8870 Westmere
#sockets	4
#cores	40
#threads	80
RAM capacity	256 GByte
Cluster	
#server within cluster	4
#cores within cluster	160
Software	HDS UCP Large SMP
Operating System	Oracle Linux 6.1
Oracle Database System	11.2.0.4
Benchware Performance Suite	8.5 Build 131015


# System Configuration



## Storage System Architecture and Configuration

FC attached Storage System	HDS HUS VM All Flash Array
Host connectivity	32 x 8 Gbps
DRAM Cache	256 GByte
Flash storage <ul style="list-style-type: none"><li>#modules</li><li>Capacity raw</li><li>Capacity formatted</li></ul>	32 32 x 1.6 = 51.2 TByte 25.6 TByte (Raid-10)
Number of attached database server <ul style="list-style-type: none"><li>#hba's single db server</li><li>#ports single db server</li></ul>	4 4 8

# System Configuration




## Volume and File Management

Volume Manager	Oracle ASM
#LUN's for data	64
#LUN's for redo	32
Concatenated oder Striped	striped
Queue depth	128
Special parameters <ul style="list-style-type: none"><li>asm_au_size</li></ul>	4 MByte

File System	Oracle ASM
Block size [KByte]	-
Direct I/O	Yes
Concurrent I/O	Yes
Asynchronous I/O	Yes
Special parameters	-

# System Configuration




## Database Configuration

Database	Oracle 11g
<ul style="list-style-type: none"><li>▪ Oracle Release</li><li>▪ Partition Option</li><li>▪ Real Application Cluster</li><li>▪ Data Guard</li></ul>	11.2.0.3 Yes Yes No

Basic Configuration	Oracle 11g
<ul style="list-style-type: none"><li>▪ Block size [kByte]</li><li>▪ Archiving</li><li>▪ Force Logging</li><li>▪ Flashback</li></ul>	8 No No No

# System Configuration

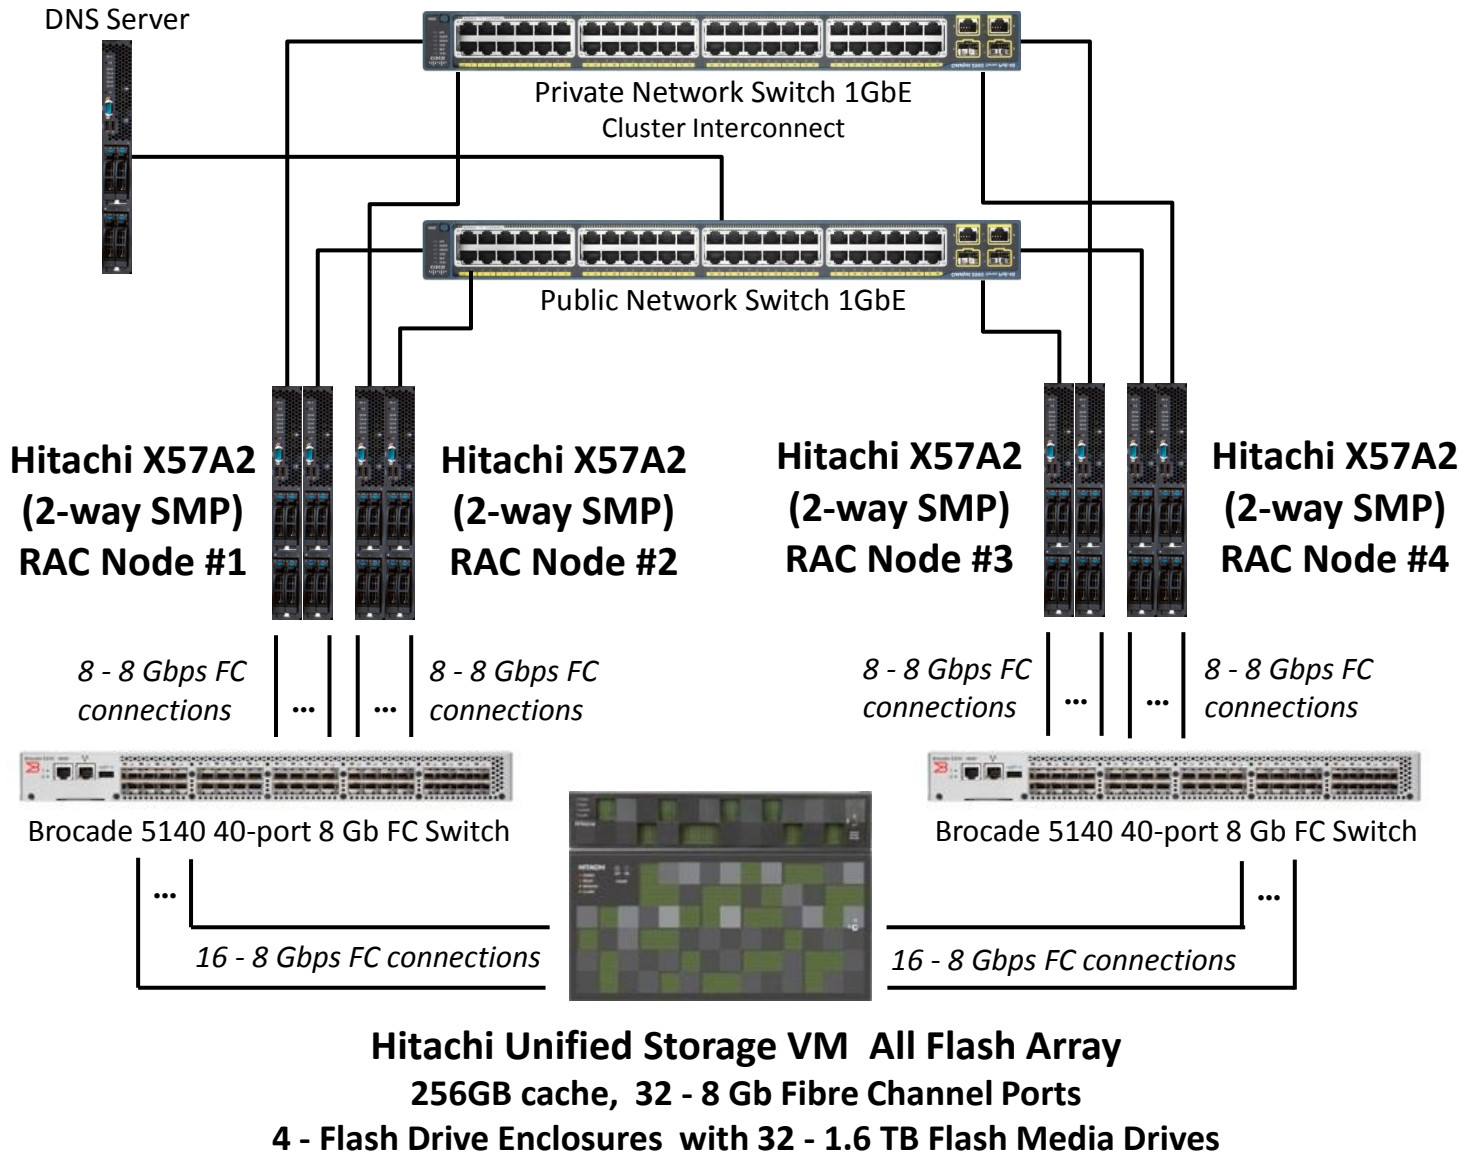


## Database Configuration


Memory Configuration	Oracle 11g
<ul style="list-style-type: none"><li>SGA target [GByte]</li></ul>	64
<ul style="list-style-type: none"><li>PGA target [GByte]</li></ul>	-
<ul style="list-style-type: none"><li>Keep pool [GByte]</li></ul>	16
<ul style="list-style-type: none"><li>Recycle pool [GByte]</li></ul>	4
<ul style="list-style-type: none"><li>Default pool [GByte]</li></ul>	16

REDO Configuration	Oracle 11g
<ul style="list-style-type: none"><li>REDO log file size [GByte]</li></ul>	8
<ul style="list-style-type: none"><li>#REDO log files groups</li></ul>	4
<ul style="list-style-type: none"><li>#REDO log file members</li></ul>	1

# System Configuration








- 1 System Configuration
- 2 Introduction to Storage Performance Tests**
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Introduction to Database Performance Tests
- 6 Database Benchmark Results – Database Load
- 7 Database Benchmark Results – OLTP Transactions
- 8 Reviewing Storage and Database Benchmark Results

# Storage Performance


---



Why measure Storage performance?

- Storage performance is essential not only for overall Oracle database performance, but also for system management tasks like backup, recovery and archiving
- Oracle uses all kinds of I/O patterns, but different o/s calls dependent upon the
  - operating system
  - system load (Oracle changes system call dependent on load)


# Storage Performance



## Why measure Storage performance?

- Oracle sequential read
  - User processes:  
full table scan, full index scan
  - Temp segment
  - Backup, restore, recovery RMAN,  
Export, Data Pump
  - ARCH:  
reading online REDO logfile
- Oracle random read
  - User processes
- Oracle sequential write
  - Temp segment
  - Backup, restore RMAN, Export,  
Data Pump
  - LWGR process: small block size
  - ARCH processes:  
writing archived REDO logfile
  - RVWR process:  
flashback log file writer
  - CTWR process: block change  
tracking file
- Oracle random write
  - DBWR processes


# Storage Performance



## What is measured?

- Storage performance from the Oracle point of view
  - Using database block size
  - I/O service time measured within Oracle
- Throughput
  - Data transfer in mega byte per second [MBps]
  - Data transfer in database blocks per second [dbps]
  - I/O Operations in O/S system calls per second [IOPS]
- Service Time
  - For random I/O operation in [ms] or [ $\mu$ s]
- Efficiency of
  - Auto-Tiering
  - RAID-level
  - Striping
  - Remote mirroring
  - Virtualization

# Storage Performance




## Overview of Storage performance tests with Benchware test codes

Oracle Storage Performance Sequential I/O	Test Code for Data in-memory	Test Code for Data on default Storage	Test Code for Data in Flash Cache	Test Code for Data in Cell Flash Cache
▪ Sequential read	-	STO-12	STO-13	STO-14
▪ Sequential write	-	STO-22	-	-
▪ Sequential write (LGWR commit stress test)	DBL-11	-	-	-


Oracle Storage Performance Random I/O	Test Code for Data in-memory	Test Code for Data on default Storage	Test Code for Data in Flash Cache	Test Code for Data in Cell Flash Cache
▪ Random read 25 I/O's per SQL	-	STO-32	STO-33	STO-34
▪ Random write (DBWR stress test)	STO-41	-	-	-
▪ Mixed random read write	-	STO-52	STO-53	STO-54
▪ More aggressive random read 25 I/O's per 125'000 rows per SQL	-	STO-62	STO-63	STO-64

# Storage Performance




Remarks on other benchmark tools . . .

- SAP, TPC, Swingbench, Hammerora, ...
  - No specific storage performance metrics
  - SAP even does not document storage configuration for its benchmarks
  
- I/O load generators like IOzone, IOgen, IOmeter, vdbench
  - Do not reflect database reality – database I/O is much more complex
  - Measure throughput and service times on operating system level, but not on database level
  
- `dbms_resource_manager.calibrate_io`
  - Does not recognize hybrid flash/disk architectures

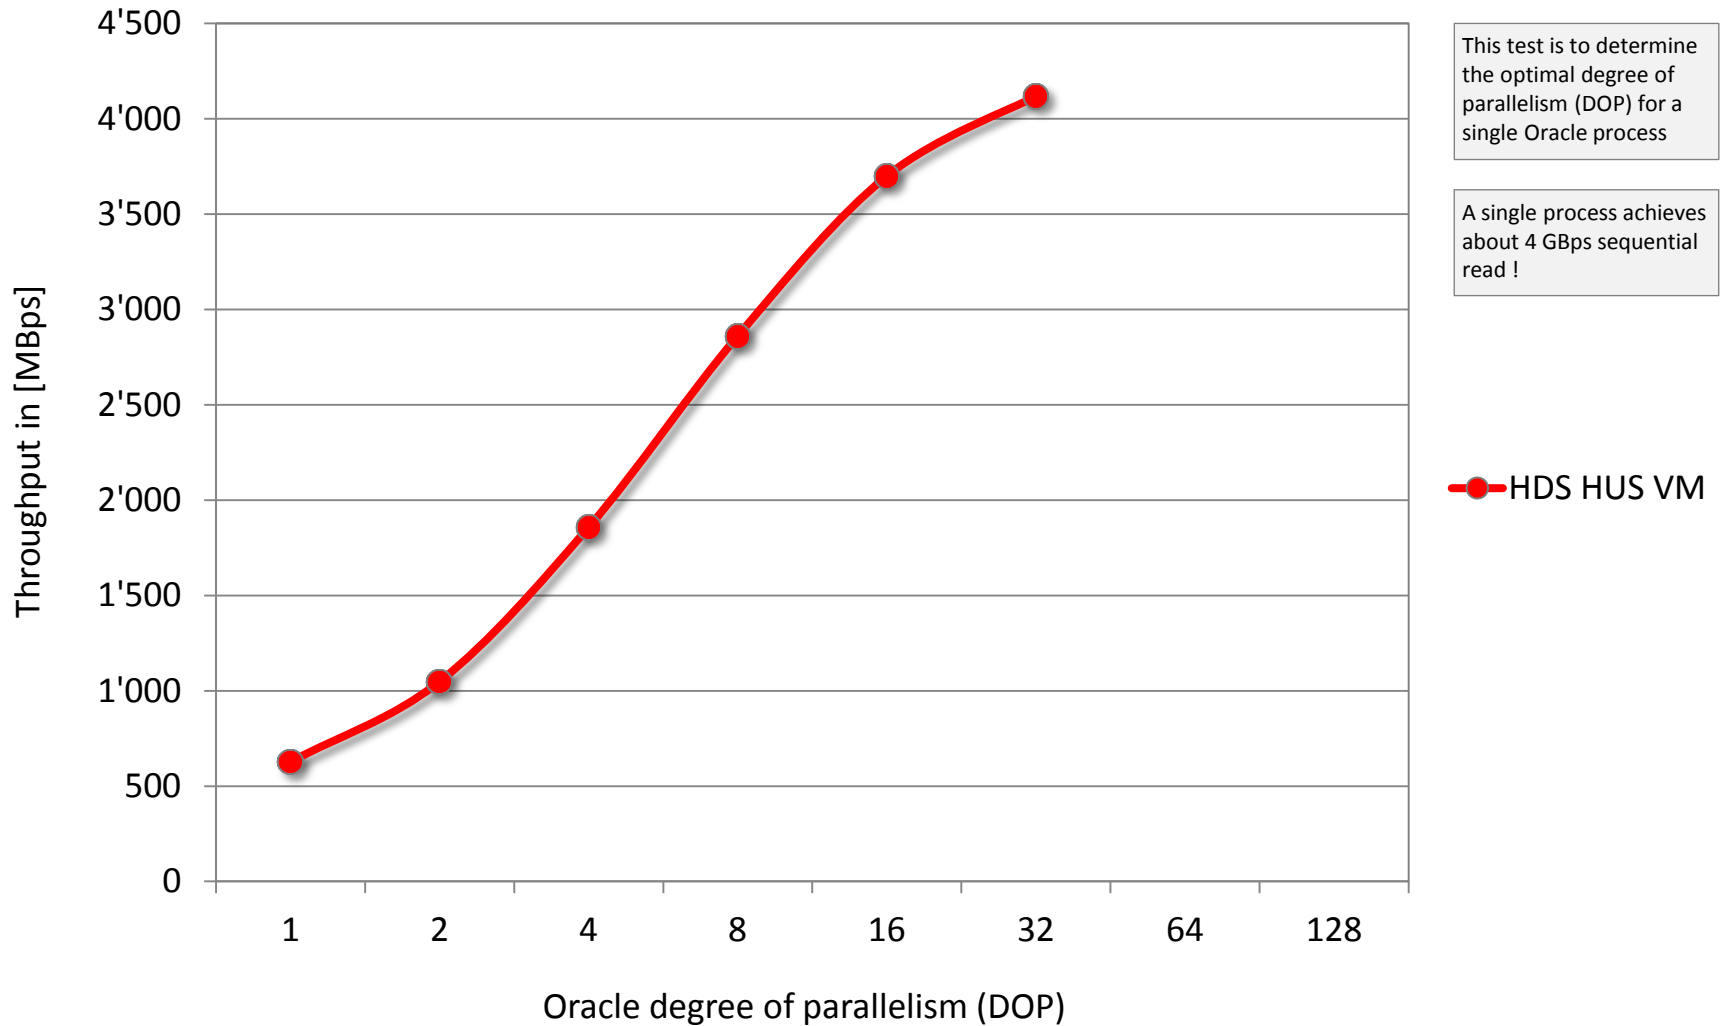


- 1 System Configuration
- 2 Introduction to Storage Performance Tests
- 3 Storage Benchmark Results – Sequential I/O**
- 4 Storage Benchmark Results – Random I/O
- 5 Introduction to Database Performance Tests
- 6 Database Benchmark Results – Database Load
- 7 Database Benchmark Results – OLTP Transactions
- 8 Reviewing Storage and Database Benchmark Results

# Storage Performance




Oracle storage performance: sequential read, single process





# Storage Performance



Oracle storage performance: sequential read, single process

HDS HUS VM  
All Flash Array


Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [bps]	Physical read [MBps]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
25	1 STO-12	4	4	1	1	0	2535	321442	2511	13	6	0	1	0	0	303
	2 STO-12	4	4	2	2	0	4226	536866	4194	11	7	0	0	0	0	311
	3 STO-12	4	4	4	4	1	7476	950941	7429	7	0	0	0	0	0	278
	4 STO-12	4	4	8	5	1	11508	1464607	11442	7	0	0	0	0	0	266
	5 STO-12	4	4	16	7	1	14885	1894018	14797	7	0	0	0	0	0	191
	6 STO-12	4	4	32	8	2	17936	2108141	16470	7	0	0	0	0	0	231

Legend:

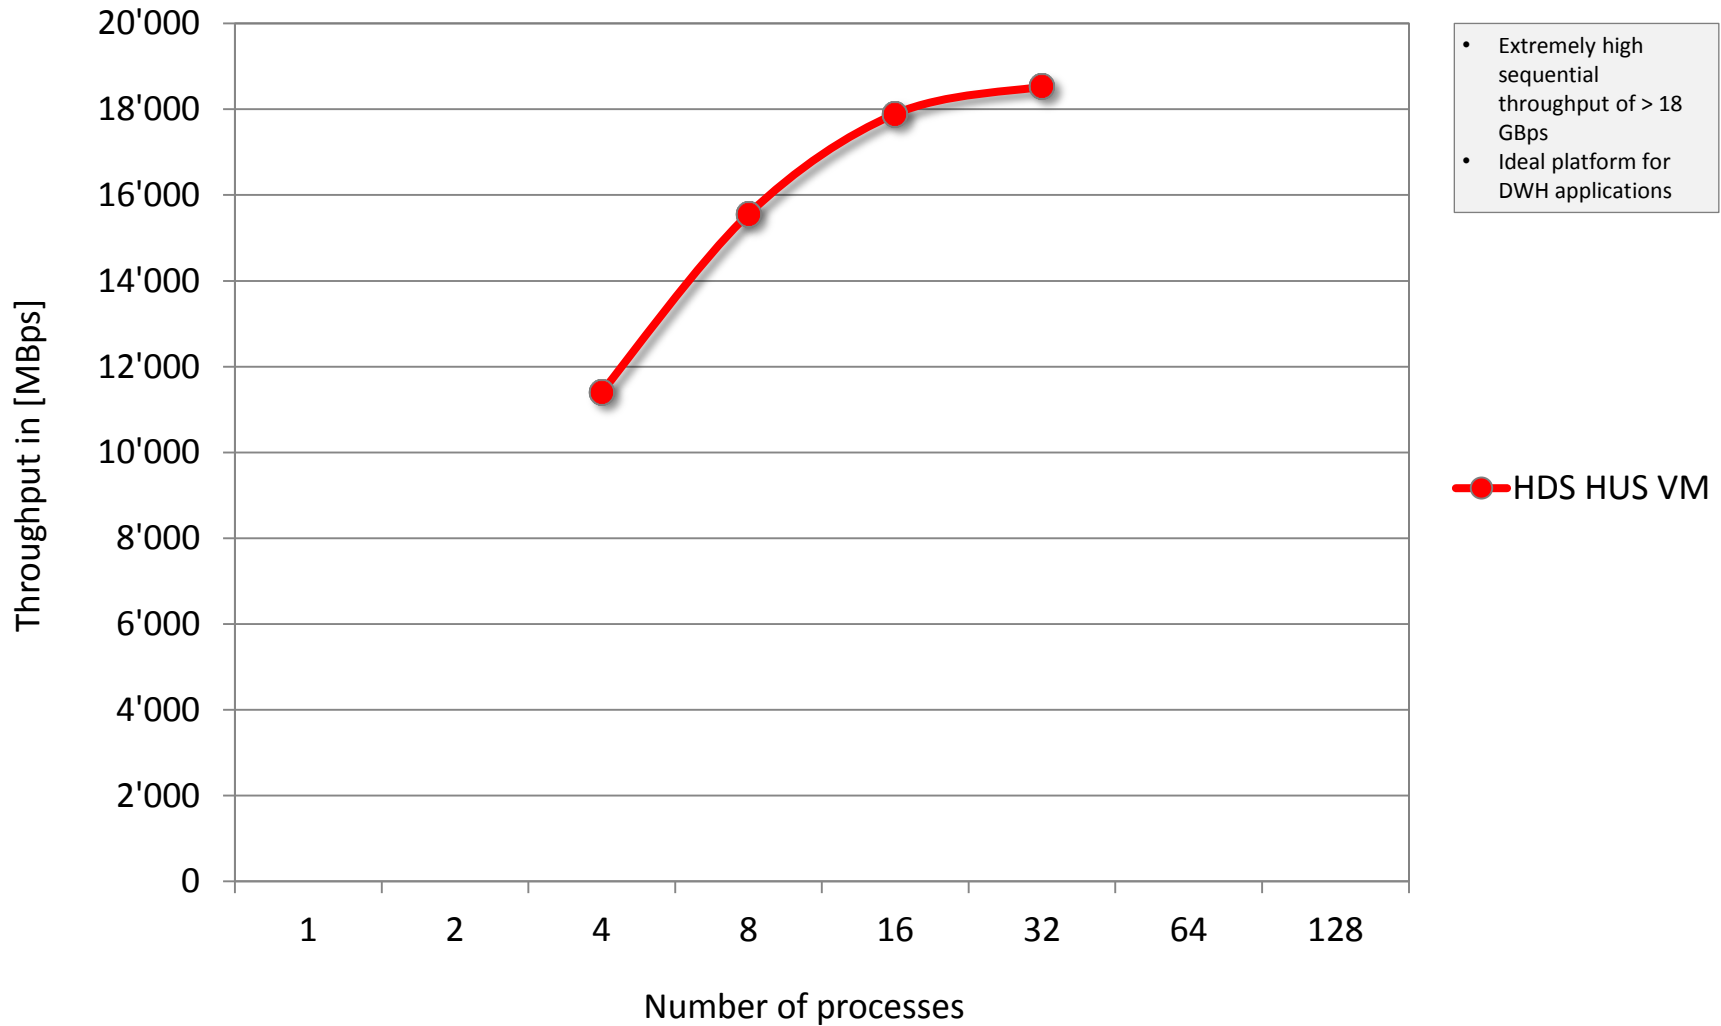
- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [s] elapsed time in seconds
- [iops] i/o operations per second
- [bps] blocks per second
- [MBps] mega byte per second

This test has been run on all four nodes – the throughput per process must therefore be divided by 4


# Storage Performance



Oracle storage performance: sequential read, multiple processes



# Storage Performance



Oracle storage performance: sequential read, multiple processes


HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [bps]	Physical read [MBps]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO write [iops]	Hitrata db flash [%]	Hitrata exa flash [%]	Elap time [s]
28	6 STO-12	4	4	8	5	1	11468	1459486	11402	7	0	0	0	0	0	286
	7 STO-12	4	8	8	7	1	15637	1990867	15554	17	12	0	1	0	0	311
	9 STO-12	4	16	8	8	2	17985	2289876	17890	17	12	0	2	0	0	319
	13 STO-12	4	32	8	8	2	18622	2371178	18525	19	11	0	2	0	0	333
	15 STO-12	4	40	8	8	2	18458	2350287	18362	18	11	0	2	0	0	333
	17 STO-12	4	48	8	8	2	17457	2222605	17364	19	11	0	2	0	0	349


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [s] elapsed time in seconds
- [iops] i/o operations per second
- [bps] blocks per second
- [MBps] mega byte per second


# Storage Performance



Oracle storage performance: sequential write



# Storage Performance




## Oracle storage performance: sequential write

HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [bps]	Physical read [MBps]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
2	1 STO-22	1	1	1	1	0	22	0	0	699	34	547	35	0	0	15
	2 STO-22	1	2	1	1	0	35	0	1	1388	68	1093	68	0	0	15
	3 STO-22	1	4	1	1	0	49	0	1	2267	113	1822	95	0	0	18
	4 STO-22	1	8	1	1	0	58	0	1	2680	136	2186	100	0	0	30
	5 STO-22	1	16	1	1	0	67	0	1	3118	167	2677	62	0	0	49
	6 STO-22	1	32	1	1	0	72	0	1	3331	206	2883	48	0	0	91
3	1 STO-22	1	48	1	1	0	77	0	1	3658	201	3225	45	0	0	122
4	1 STO-22	1	64	1	1	0	85	0	1	3940	225	3475	47	0	0	151


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [s] elapsed time in seconds
- [iops] i/o operations per second
- [bps] blocks per second
- [MBps] mega byte per second




- 1 System Configuration
- 2 Introduction to Storage Performance Tests
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O**
- 5 Introduction to Database Performance Tests
- 6 Database Benchmark Results – Database Load
- 7 Database Benchmark Results – OLTP Transactions
- 8 Reviewing Storage and Database Benchmark Results

# Storage Performance



## Oracle storage performance: random read



# Storage Performance



## Oracle storage performance: random read

HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [bps]	Physical read [MBps]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
5	84 STO-62	1	1	1	0	0	2222	2226	18	15	12	0	1	0	0	305
	85 STO-62	1	2	1	1	0	4375	4389	34	19	12	0	2	0	0	305
	86 STO-62	1	4	1	1	0	8832	8864	69	15	12	0	1	0	0	306
	87 STO-62	1	8	1	1	0	18227	18279	143	16	13	0	1	0	0	306
	88 STO-62	1	16	1	1	1	36825	36896	288	17	14	0	1	0	0	307
	89 STO-62	1	32	1	3	1	70686	70759	553	21	16	0	2	0	0	308
	90 STO-62	1	64	1	4	1	134721	134796	1053	30	20	0	3	0	0	307
	91 STO-62	1	128	1	11	7	205312	205392	1605	31	18	0	3	0	0	288
5	27 STO-62	2	256	1	14	8	416264	416434	3254	57	31	0	8	0	0	284
5	44 STO-62	3	384	1	17	8	609607	609860	4765	81	44	0	11	0	0	291
13	11 STO-62	4	736	1	27	13	806051	807746	6311	134	73	1	20	0	0	313


Legend:

#N number of RAC nodes  
 #J number of jobs  
 #T number of threads (PX)  
 [s] elapsed time in seconds


[rps] rows per second  
 [tps] transactions per second  
 [dbps] database blocks per second




# Storage Performance



## Oracle storage performance: random write



# Storage Performance



## Oracle storage performance: random write

HDS HUS VM  
All Flash Array


Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Physical read [iops]	Physical read [dbps]	Physical read [MBps]	Physical write [iops]	Physical write [dbps]	Physical write [MBps]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
29	1 STO-41	4	4	1	2	1	29	80	1	22966	11526	354	6563	0	0	306
	2 STO-41	4	8	1	4	1	46	149	1	50749	40885	895	6079	0	0	280
	3 STO-41	4	16	1	8	1	52	165	2	111746	358108	3873	2703	0	0	301
	4 STO-41	4	32	1	14	2	88	448	4	532985	591072	6405	365	0	0	316
	5 STO-41	4	64	1	15	2	104	576	5	498419	538273	6189	188	0	0	314
	6 STO-41	4	128	1	15	2	273	1669	13	491759	528837	5894	102	0	0	330

Legend:


#N number of RAC nodes  
 #J number of jobs  
 #T number of threads (PX)  
 [s] elapsed time in seconds

[rps] rows per second  
 [tps] transactions per second  
 [dbps] database blocks per second

Table size 125'000 rows




- 1 System Configuration
- 2 Introduction to Storage Performance Tests
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Introduction to Database Performance Tests**
- 6 Database Benchmark Results – Database Load
- 7 Database Benchmark Results – OLTP Transactions
- 8 Reviewing Storage and Database Benchmark Results



Why measure Database performance?

- Projects need understandable key performance metrics for capacity planning
  - Data load
  - Data scan
  - Data aggregation
  - OLTP transactions
  - Time windows for certain operations


# Database Performance



What is measured?

- Typical Oracle database operations
- Speed of single thread
  - Rows per second [rps]
  - Transactions per second [tps]
  - SQL response time [s]
- Maximum throughput of system
  - Rows per second [rps]
  - Transactions per second [tps]
  - SQL response time [s]
  - REDO rate [MBps]
  - REDO service time [s]
- Scalability
  - Throughput per process for  $n = \{1, 2, 4, 8, \dots, n\}$
- Efficiency of
  - All platform layers

# Database Performance




## Overview of Database load performance tests with Test Codes

Database Performance Data load un-compressed	Test Code for Data Load via buffer cache	Test Code for Data Load direct	
▪ Conventional data load (LGWR commit stress test)	DBL-11	-	-
▪ Bulk load	-	DBL-21	-

Database Performance Data load compressed			Test Code for Data Load direct compressed
▪ Bulk load, compress BASIC			DBL-31
▪ Bulk load, compress OLTP			DBL-32
▪ Bulk load, compress HCC query low			DBL-33
▪ Bulk load, compress HCC archive low			DBL-34

# Database Performance




## Overview of Database aggregation performance tests with Test Codes

Database Performance Data Aggregation	Test Code for Data on default Storage		
▪ Create unique b-tree index	DBA-12		
▪ Create non-unique b-tree index	DBA-22		

Database Performance Data Scan	Test Code for Data on default Storage	Test Code for Data in Flash Cache	Test Code for Data in Cell Flash Cache
▪ Full table scan	DBS-12	DBS-13	DBS-14


# Database Performance



## Overview of Database OLTP performance tests with Test Codes


Database Performance OLTP Transactions 1 hit per transaction	Test Code for Data on default Storage	Test Code for Data in Flash Cache	Test Code for Data in Cell Flash Cache
▪ Select transaction	DBX-12	DBX-13	DBX-14
▪ Update transaction	DBX-22	DBX-23	DBX-24






Remarks on other benchmark tools . . .

- SAP, TPC, Swingbench, Hammerora, ...
  - No specific database performance metrics for data scan
  - No specific database performance metrics for data load
  - No specific database performance metrics for data aggregation
  - No explicit support for Oracle flash cache or Oracle cell flash cache
  - No explicit support for Oracle compression techniques

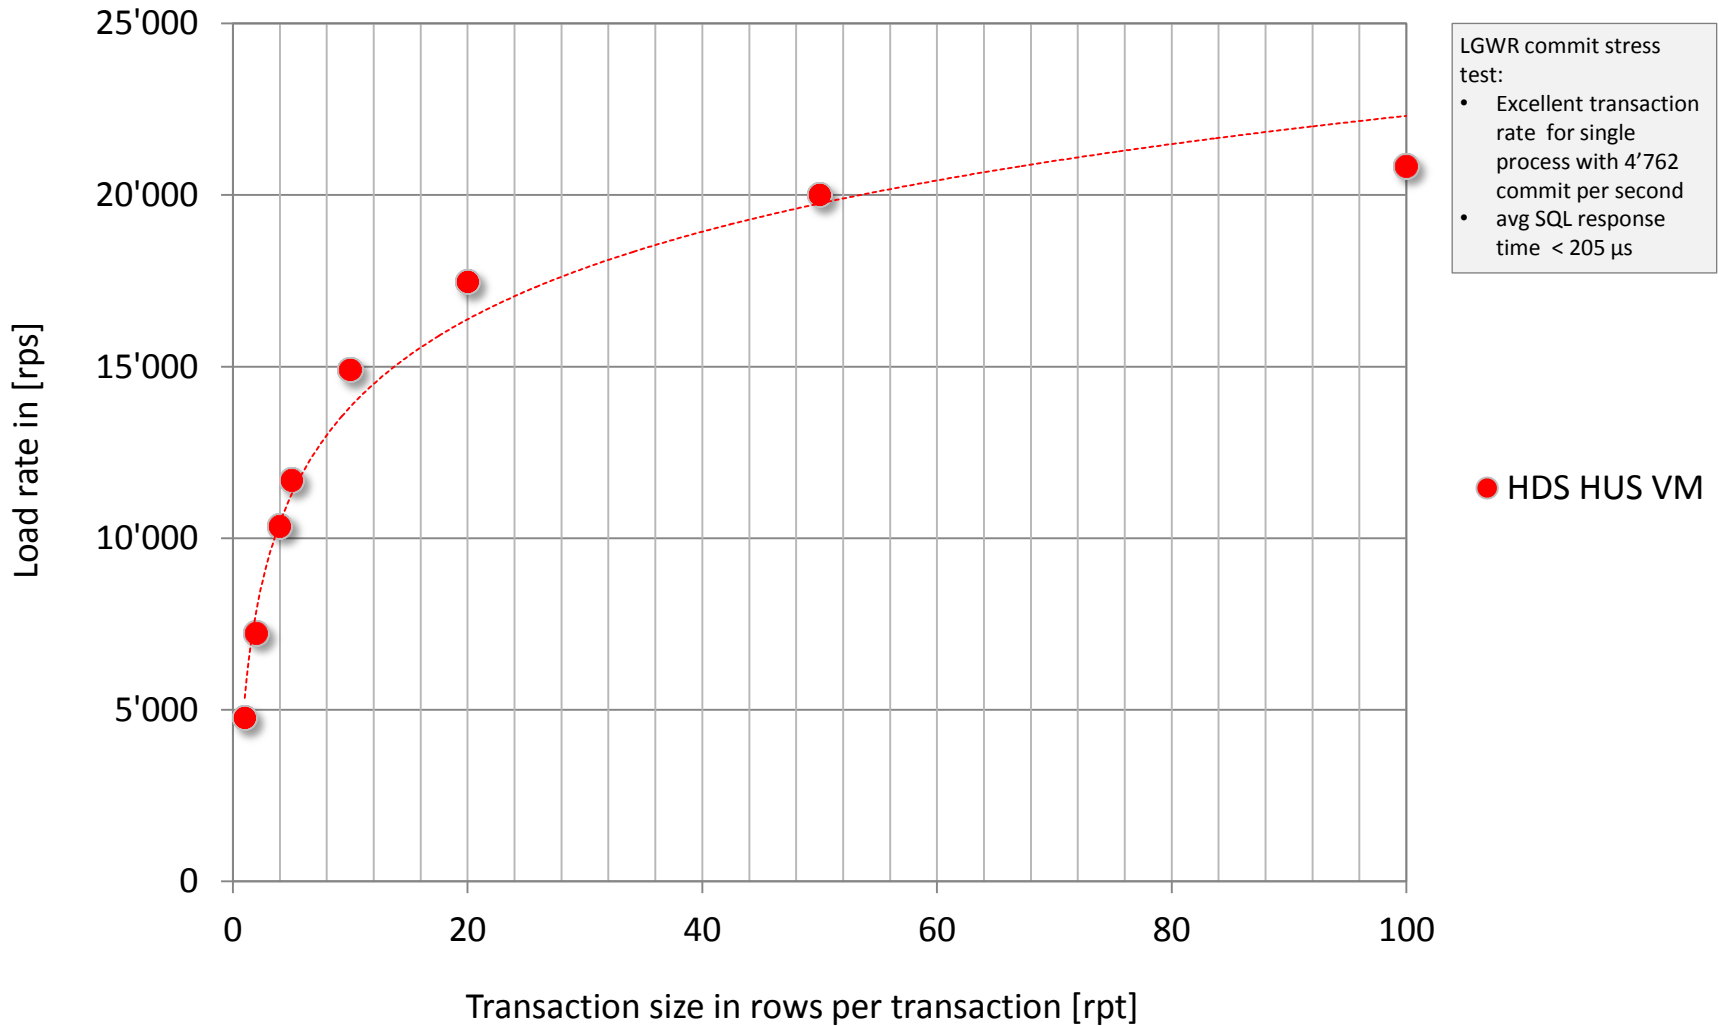


- 1 System Configuration
- 2 Introduction to Storage Performance Tests
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Introduction to Database Performance Tests
- 6 Database Benchmark Results – Database Load**
- 7 Database Benchmark Results – OLTP Transactions
- 8 Reviewing Storage and Database Benchmark Results


# Database Performance



Database transactional load, single process, different transaction size



# Database Performance



Database transactional load, single process, different transaction size


HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	TX size [rpt]	CPU busy [%]	Throughput rows/sec [rps]	Throughput txn/sec [tps]	SQL service time [s]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO size [MBps]	REDO writes [iops]	REDO svt [ms]	REDO sync writes	REDO sync svt [us]	Elap time [s]
43	1 DBL-11	1	1	1	1	1	4.762E+03	4.762E+03	2.052E-04	6206	640	33	8	2026	119	2	1114	315
	2 DBL-11	1	1	1	2	1	7.222E+03	3.611E+03	2.687E-04	5944	847	39	10	1929	86	3	499	225
	3 DBL-11	1	1	1	4	1	1.035E+04	2.588E+03	3.678E-04	5394	1171	48	12	1716	62	1	2105	157
	4 DBL-11	1	1	1	5	1	1.169E+04	2.338E+03	4.105E-04	5187	1306	53	13	1651	56	1	766	139
	5 DBL-11	1	1	1	10	1	1.491E+04	1.491E+03	6.320E-04	4452	1607	63	16	1376	46	1	1415	109
	6 DBL-11	1	1	1	20	1	1.747E+04	8.740E+02	1.054E-03	3079	1829	70	18	872	44	2	604	93
	7 DBL-11	1	1	1	50	1	2.006E+04	4.010E+02	2.299E-03	1813	2053	78	20	409	52	1	472	81
	8 DBL-11	1	1	1	100	1	2.083E+04	2.080E+02	4.378E-03	1167	2157	80	20	214	70	1	515	78


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [iops] i/o operations per second
- [s] time in seconds
- [ms] time in milli seconds
- [us] time in micro seconds


# Database Performance



Database transactional load, 2 rows per transaction



# Database Performance



Database transactional load, 2 rows per transaction


HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	TX size [rpt]	CPU busy [%]	Throughput rows/sec [rps]	Throughput txn/sec [tps]	SQL service time [s]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO size [MBps]	REDO writes [iops]	REDO svt [ms]	REDO sync writes	REDO sync svt [us]	Elap time [s]
43	9 DBL-11	1	1	1	2	1	7.287E+03	3.643E+03	2.650E-04	5969	895	40	10	1934	85	1	3147	223
	10 DBL-11	1	2	1	2	1	1.484E+04	7.420E+03	2.595E-04	10551	1967	81	20	1978	89	5	330	219
	11 DBL-11	1	4	1	2	2	3.023E+04	1.512E+04	2.578E-04	14276	4170	164	41	1642	109	5	727	215
	12 DBL-11	1	8	1	2	3	5.856E+04	2.928E+04	2.607E-04	17877	9646	326	80	996	194	9	981	222
	13 DBL-11	1	16	1	2	5	1.074E+05	5.372E+04	2.703E-04	17083	35605	734	147	455	473	16	653	242
	14 DBL-11	1	32	1	2	11	1.857E+05	9.286E+04	3.218E-04	83686	126915	1781	254	241	1059	1179	25916	280
	15 DBL-11	1	64	1	2	16	2.477E+05	1.238E+05	4.686E-04	145357	167971	2366	340	127	2158	76897	28997	323
	22 DBL-11	2	128	1	2	33	5.496E+05	2.748E+05	3.922E-04	133501	155688	3541	749	401	1380	22529	28818	320
	29 DBL-11	4	256	1	2	34	5.373E+05	2.686E+05	7.685E-04	324858	373833	5209	737	500	2540	#####	74995	352

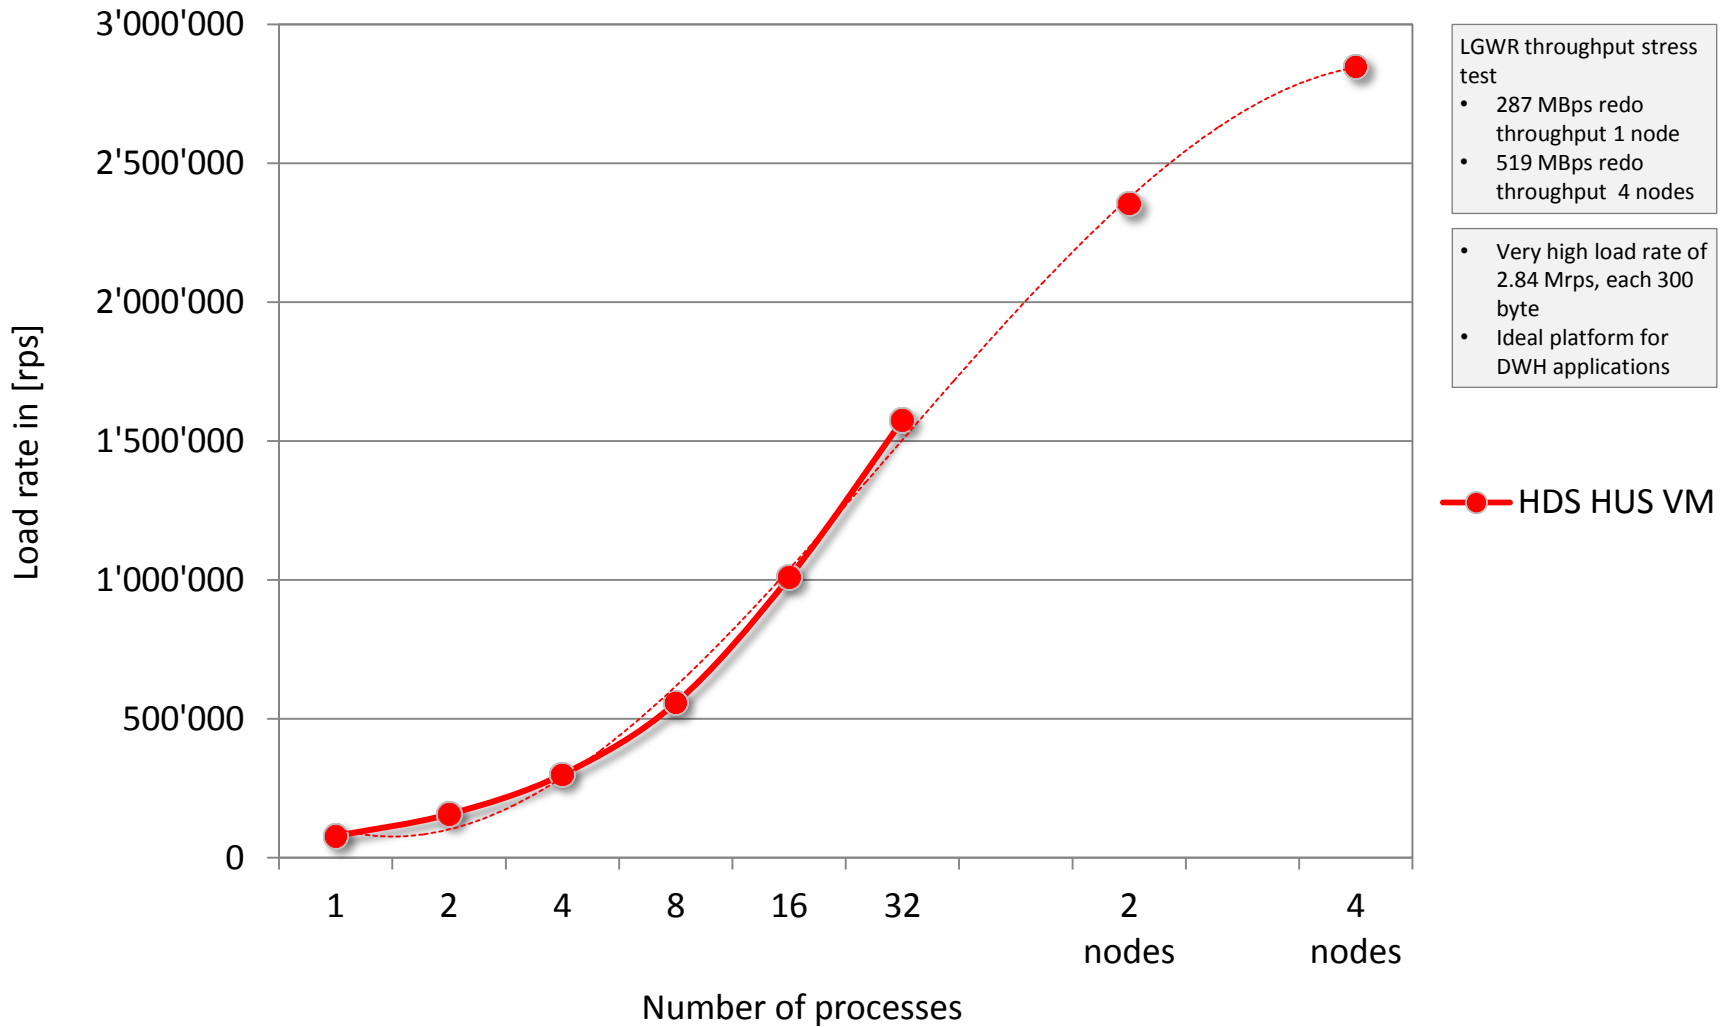
Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [iops] i/o operations per second
- [s] time in seconds
- [ms] time in milli seconds
- [μs] time in micro seconds

# Database Performance



## Database bulk load



# Database Performance



## Database bulk load


HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	TX size [rpt]	CPU busy [%]	Throughput rows/sec [rps]	Throughput txn/sec [tps]	SQL service time [s]	Physical write [iops]	Physical write [bps]	Physical write [MBps]	REDO size [MBps]	REDO writes [iops]	REDO svt [ms]	REDO sync writes	REDO sync svt [us]	Elap time [s]
44	1 DBL-21	1	1	1	0	1	7.767E+04	0.000E+00	9.937E+01	519	7036	99	14	21	661	2	434	206
	2 DBL-21	1	2	1	0	1	1.561E+05	0.000E+00	9.988E+01	964	14051	198	28	29	648	3	894	205
	3 DBL-21	1	4	1	0	2	2.991E+05	0.000E+00	1.030E+02	2094	27739	386	54	42	749	6	2322	214
	4 DBL-21	1	8	1	0	3	5.565E+05	0.000E+00	1.100E+02	4162	54027	736	101	49	1116	8	3670	230
	5 DBL-21	1	16	1	0	5	1.008E+06	0.000E+00	1.221E+02	14974	98745	1340	184	40	1861	17	11245	254
	6 DBL-21	1	32	1	0	8	1.575E+06	0.000E+00	1.540E+02	41355	148529	2049	287	44	2746	32	26631	325
12	DBL-21	2	64	1	0	12	2.354E+06	0.000E+00	1.986E+02	35881	227379	3105	429	159	2593	66	7887	435
18	DBL-21	4	128	1	0	14	2.848E+06	0.000E+00	3.447E+02	41633	273747	3746	519	241	5783	10943	533631	719

Legend:


- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [iops] i/o operations per second
- [s] time in seconds
- [ms] time in milli seconds
- [μs] time in micro seconds






- 1 System Configuration
- 2 Introduction to Storage Performance Tests
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Introduction to Database Performance Tests
- 6 Database Benchmark Results – Database Load
- 7 Database Benchmark Results – OLTP Transactions**
- 8 Reviewing Storage and Database Benchmark Results


# Database Performance



Oracle OLTP select performance, 1 row per transaction



# Database Performance



Oracle OLTP select performance, 1 row per transaction


HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Throughput rows/sec [rps]	Throughput txn/sec [tps]	SQL service time [s]	Physical read [iops]	Physical write [iops]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Physical read [MBps]	Physical write [MBps]	Elap time [s]
7	1 DBX-12	1	1	1	0	0	8.050E+02	8.050E+02	1.194E-03	1652	50	0	0	0	13	1	82
	2 DBX-12	1	2	1	1	0	1.535E+03	1.535E+03	1.224E-03	3101	63	1	0	0	24	1	86
	3 DBX-12	1	4	1	1	0	3.070E+03	3.070E+03	1.228E-03	6071	65	1	0	0	48	1	86
	4 DBX-12	1	8	1	1	0	6.212E+03	6.212E+03	1.185E-03	11982	55	1	0	0	94	1	85
	5 DBX-12	1	16	1	2	1	1.427E+04	1.427E+04	1.051E-03	26563	78	2	0	0	208	1	74
	6 DBX-12	1	32	1	4	1	3.352E+04	3.352E+04	8.679E-04	59094	86	12	0	0	462	1	63
	7 DBX-12	1	64	1	6	2	6.925E+04	6.925E+04	8.052E-04	111404	110	8	0	0	871	1	61
	8 DBX-12	1	128	1	10	3	1.320E+05	1.320E+05	8.139E-04	186584	123	13	0	0	1458	1	64
	9 DBX-12	1	256	1	19	9	1.690E+05	1.690E+05	1.280E-03	208241	96	17	0	0	1627	1	100
	10 DBX-12	1	512	1	21	7	2.209E+05	2.209E+05	1.884E-03	246663	109	20	0	0	1927	1	153
	20 DBX-12	2	1024	1	38	11	4.305E+05	4.305E+05	1.912E-03	480650	173	49	0	0	3755	1	157
8	5 DBX-12	4	2048	1	50	9	6.174E+05	6.174E+05	2.774E-03	660026	195	57	0	0	5157	1	272


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [iops] i/o operations per second
- [s] time in seconds


# Database Performance



Oracle OLTP update performance, 1 row per transaction



# Database Performance




## Oracle OLTP update performance, 1 row per transaction

HDS HUS VM  
All Flash Array

Run	Tst Code	#N	#J	#T	CPU busy [%]	CPU sys [%]	Throughput rows/sec [rps]	Throughput txn/sec [tps]	SQL service time [s]	Physical read [iops]	Physical write [iops]	REDO write [iops]	Hitrate db flash [%]	Hitrate exa flash [%]	Physical read [MBps]	Physical write [MBps]	Elap time [s]
12	1 DBX-22	1	1	1	1	0	7.380E+02	7.380E+02	1.269E-03	1563	3018	739	0	0	12	9	91
	2 DBX-22	1	2	1	1	0	1.493E+03	1.493E+03	1.263E-03	3061	5881	1434	0	0	24	17	90
	3 DBX-22	1	4	1	1	0	3.126E+03	3.126E+03	1.201E-03	6255	9244	1903	0	0	49	34	86
	4 DBX-22	1	8	1	2	1	6.556E+03	6.556E+03	1.135E-03	12803	14356	1869	0	0	100	68	82
	5 DBX-22	1	16	1	3	1	1.453E+04	1.453E+04	1.027E-03	27265	24587	1697	0	0	213	148	74
	6 DBX-22	1	32	1	4	1	3.029E+04	3.029E+04	9.633E-04	53683	42125	1354	0	0	420	305	71
	7 DBX-22	1	64	1	6	1	5.060E+04	5.060E+04	1.107E-03	81537	60329	614	0	0	637	511	85
	8 DBX-22	1	128	1	9	2	7.680E+04	7.680E+04	1.512E-03	108433	84646	354	0	0	847	782	112
	9 DBX-22	1	256	1	10	2	8.559E+04	8.559E+04	2.705E-03	105432	92928	266	0	0	824	878	201
19	DBX-22	2	512	1	13	2	1.139E+05	1.139E+05	4.099E-03	140259	125739	194	0	0	1096	1198	302
29	DBX-22	4	1024	1	17	3	1.183E+05	1.183E+05	8.011E-03	163170	128838	276	0	0	1275	1227	325
30	DBX-22	4	2048	1	18	3	1.207E+05	1.207E+05	1.429E-02	162494	131039	352	0	0	1270	1247	358


Legend:

- #N number of RAC nodes
- #J number of jobs
- #T number of threads (PX)
- [rps] rows per second
- [tps] transactions per second
- [iops] i/o operations per second
- [s] time in seconds



- 1 System Configuration
- 2 Introduction to Storage Performance Tests
- 3 Storage Benchmark Results – Sequential I/O
- 4 Storage Benchmark Results – Random I/O
- 5 Introduction to Database Performance Tests
- 6 Database Benchmark Results – Database Load
- 7 Database Benchmark Results – OLTP Transactions
- 8 Reviewing Storage and Database Benchmark Results**

# Benchmark Results



## Summary Storage Performance

	Metric		HDS HUS VM All Flash Array
Sequential I/O			
▪ Read	[MBps]		18'525
▪ Write	[MBps]		3'475 (6'400)
Random I/O read			
▪ Throughput	[IOPS]		806'000
▪ Service time	[ms]		0.840
Random I/O write			
▪ Throughput	[IOPS]		532'985
▪ Throughput	[dbps]		591'072
▪ Service time, DBWR multi block	[ms]		1.1
▪ Service time, LGWR multi block	[ms]		10

Legend:


[MBps] mega byte per second

[IOPS] I/O operations per second

[dbps] database blocks per second

[ms] milli seconds

# Benchmark Results




## Summary Database Performance

Database Performance Data Load	Metric		HDS UCP & HDS HUS VM All Flash Array
Un-Compressed			
▪ Conventional OLTP insert [rps]	[rps]		549'000
▪ Bulk load [rps]	[rps]		2'840'000

Database Performance OLTP Transactions 1 hit per transaction	Metric		HDS UCP & HDS HUS VM All Flash Array
Select transaction			
▪ Throughput	[tps]		617'000
▪ Service time	[ms]		2.8
Update transaction			
▪ Throughput	[tps]		120'000
▪ Service Time	[ms]		14




# Benchmark Results



## Reviewing Storage and Database Performance

- HDS HUS VM All Flash Array combines the best of all worlds
  - Seamless integration in existing SAN infrastructures
  - Proven scalability and performance for all workloads
  - Extremely high I/O throughput with microseconds service time
  - Highly efficient DRAM storage cache in addition to flash technology
  - Rich and mature storage management software portfolio (cloning, snapshots, replication, dynamic provision, dynamic tiering etc.)
  
- Superior Oracle Database Platform based on HDS UCP & HUS VM
  - Extremely efficient transactional and bulk load
  - Highly scalable transaction rates
  - Consistent performance for ALL database objects

# Benchmark Results



## Reviewing Storage and Database Performance

- Closing the I/O gap
  - Gap is widening between high performance servers and slow storage
  - Extremely high I/O performance and low response time is required to maximize utilization of server processing power
  - HUS VM All Flash Array reduces the I/O gap and allows Oracle platforms to fully utilize cpu capacity without waiting for I/O operations
  - Increasing server cpu utilization improves return on assets for customers and lowers CAPEX for server platforms

**BENCHWARE**

*swiss precision in performance measurement*

*[www.benchmarkware.ch](http://www.benchmarkware.ch)*

*[info@benchmarkware.ch](mailto:info@benchmarkware.ch)*