Oracle Database 12*c* Performance Characterization

Tegile IntelliFlash T4700 All-Flash Storage Array

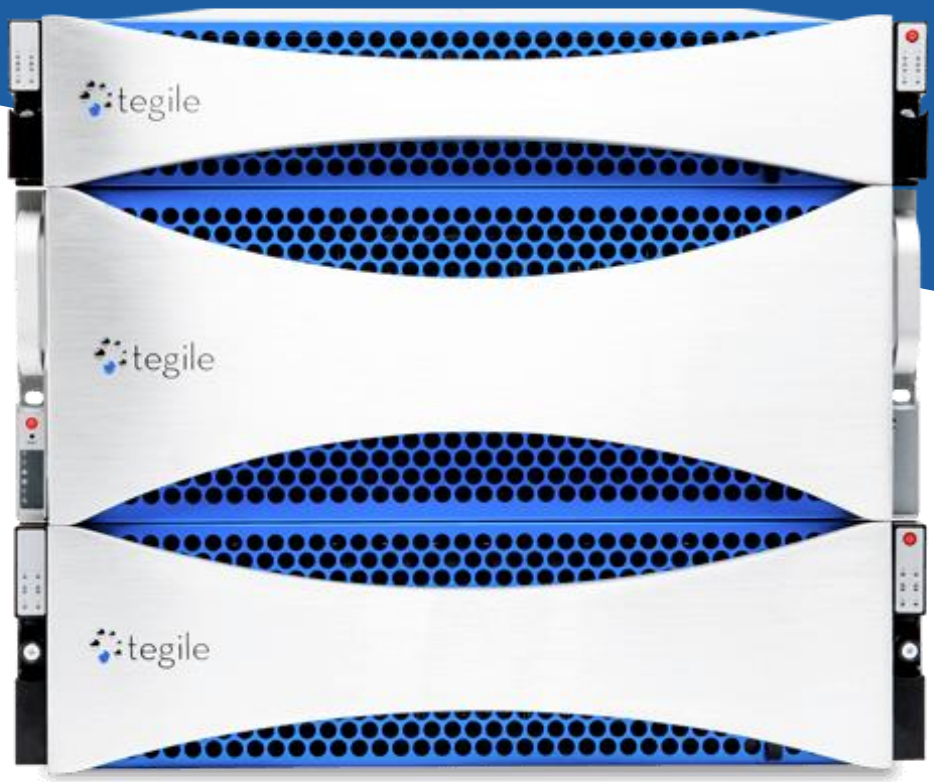

Executive Summary

This white paper documents testing to characterize the performance of Oracle Database 12*c* on the Tegile IntelliFlash T4700 All-Flash Storage Array using the HammerDB tool with a TPC-C workload.

For these tests, an aggregate HammerDB score of **8,392,994 Oracle TPM** and **2,877,715 NOPM** was achieved using two Oracle instances connected to separate controllers on the Tegile storage array.

Node A score of 4,189,540 Oracle TPM and 1,435,474 NOPM.
Node B score of 4,203,454 Oracle TPM and 1,442,241 NOPM.

These are not official TPC-C results.

These results highlight the T4700 storage array as a high-performance solution for Oracle Database 12*c* deployments, and demonstrate the benefits of Tegile's active/active controller technology.

Tegile all-flash arrays scale to meet the performance requirements of a wide range of database workloads, providing very low latency to dramatically reduce I/O wait times. Furthermore, all Tegile arrays are built with enterprise class SSDs with 10X more endurance, ensuring sustained performance over time.

These results demonstrate outstanding performance and provide proof that customers can deploy Oracle Database 12*c* on Tegile all-flash arrays with confidence.

Disclaimers

- The objective of this testing is to measure the performance of Oracle Database 12*c* on Tegile all-flash storage arrays for internal reference and informational purposes.
- These test results were NOT audited and approved by TPC council or HammerDB support.
- Some tuning parameters used to maximize the database performance in order to stress the T4700 storage array are not typically used in production environments. This is a common practice for all TPC testing.
- Comparative HammerDB results from other vendors on public websites may be removed or changed at any time.

Understanding HammerDB (www.hammerdb.com/about)

HammerDB is a graphical open-source database load testing and benchmarking tool for Windows and Linux platforms. Originally developed for Oracle Database, HammerDB can now be run with many different databases. The HammerDB tool creates a test schema, loads it with data, and simulates a workload of multiple virtual uses against the database.

HammerDB includes a workload derived from the well-known TPC-C benchmark. HammerDB does not implement the full TPC-C benchmark. Official Audited TPC-C benchmarks are extremely costly, time consuming and complex to establish and maintain. HammerDB takes the essence of this benchmark and implements it in a way that can be run at a fraction of the cost. Unlike the TPC-C benchmark, HammerDB does not have a formal auditing procedure for publication. For this reason, HammerDB results should not be compared against published TPC-C benchmarks.

HammerDB produces two primary metrics called Transactions Per Minute (TPM) and New Orders Per Minute (NOPM).

Understanding TPC-C (www.tpc.org/tpcc)

TPC-C is an industry-standard benchmark which generates an On-Line Transaction Processing (OLTP) workload that simulates a wholesale supplier with sales districts and associated warehouses. A warehouse covers 10 districts, and a district covers 3,000 customers. All warehouses maintain the entire stock of 100,000 items sold by the company.

The TPC-C business flow includes order entry, order processing, order delivery, payment recording, order status checking, and warehouse inventory management. Orders are processed from the local warehouse but can be processed from other warehouses in case of insufficient stock. Order cancellations and sales returns are possible. While the benchmark portrays the activity of a wholesale supplier, TPC-C is not limited to any particular business segment. Rather, it represents any industry that must manage, sell, or distribute a product or service.

From a database perspective, the TPC-C benchmark involves a mix of five concurrent transactions of different types and complexity. The database is comprised of nine types of tables with a wide range of record and population sizes.

TPC-C results are measured in transactions per minute (tpmC).

Benchmark Topology

The system under test has the following configuration:

Software Environment

Oracle Linux 7.3 x86_64 with the Unbreakable Enterprise Kernel 4.1.12-94.5.9.el6uek.x86_64

Oracle Database 12*c* Enterprise Edition Release 12.2.0.1.0 – 64 bit Production

Oracle Grid Infrastructure 12*c* Release 1 (12.2.0.1.0) for Linux x86-64

HammerDB Release 2.23 for Linux 64-bit

Hardware Environment

Figure 1 - Tegile IntelliFlash T4700 All-Flash Storage Array

1 x Tegile IntelliFlash T4700 All-Flash Storage Array with

4 x Intel® Xeon® E5-2680 v3 @ 2.50 GHz (12 cores per processor / 48 cores total)

464 GB memory

24 x 1 TB SSD

IntelliFlash version 3.7.0.0.170808(GA)

Figure 2 - Dell PowerEdge R930 Server

2 x Dell PowerEdge R930 Server with

- 4 x Intel® Xeon® E7-4830 v3 @ 2.10 GHz with multithreading enabled
- (12 cores per processor / 48 cores total / 96 threads)
- 768 GB memory (48 x 16 GB DDR-4 dual-rank DIMMs @ 1333 MHz)
- 2 x 278 GB internal SAS HDD
- 1 x Integrated Dell PowerEdge RAID controller (PERC) H730P adapter (embedded)
- 1 x Integrated Broadcom BRCM5720 GbE quad-port Ethernet controller (embedded)
- 2 x QLogic QLE2662 16 Gb dual-port fibre channel adapter
- BIOS version 1.1.5
- Firmware version 2.21.21.21

Figure 3 - Brocade 6505 Fibre Channel Switch

1 x Brocade 6505 16 Gbps 24-port fibre channel switch

Figure 4 - Cisco Nexus 5020 Ethernet Switch

1 x Cisco Nexus 5020 10 Gbps Ethernet switch

BIOS Settings

In the Dell System BIOS Settings, Node Interleaving was disabled for non-uniform memory access (NUMA).

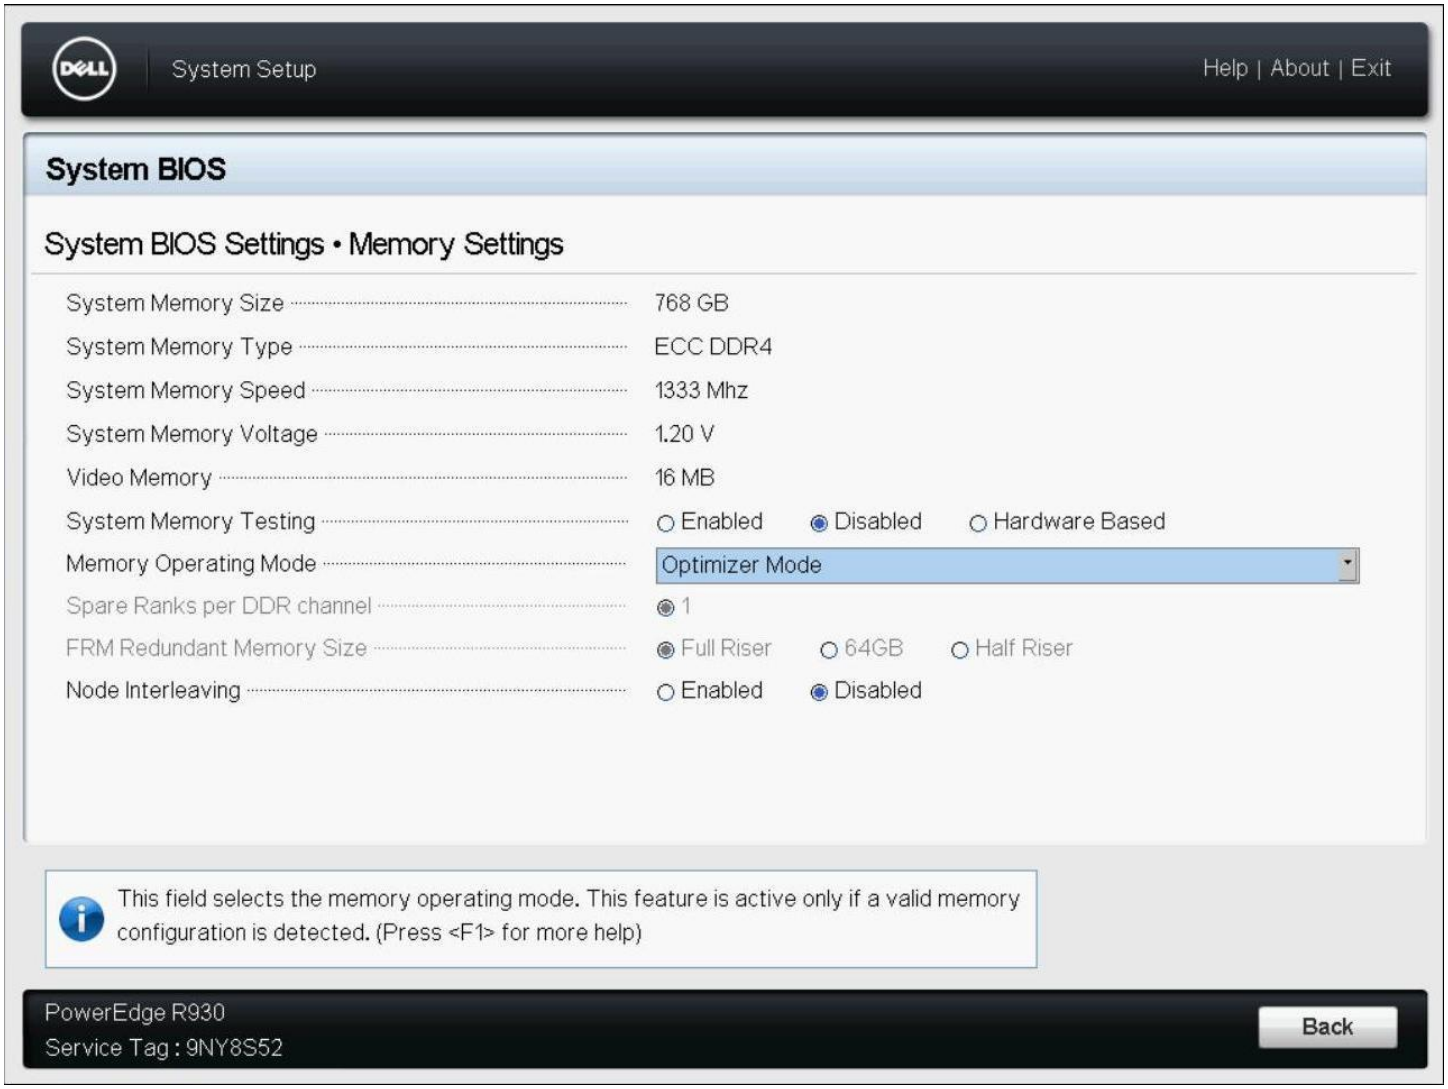

Figure 5 - System BIOS Memory Settings

The Performance System Profile was selected to maximize processor frequency and disable certain power-saving features such as C-states.

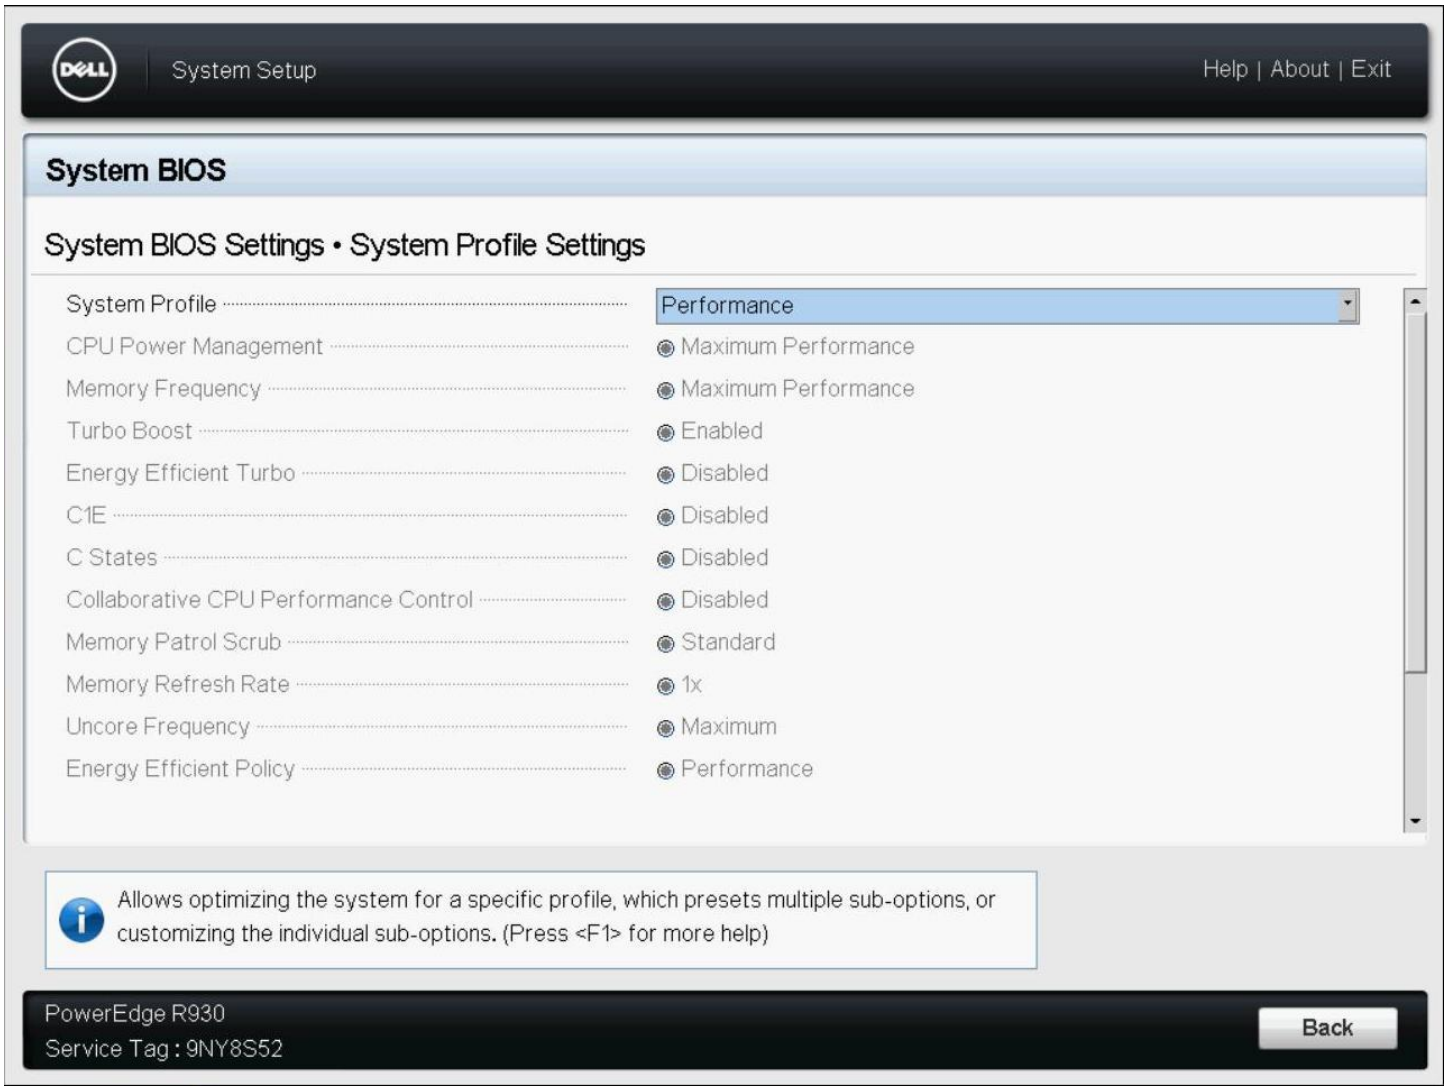

Figure 6 - System BIOS Profile Settings

Linux Kernel Settings

For optimal performance, Oracle Linux was configured with the following kernel settings in /etc/sysctl.conf as follows:

```
fs.aio-max-nr = 1048576
fs.file-max = 6815744
kernel.sem = 250 32000 100 128
kernel.shmmni = 4096
kernel.shmall = 1073741824
kernel.shmmax = 4398046511104
kernel.panic_on_oops = 1
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
net.ipv4.conf.all.rp_filter = 2
net.ipv4.conf.default.rp_filter = 2
net.ipv4.ip_local_port_range = 9000 65500
```

Storage Configuration

The T4700 storage array was configured with the following settings for all Storage Pools, Project, and LUNs:

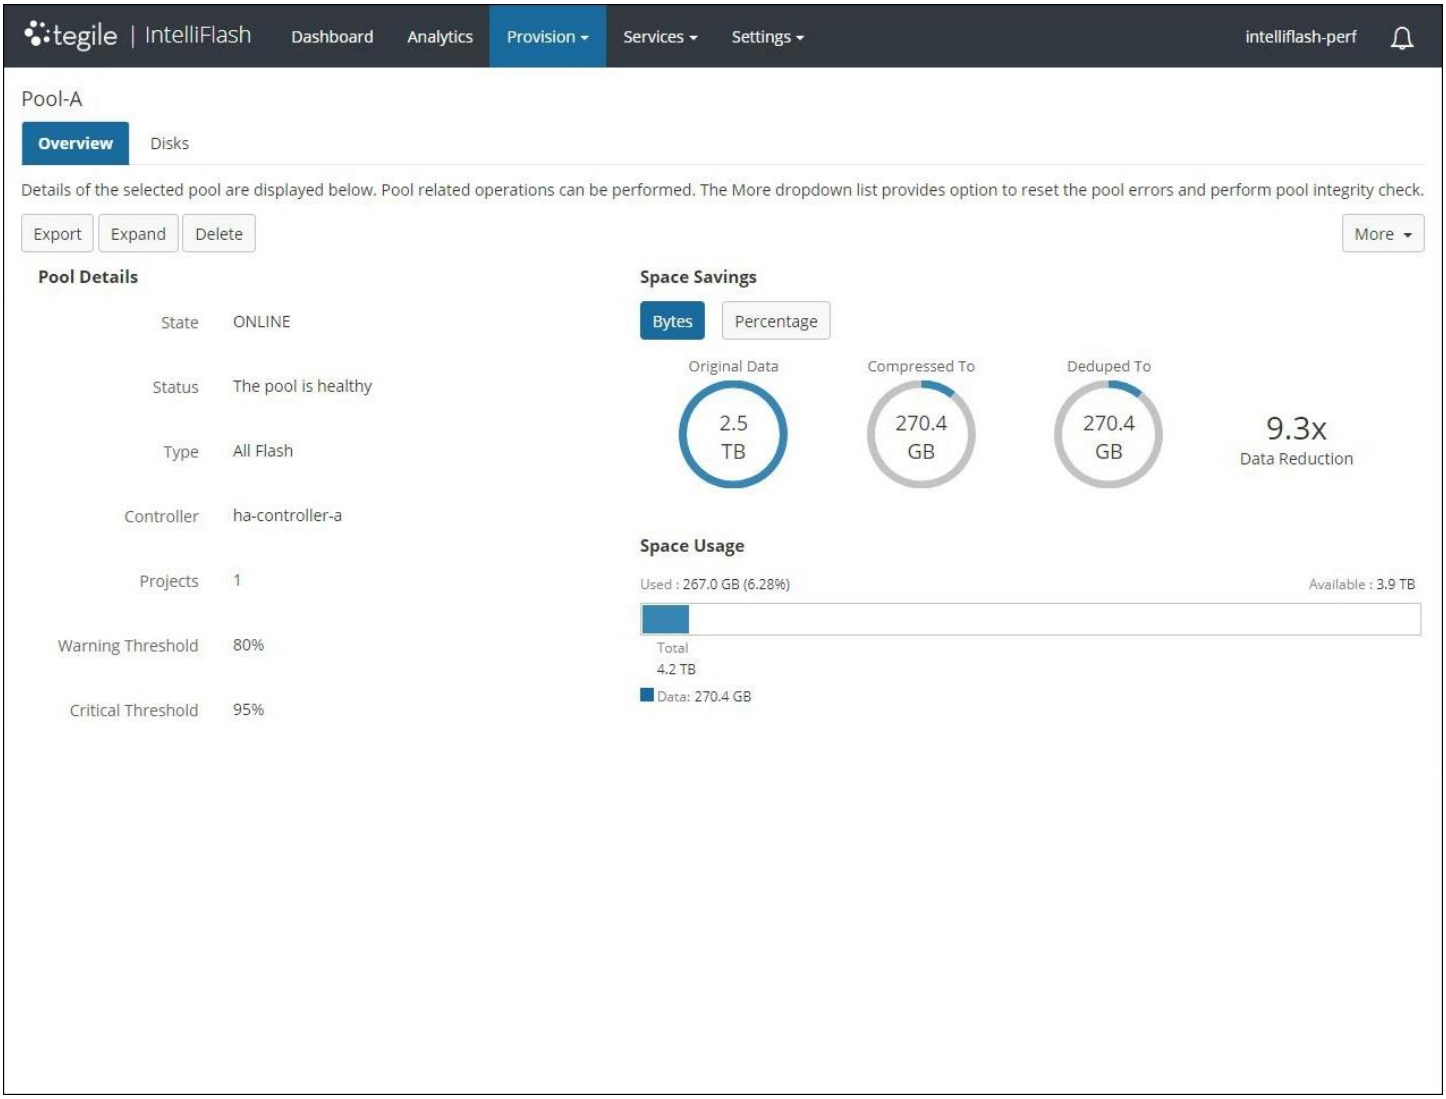

Figure 7 - Pool Settings

Pool A Controller Configuration: Active/Active
Disks: 12 x 1 TB SSD
RAID: 2-way mirror
Resource Group A

Pool B Controller Configuration: Active/Active
Disks: 12 x 1 TB SSD
RAID: 2-way mirror
Resource Group B

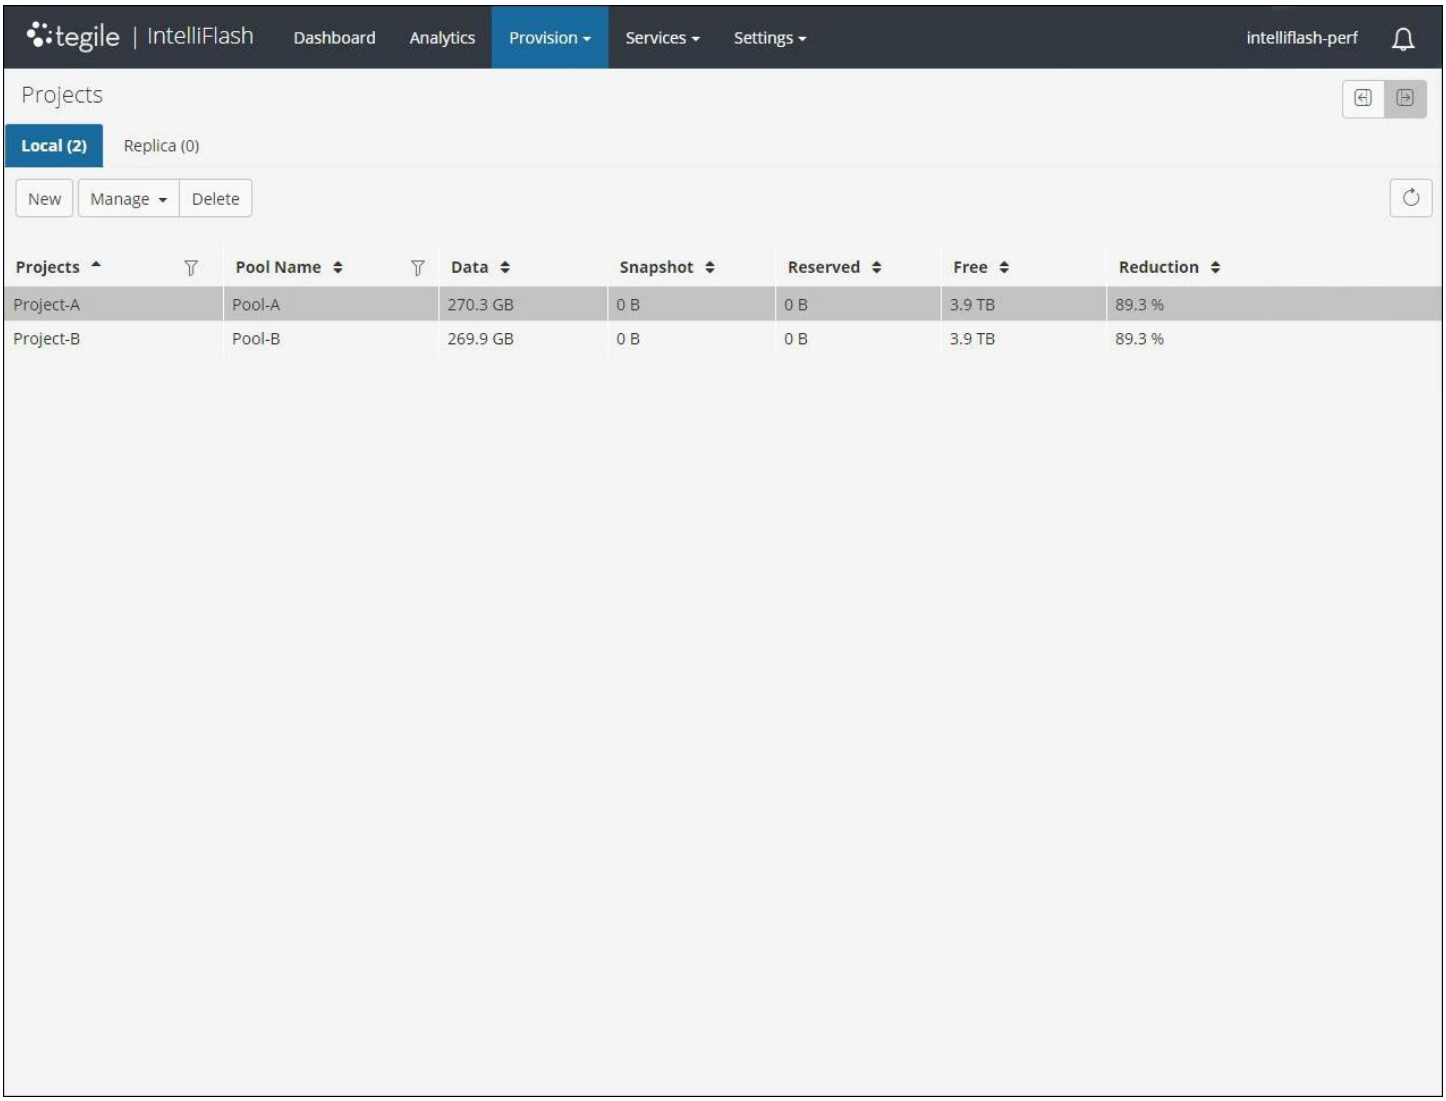

Figure 8 - Project Settings

Project A

Purpose: Generic
Provisioning: Thin
Compression: lz4
Deduplication: Off
Snapshots: None

Project B

Purpose: Generic
Provisioning: Thin
Compression: lz4
Deduplication: Off
Snapshots: None

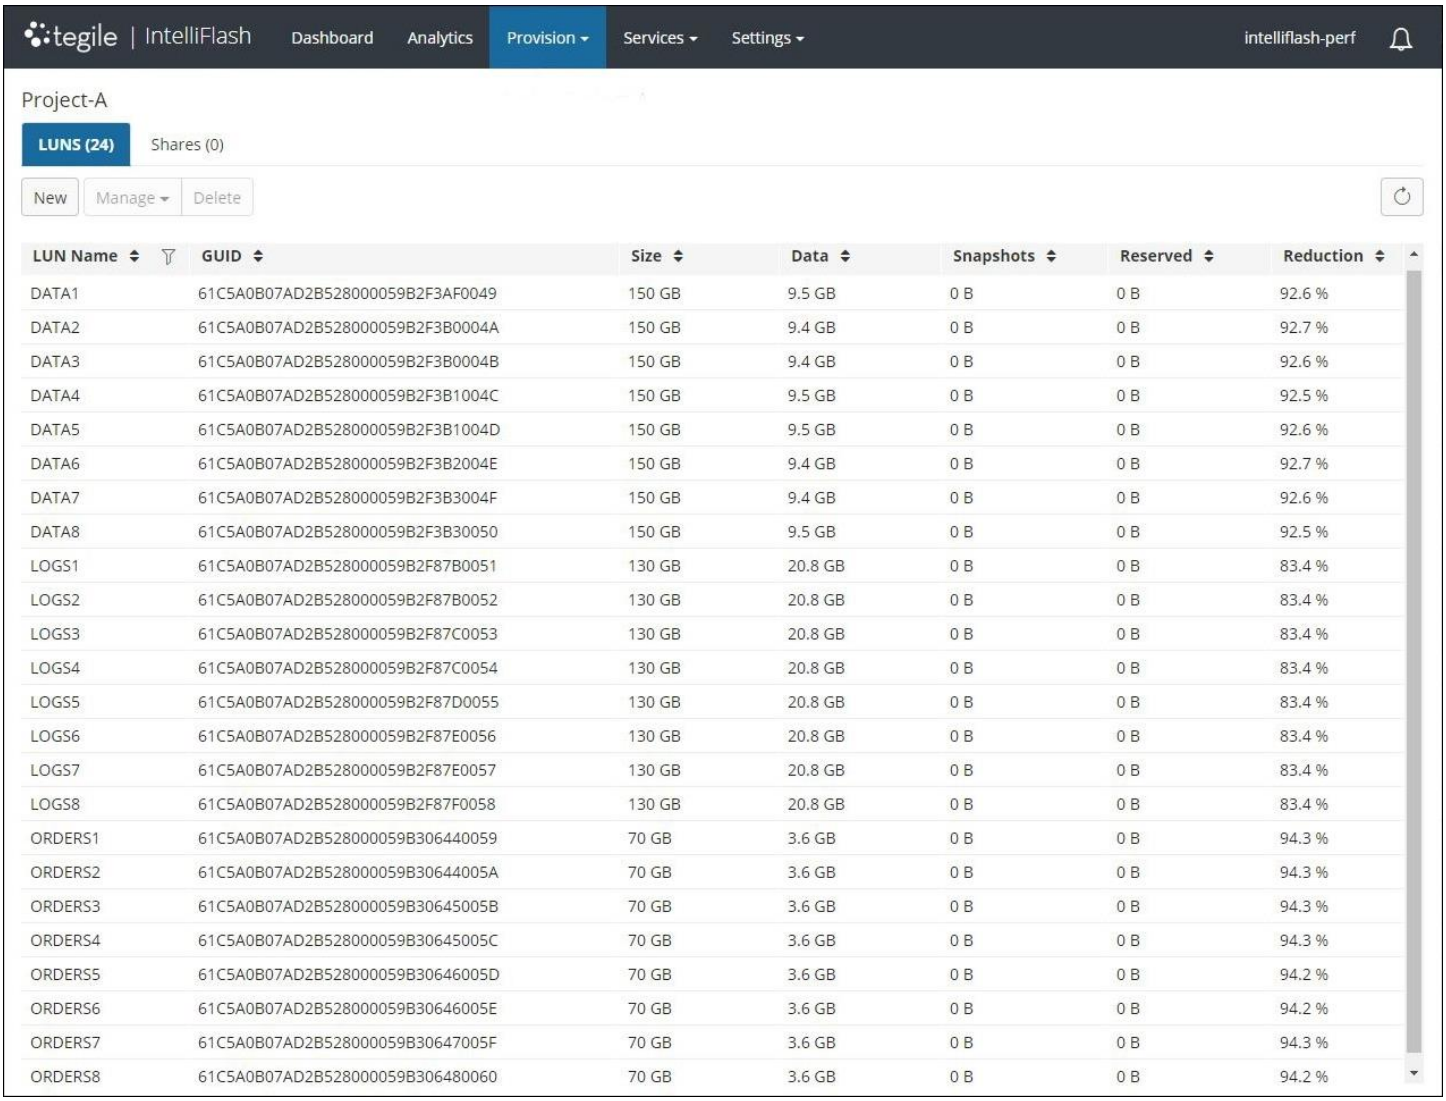

LUN Name	GUID	Size	Data	Snapshots	Reserved	Reduction
DATA1	61C5A0B07AD2B528000059B2F3AF0049	150 GB	9.5 GB	0 B	0 B	92.6 %
DATA2	61C5A0B07AD2B528000059B2F3B0004A	150 GB	9.4 GB	0 B	0 B	92.7 %
DATA3	61C5A0B07AD2B528000059B2F3B0004B	150 GB	9.4 GB	0 B	0 B	92.6 %
DATA4	61C5A0B07AD2B528000059B2F3B1004C	150 GB	9.5 GB	0 B	0 B	92.5 %
DATA5	61C5A0B07AD2B528000059B2F3B1004D	150 GB	9.5 GB	0 B	0 B	92.6 %
DATA6	61C5A0B07AD2B528000059B2F3B2004E	150 GB	9.4 GB	0 B	0 B	92.7 %
DATA7	61C5A0B07AD2B528000059B2F3B3004F	150 GB	9.4 GB	0 B	0 B	92.6 %
DATA8	61C5A0B07AD2B528000059B2F3B30050	150 GB	9.5 GB	0 B	0 B	92.5 %
LOGS1	61C5A0B07AD2B528000059B2F87B0051	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS2	61C5A0B07AD2B528000059B2F87B0052	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS3	61C5A0B07AD2B528000059B2F87C0053	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS4	61C5A0B07AD2B528000059B2F87C0054	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS5	61C5A0B07AD2B528000059B2F87D0055	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS6	61C5A0B07AD2B528000059B2F87E0056	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS7	61C5A0B07AD2B528000059B2F87E0057	130 GB	20.8 GB	0 B	0 B	83.4 %
LOGS8	61C5A0B07AD2B528000059B2F87F0058	130 GB	20.8 GB	0 B	0 B	83.4 %
ORDERS1	61C5A0B07AD2B528000059B306440059	70 GB	3.6 GB	0 B	0 B	94.3 %
ORDERS2	61C5A0B07AD2B528000059B30644005A	70 GB	3.6 GB	0 B	0 B	94.3 %
ORDERS3	61C5A0B07AD2B528000059B30645005B	70 GB	3.6 GB	0 B	0 B	94.3 %
ORDERS4	61C5A0B07AD2B528000059B30645005C	70 GB	3.6 GB	0 B	0 B	94.3 %
ORDERS5	61C5A0B07AD2B528000059B30646005D	70 GB	3.6 GB	0 B	0 B	94.2 %
ORDERS6	61C5A0B07AD2B528000059B30646005E	70 GB	3.6 GB	0 B	0 B	94.2 %
ORDERS7	61C5A0B07AD2B528000059B30647005F	70 GB	3.6 GB	0 B	0 B	94.3 %
ORDERS8	61C5A0B07AD2B528000059B306480060	70 GB	3.6 GB	0 B	0 B	94.2 %

Figure 9 – LUN Settings

Node A LUNs
Access: 16Gb Fibre Channel
Purpose: Database
Block Size: 16 KB
8 x 150 GB LUNs for DATA tablespace
8 x 130 GB LUNs for redo logs
8 x 70 GB LUNs for ORDERS tablespace

Node B LUNs
Access: 16Gb Fibre Channel
Purpose: Database
Block Size: 16 KB
8 x 150 GB LUNs for DATA tablespace
8 x 130 GB LUNs for redo logs
8 x 70 GB LUNs for ORDERS tblespace

Oracle ASM Configuration

For optimal performance, a database on each node was created using Oracle ASM with the following configuration. Each disk group is comprised of eight LUNs for optimal performance.

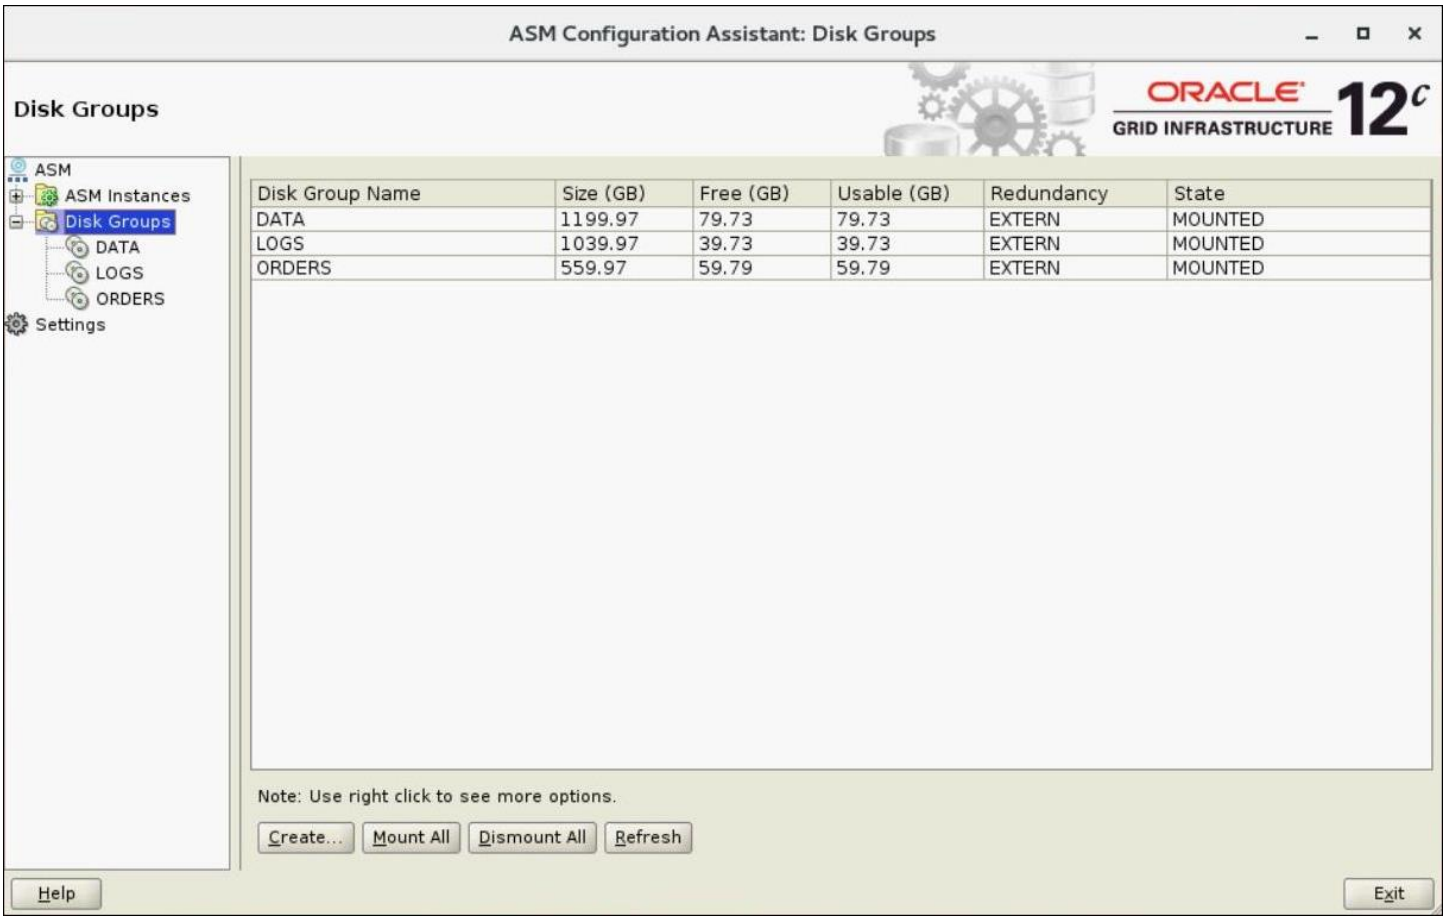

Figure 10 - ASM Configuration Assistant

Oracle Listener Configuration

The database was configured with four listener processes, each listening on a dedicated 10 Gb network IP connection with a separate port number for each listener. The listener.ora configuration follows:

```
$ cat listener.ora
LISTENER_TPCC_4 = (DESCRIPTION_LIST = (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1525))))
LISTENER_TPCC_3 = (DESCRIPTION_LIST = (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1524))))
LISTENER_TPCC_2 = (DESCRIPTION_LIST = (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1523))))
LISTENER_TPCC_1 = (DESCRIPTION_LIST = (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1522))
    (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1522))))
SID_LIST_LISTENER_TPCC_4 = (SID_LIST = (SID_DESC =
    (ORACLE_HOME = /oracle/12c/database) (SID_NAME = HDB)))
SID_LIST_LISTENER_TPCC_3 = (SID_LIST = (SID_DESC =
    (ORACLE_HOME = /oracle/12c/database) (SID_NAME = HDB)))
SID_LIST_LISTENER_TPCC_2 = (SID_LIST = (SID_DESC =
    (ORACLE_HOME = /oracle/12c/database) (SID_NAME = HDB)))
SID_LIST_LISTENER_TPCC_1 = (SID_LIST = (SID_DESC =
    (ORACLE_HOME = /oracle/12c/database) (SID_NAME = HDB)) )

ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_TPCC_1=ON
VALID_NODE_CHECKING_REGISTRATION_LISTENER_TPCC_1=SUBNET
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_TPCC_2=ON
VALID_NODE_CHECKING_REGISTRATION_LISTENER_TPCC_2=SUBNET
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_TPCC_3=ON
VALID_NODE_CHECKING_REGISTRATION_LISTENER_TPCC_3=SUBNET
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_TPCC_4=ON
VALID_NODE_CHECKING_REGISTRATION_LISTENER_TPCC_4=SUBNET
```

To equally distribute the workload across the four listener processes, load balancing was enabled using the tnsnames.ora configuration as follows:

```
$ cat tnsnames.ora
HDB =  (DESCRIPTION =
            (LOAD_BALANCE=on)
            (FAILOVER=on)
            (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1522))
                (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1523))
                (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1524))
                (ADDRESS = (PROTOCOL = TCP) (HOST = db_10) (PORT = 1525))
            )
        (CONNECT_DATA = (SERVICE_NAME = HDB))
        )
```

Oracle Database Creation

The Oracle Database was created with the following tablespaces:

```
$ cat createHDB.sql

CREATE DATABASE HDB
  CONTROLFILE REUSE
  USER sys IDENTIFIED BY tegile
  USER system IDENTIFIED BY tegile
  SET DEFAULT BIGFILE TABLESPACE
  LOGFILE GROUP 1 '+LOGS/log1' SIZE 500G REUSE,
            GROUP 2 '+LOGS/log2' SIZE 500G REUSE
  MAXDATAFILES 200
  DATAFILE '+DATA/sys' SIZE 10G REUSE
  SYSAUX DATAFILE '+DATA/sysaux' SIZE 10G REUSE
  DEFAULT TABLESPACE data DATAFILE '+DATA/data' SIZE 500G REUSE
  DEFAULT TEMPORARY TABLESPACE temp TEMPFILE '+DATA/temp' SIZE 100G REUSE
  UNDO TABLESPACE undo DATAFILE '+DATA/undo' SIZE 500G REUSE;

CREATE TABLESPACE orders DATAFILE '+ORDERS/orders' SIZE 500G REUSE;

@/u01/app/oracle/product/12.1.0.2/db_1/rdbms/admin/catalog.sql
@/u01/app/oracle/product/12.1.0.2/db_1/rdbms/admin/catexp.sql
@/u01/app/oracle/product/12.1.0.2/db_1/rdbms/admin/catldr.sql
@/u01/app/oracle/product/12.1.0.2/db_1/rdbms/admin/catproc.sql
@/u01/app/oracle/product/12.1.0.2/db_1/rdbms/admin/catparr.sql
```

Oracle Database Tuning

For optimal performance, the following database tunings were applied:

```
$ cat initHDB.ora
_ash_size=2013265920
_disk_sector_size_override=TRUE
_enable_reliable_latch_waits=TRUE
_fast_cursor_reexecute=TRUE
_fg_sync_sleep_usecs=500
_undo_autotune=FALSE
aq_tm_processes=0
audit_file_dest='/home/oracle/database/audit'
audit_trail='NONE'
commit_logging='BATCH'
commit_wait='NOWAIT'
compatible='12.2.0.1.0'
control_files='+DATA/control'
db_block_checking=FALSE
db_block_checksum=FALSE
db_block_size=16384
db_cache_advice=off
db_cache_size=100G
db_create_file_dest='+DATA'
db_domain=''
db_name='HDB'
db_recovery_file_dest='+DATA'
db_recovery_file_dest_size=4300000000000
db_unrecoverable_scn_tracking=FALSE
db_writer_processes=16
ddl_lock_timeout=30
deferred_segment_creation=FALSE
diagnostic_dest='/home/oracle/database'
disk_asynch_io=TRUE
dispatchers='(PROTOCOL=TCP) (SERVICE=HDB)'
dml_locks=1000
```

filesystemio_options=setall

inmemory_size=0

lock_sga=TRUE

log_checkpoint_interval=0

log_checkpoint_timeout=0

log_checkpoints_to_alert=TRUE

open_cursors=2400

optimizer_capture_sql_plan_baselines=FALSE

optimizer_dynamic_sampling=4

parallel_degree_policy='AUTO'

processes=8000

query_rewrite_enabled=FALSE

remote_login_passwordfile='EXCLUSIVE'

replication_dependency_tracking=FALSE

sga_max_size=250G

shared_pool_size=32G

statistics_level=BASIC

timed_statistics=FALSE

trace_enabled=FALSE

transactions=12000

transactions_per_rollback_segment=1

undo_retention=30

undo_tablespace='undo'

HammerDB Settings

HammerDB settings were configured as follows:

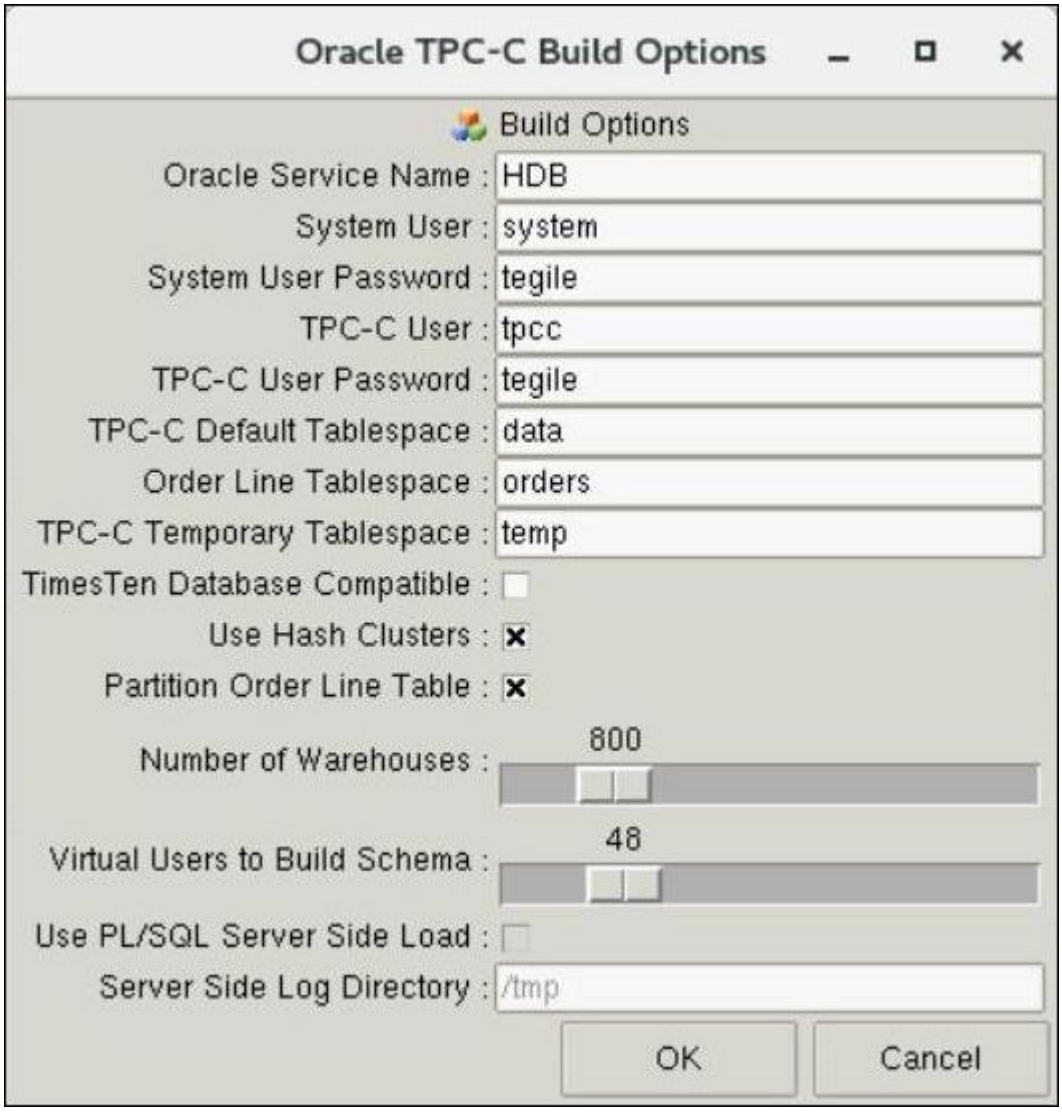

Figure 11 - HammerDB Oracle TPC-C Build Options

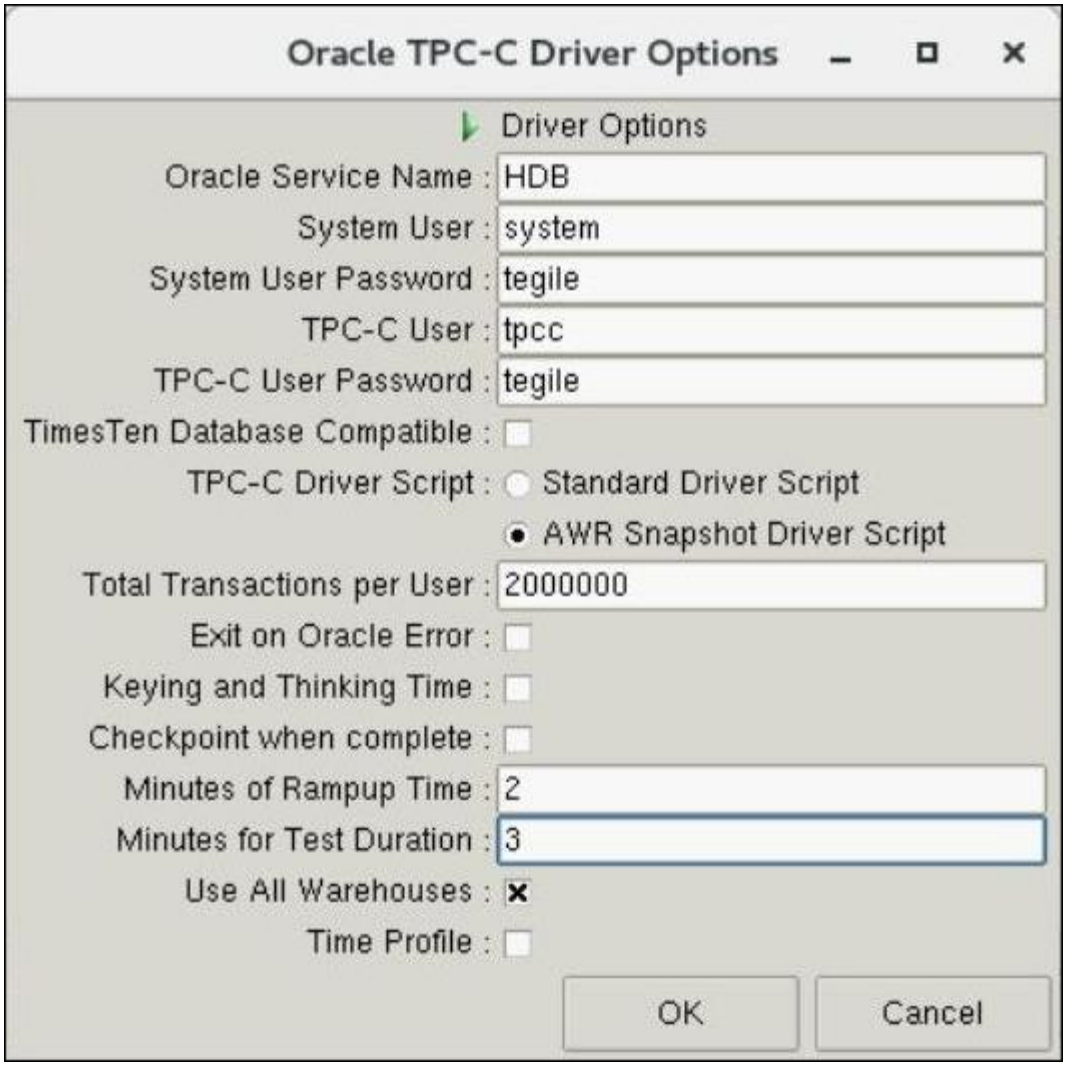

Figure 12 - HammerDB Oracle TPC-C Driver Options

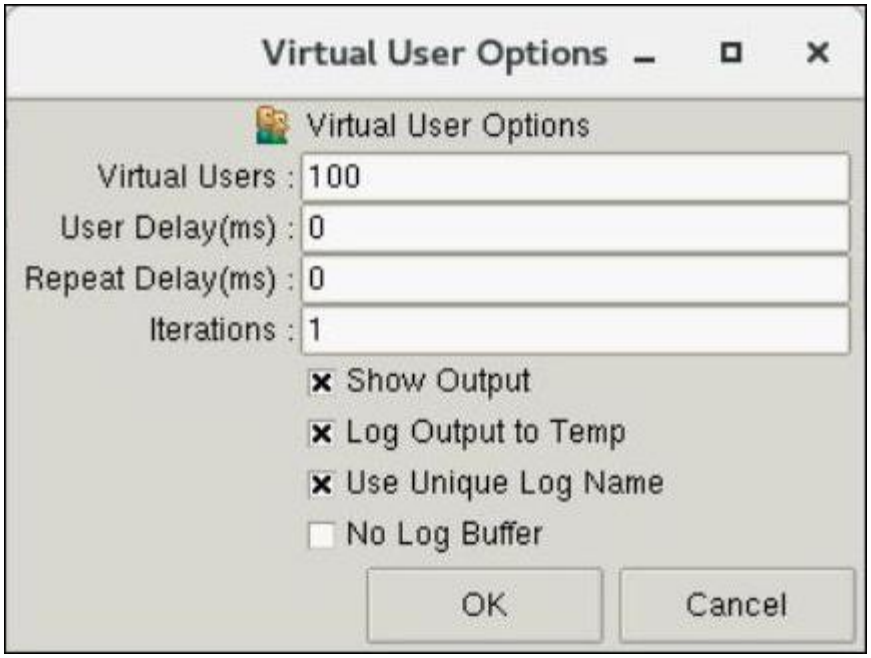

Figure 13 - HammerDB Virtual User Options

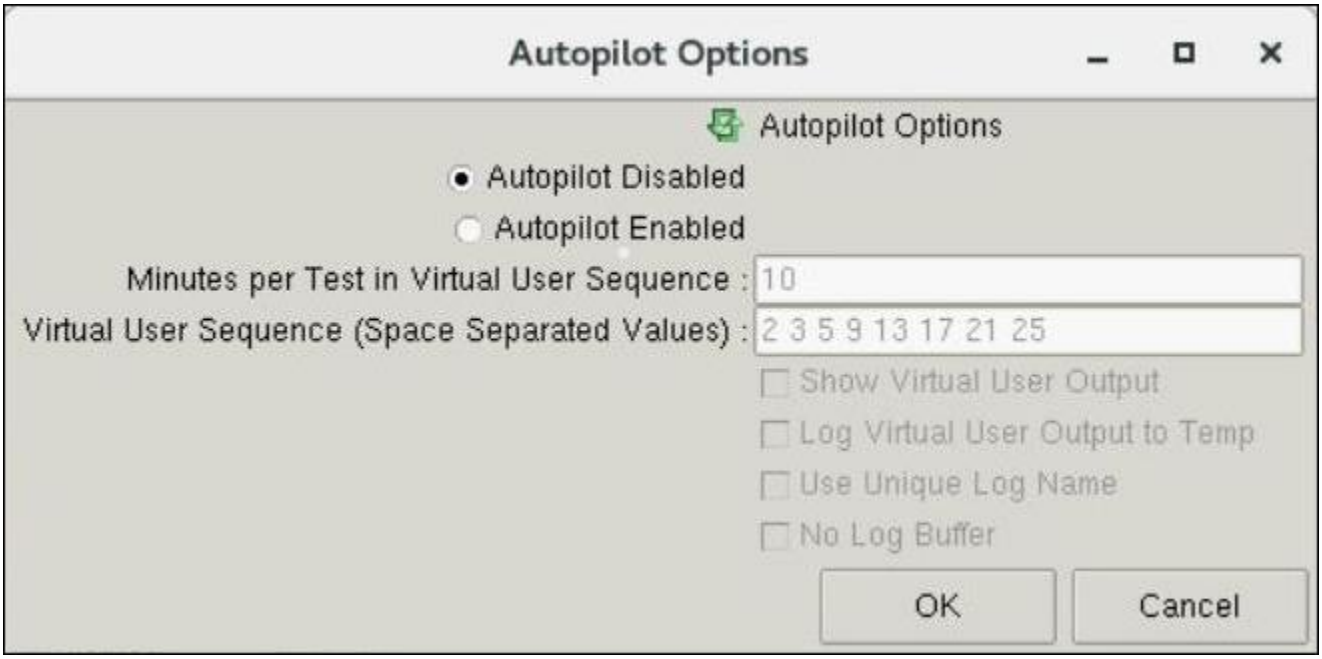

Figure 14 - HammerDB Autopilot Options

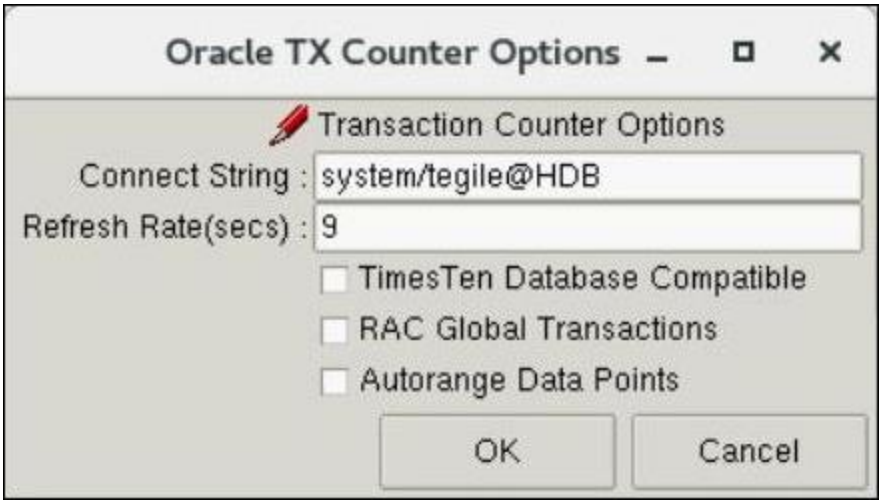

Figure 15 - HammerDB Oracle Transaction Counter Options

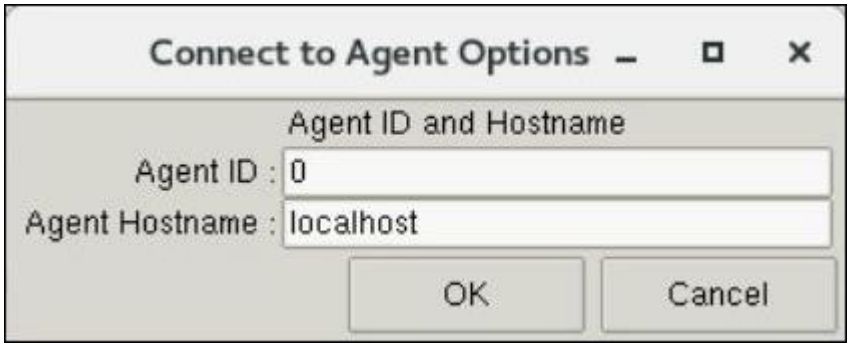

Figure 16 - HammerDB Connect to Agent Options

Results Summary

An aggregate HammerDB score of **8,392,994 Oracle TPM** and **2,877,715 NOPM** was achieved using two Oracle instances connected to separate controllers on the Tegile storage array.

Node A score of 4,189,540 Oracle TPM and 1,435,474 NOPM.
Node B score of 4,203,454 Oracle TPM and 1,442,241 NOPM.

These are not official TPC-C results.

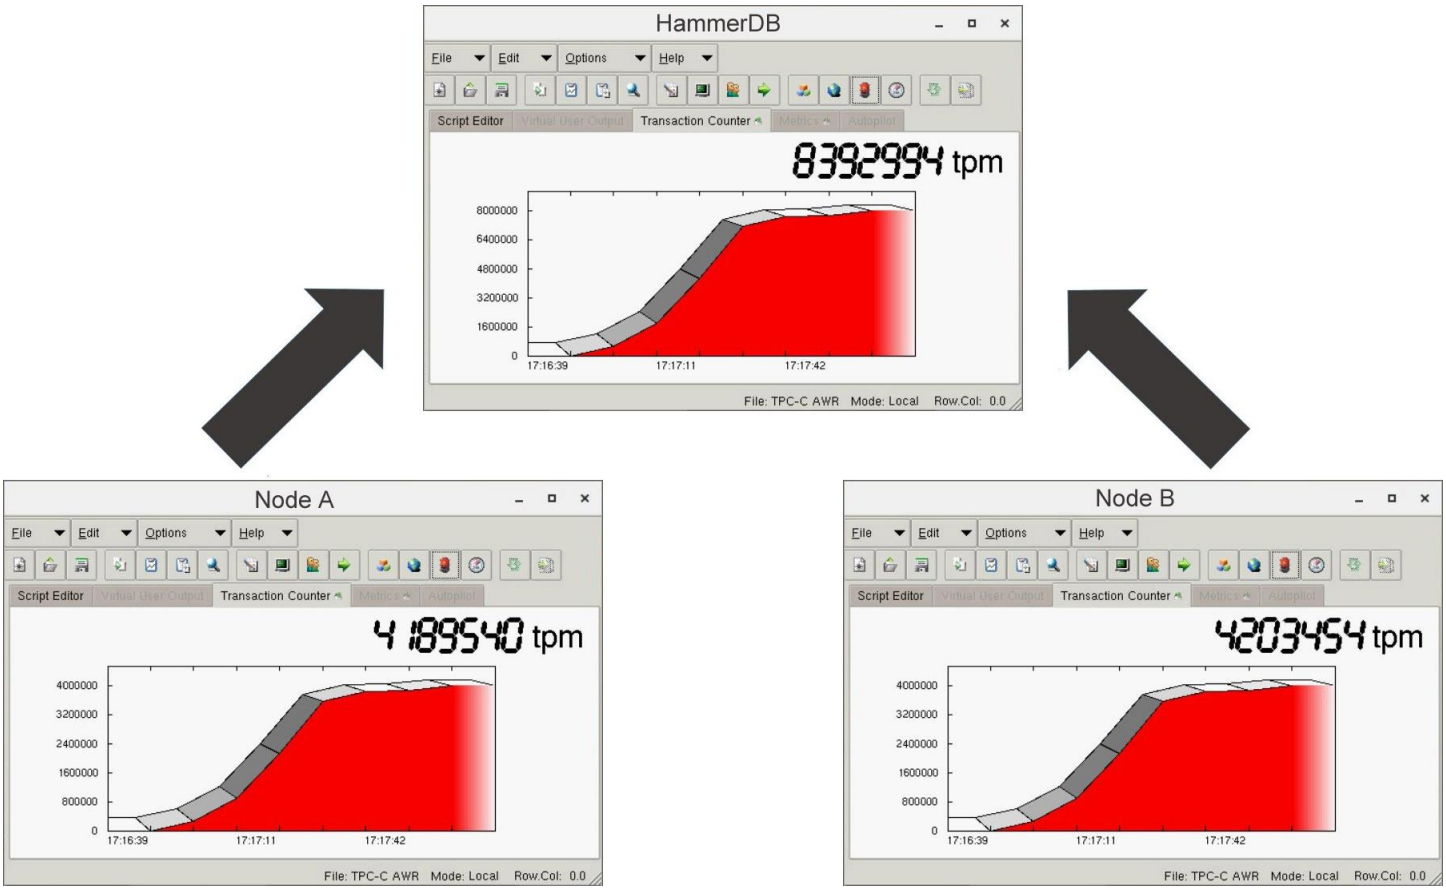

Figure 17 – Aggregate HammerDB Result

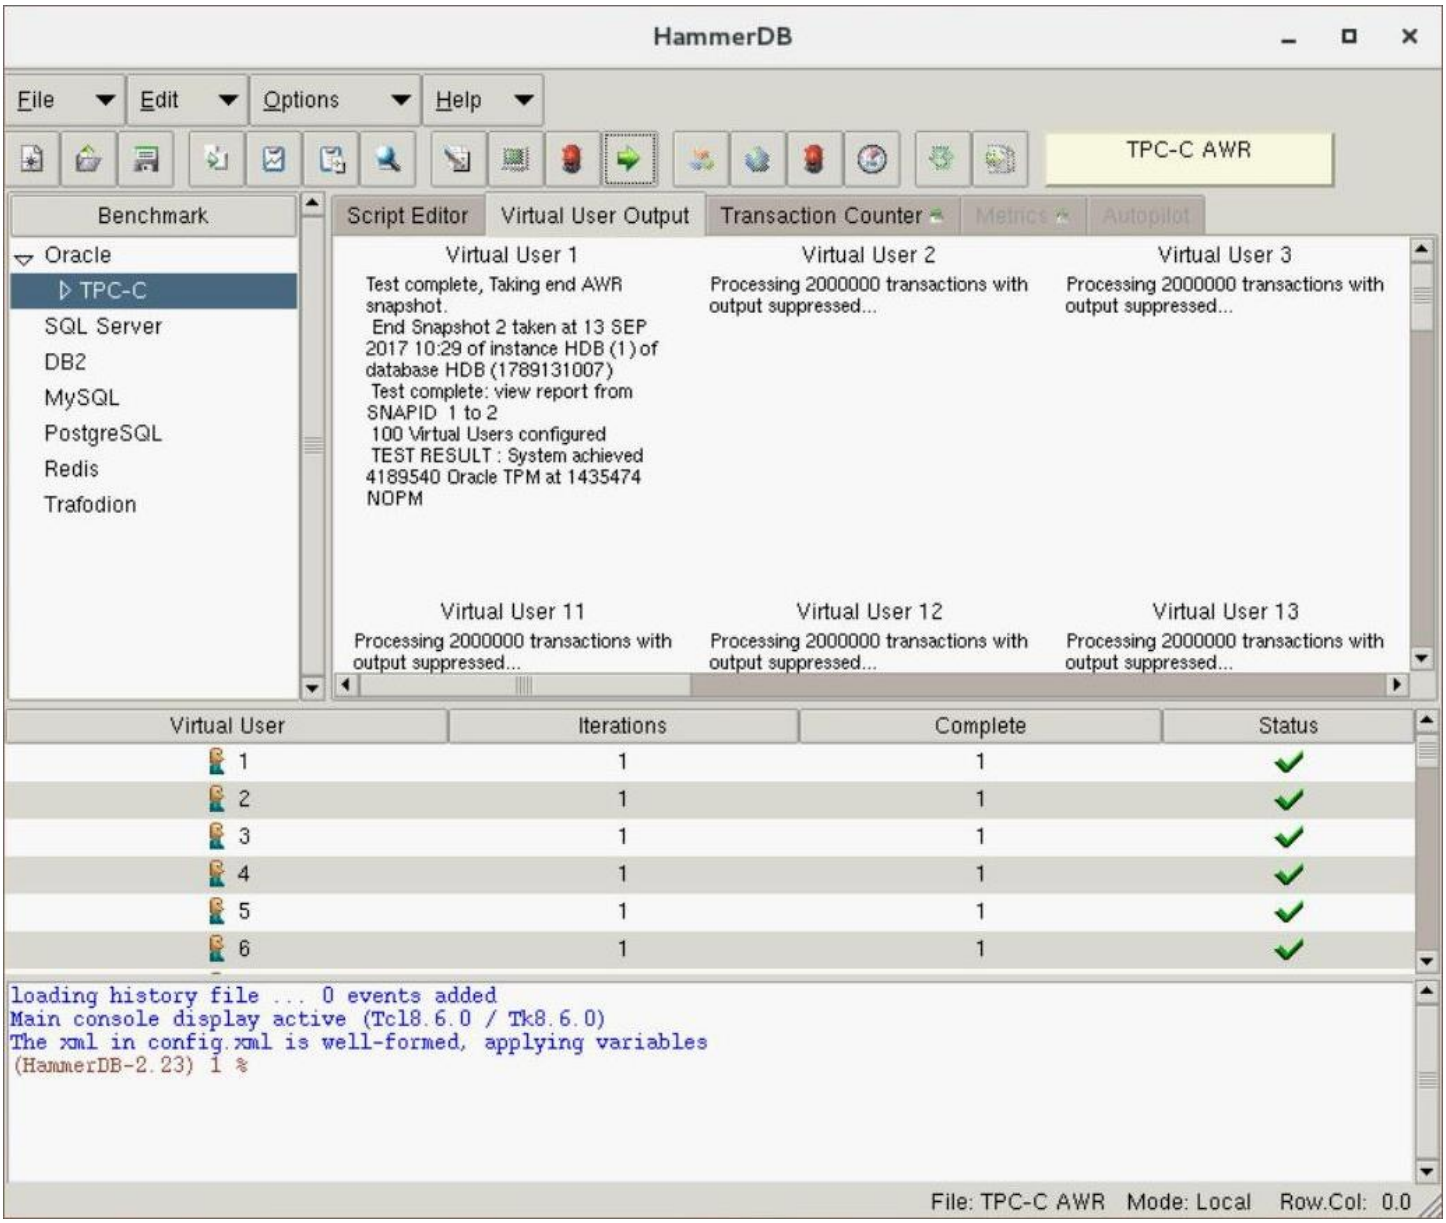

Figure 18 – Results Summary for Node A

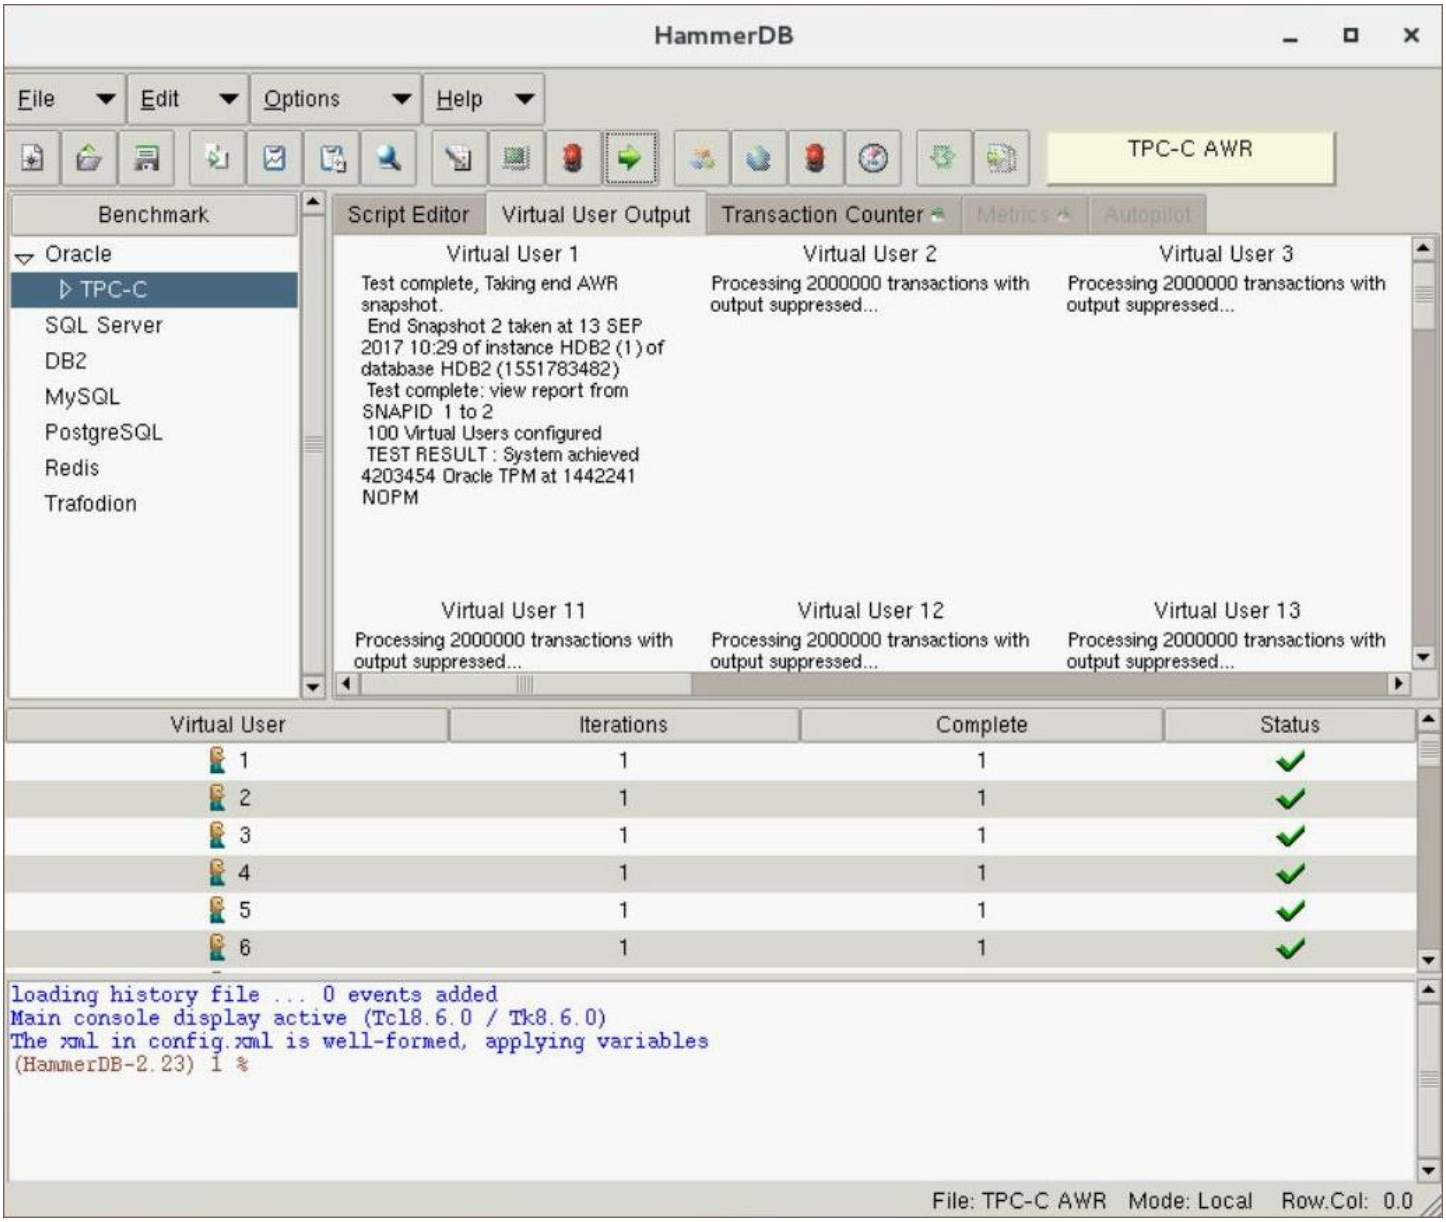

Figure 19 – Results Summary for Node B

Please note that the HammerDB tool has a GUI transaction counter that monitors progress during execution of the test. This counter generates fluctuating numbers over the course of the measurement window and therefore should not be used for reporting the final score.

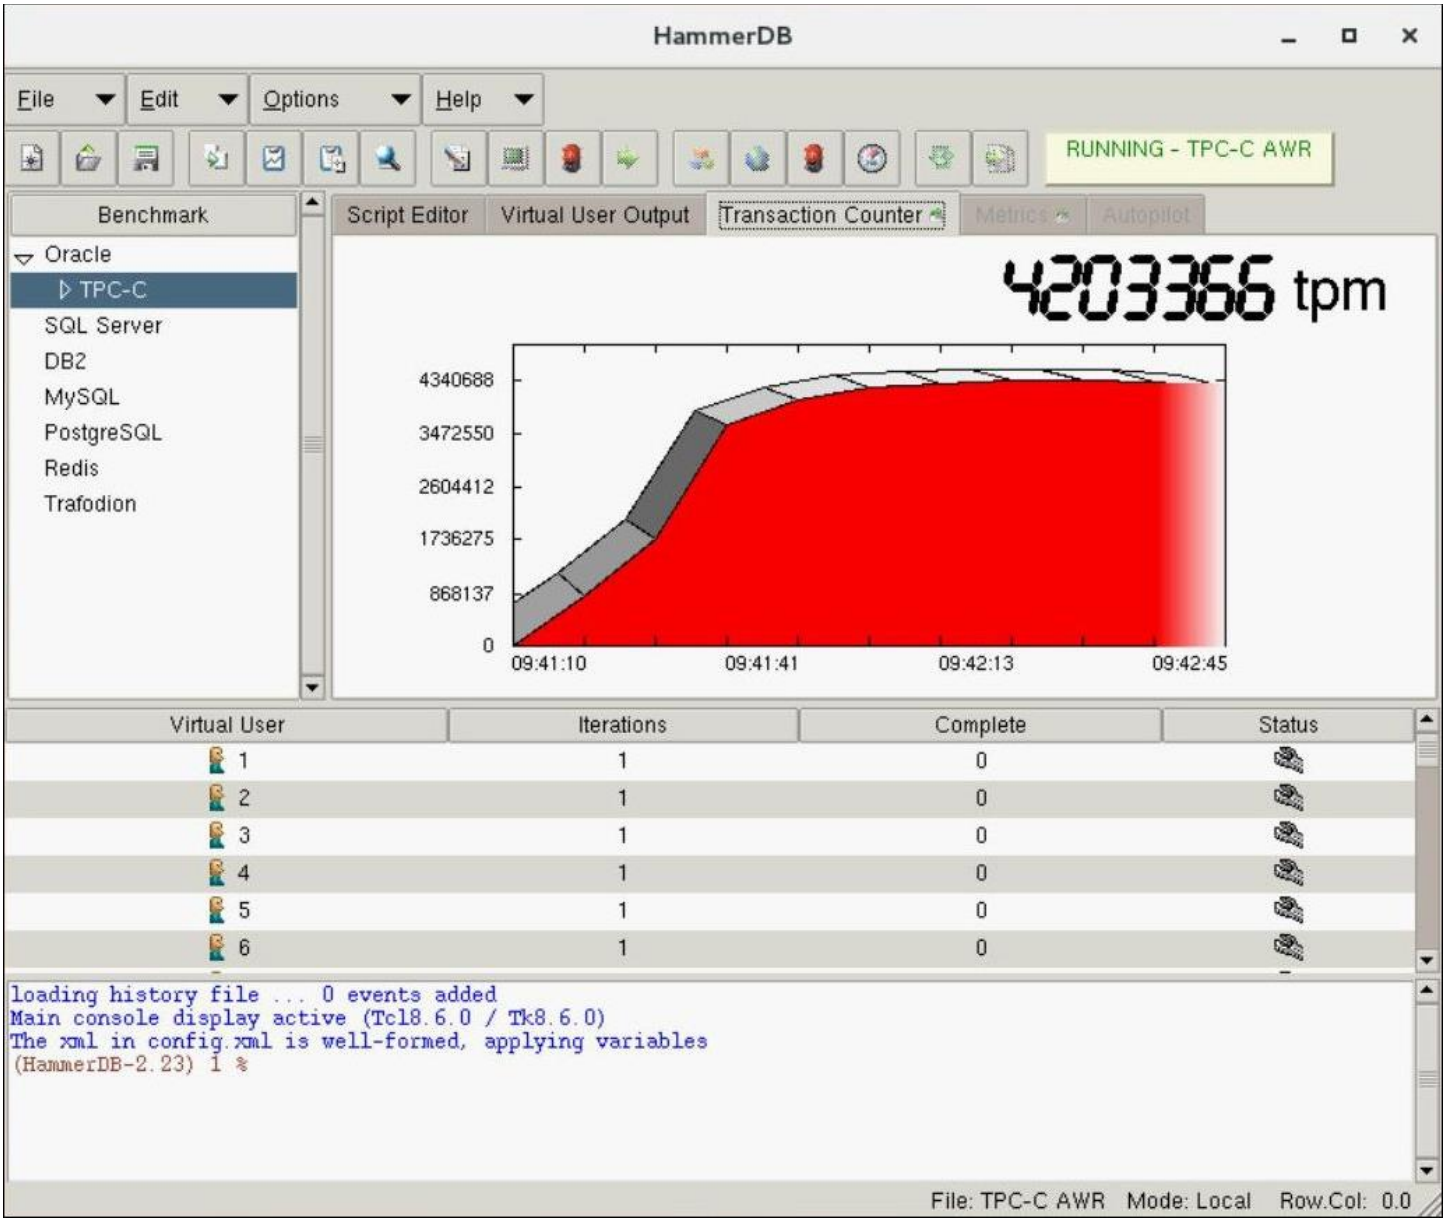

Figure 20 - HammerDB Transaction Counter

The TPM number was calculated using the Transactions metric from the Oracle Automatic Workload Repository (AWR) snapshots collected at the beginning and end of the measurement window. The Transactions metric (number of transactions per second) recorded in the Workload Repository Report Summary is multiplied by 60 to arrive at Transactions Per Minute (TPM).

Report Summary

Load Profile

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	0.0	0.0	0.00	0.00
DB CPU(s):	0.0	0.0	0.00	0.00
Background CPU(s):	0.0	0.0	0.00	0.00
Redo size (bytes):	369,349,705.0	5,289.6		
Logical read (blocks):	3,976,037.9	56.9		
Block changes:	2,177,238.4	31.2		
Physical read (blocks):	1,719.4	0.0		
Physical write (blocks):	771.2	0.0		
Read IO requests:	988.4	0.0		
Write IO requests:	115.0	0.0		
Read IO (MB):	26.9	0.0		
Write IO (MB):	12.1	0.0		
IM scan rows:	0.0	0.0		
Session Logical Read IM:	0.0	0.0		
User calls:	53,512.9	0.8		
Parses (SQL):	32,307.9	0.5		
Hard parses (SQL):	0.8	0.0		
SQL Work Area (MB):	0.0	0.0		
Logons:	0.1	0.0		
Executes (SQL):	1,440,724.5	20.6		
Rollbacks:	117.0	0.0		
Transactions:	69,825.7			

Figure 21 – Workload Repository Report Summary for Node A

Report Summary

Load Profile

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	0.0	0.0	0.00	0.00
DB CPU(s):	0.0	0.0	0.00	0.00
Background CPU(s):	0.0	0.0	0.00	0.00
Redo size (bytes):	370,626,636.8	5,290.3		
Logical read (blocks):	3,991,868.8	57.0		
Block changes:	2,185,417.1	31.2		
Physical read (blocks):	1,750.7	0.0		
Physical write (blocks):	8.4	0.0		
Read IO requests:	1,014.0	0.0		
Write IO requests:	1.4	0.0		
Read IO (MB):	27.4	0.0		
Write IO (MB):	0.1	0.0		
IM scan rows:	0.0	0.0		
Session Logical Read IM:	0.0	0.0		
User calls:	53,725.4	0.8		
Parses (SQL):	32,426.5	0.5		
Hard parses (SQL):	0.7	0.0		
SQL Work Area (MB):	0.0	0.0		
Logons:	0.1	0.0		
Executes (SQL):	1,445,467.2	20.6		
Rollbacks:	118.0	0.0		
Transactions:	70,057.6			

Figure 22 – Workload Repository Report Summary for Node B

Storage Array Performance

The Tegile T4700 storage array delivered excellent performance for the HammerDB workload.

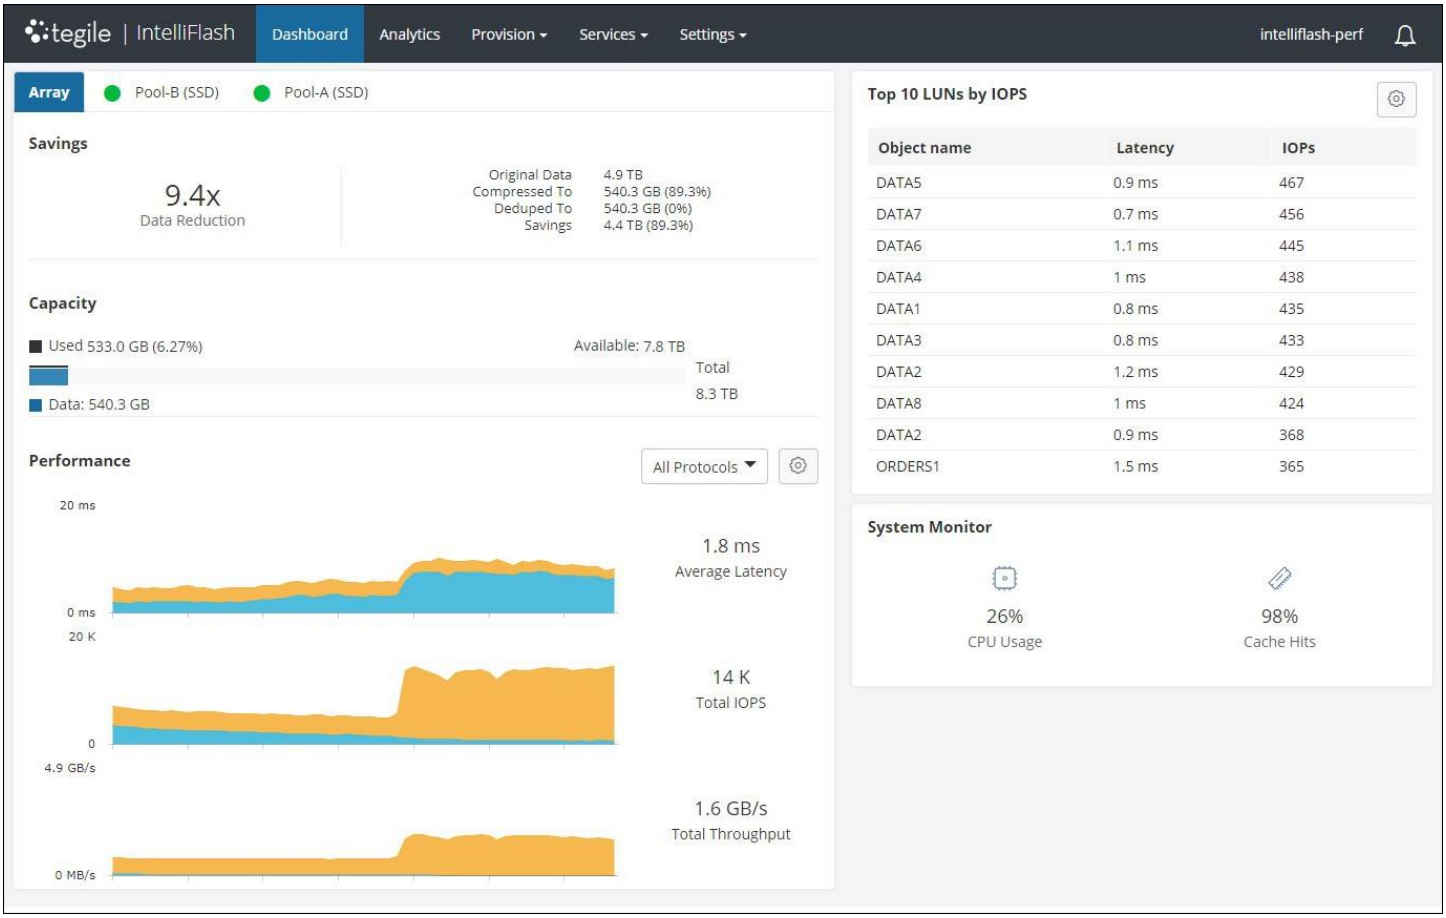

Figure 23 – Tegile T4700 Storage Array Performance

Browse the publications at hammerdb.com/benchmarks.html to learn about the performance levels achieved using HammerDB.

Please note that publications on public websites may be removed or changed at any time.

About Tegile Systems, Inc.

Tegile Systems is pioneering a new generation of flash-driven enterprise storage arrays that balance performance, capacity, features and price for virtualization, file services and database applications. With Tegile's line of all-flash and hybrid storage arrays, the company is redefining the traditional approach to storage by providing a family of arrays that accelerate business critical enterprise applications and allow customers to significantly consolidate mixed workloads in virtualized environments.

Tegile's patented IntelliFlash™ technology accelerates performance and enables inline deduplication and compression of data so each array has a usable capacity far greater than its raw capacity. Tegile's award-winning solutions enable customers to better address the requirements of server virtualization, virtual desktop integration and database integration than any other offerings. Featuring both NAS and SAN connectivity, Tegile arrays are easy-to-use, fully redundant, and highly scalable. They come complete with built-in snapshot, remote-replication, near-instant recovery, onsite or offsite failover, and VM-aware features. Follow us on Twitter @tegile or visit us at www.tegile.com