BOOSTING PERFORMANCE WITH THE DELL ACCELERATION APPLIANCE FOR DATABASES

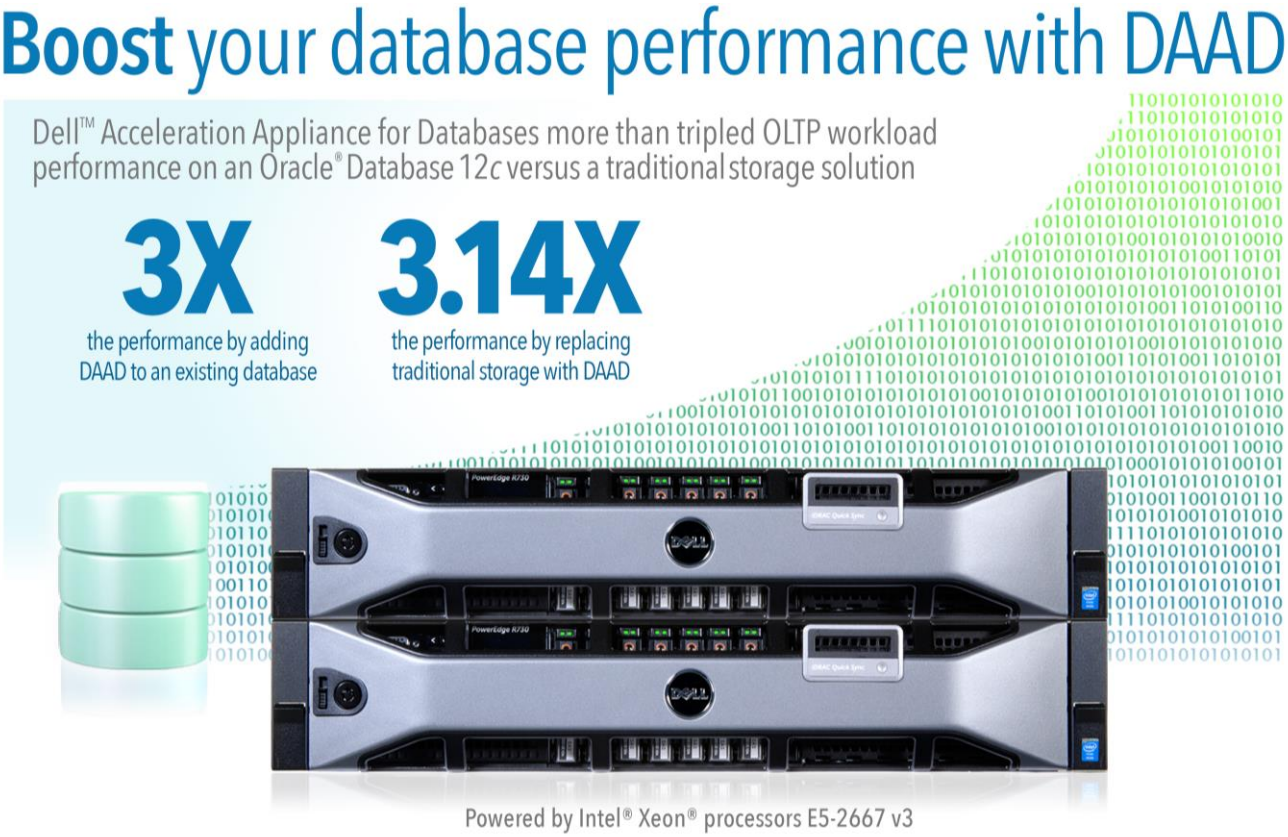

Anything you can do to help your customers place an order or access the information they're looking for more quickly improves your business. Just because you have powerful database servers doesn't mean you should overlook other ways to support expanding user counts as your business grows.

The new highly available Dell Acceleration Appliance for Databases (DAAD) 2.0 is a standalone appliance built on Intel Xeon processors E5-2667 v3 that is easily integrated into existing infrastructure. By eliminating OLTP database performance bottlenecks with this storage appliance, you can avoid the costly process of replacing sprawling legacy storage or revamping your entire infrastructure. A highly available DAAD 2.0 solution pairs two servers for redundancy to reduce the possibility of costly downtime.

In the Principled Technologies datacenter, we set up an Oracle Database 12*c* environment on a Dell PowerEdge R720 server paired with a traditional external storage array. We found that adding a highly available DAAD 2.0 solution yielded 3.01 times the database performance. Replacing traditional storage completely with the DAAD solution yielded 3.14 times the performance, so it can also help you reclaim precious datacenter space that traditional storage takes up. If you need your servers to support an expanding user count, our tests show that the Dell Acceleration Appliance for Databases solution can give your important databases a significant boost.

MAY 2015
A PRINCIPLED TECHNOLOGIES TEST REPORT
Commissioned by Dell Inc.

HOW IT WORKS

The Dell Acceleration Appliance for Databases 2.0 is a pre-integrated server/software combination consisting of one or two Dell PowerEdge R730 servers with Intel Xeon processors E5-2667 v3 , DAAD ION ION Accelerator® software, and Fusion SX300 ioMemory® drives, which are high-capacity NAND Flash PCIe® SSDs. Available in 12.8TB and 25.6TB capacities, DAAD 2.0 supports Fibre Channel, iSCSI, and Infiniband/SRP front-end fabrics. Designed to work with many database platforms and front-end servers, DAAD makes appliance resources available to the application server and uses I/O acceleration to optimize database requests, leveraging DRAM to improve performance. Boosting database performance with DAAD can help you work within your existing infrastructure, so you don't have to purchase more disk or compute resources, incur additional database licensing costs, or take up vast amounts of extra datacenter space. To learn more about the Dell Acceleration Appliance for Databases, visit en.community.dell.com/techcenter/enterprise-solutions/m/oracle_db_gallery/20441362/download.

In this study we configured two 25.6TB DAAD appliance nodes as an HA cluster pair. For information about the DAAD and our other test components, see Appendix A. For detailed system configuration, see Appendix B. For step-by-step testing details, see Appendix C.

WHAT WE FOUND

About the results

In our tests, we found that the highly available Dell Acceleration Appliances for Databases solution worked as advertised: Adding the DAAD to our traditional storage environment increased our Oracle Database 12*c* performance by 3.01 times, and replacing the traditional storage solution with the DAAD increased performance by 3.14 times (see Figure 1).

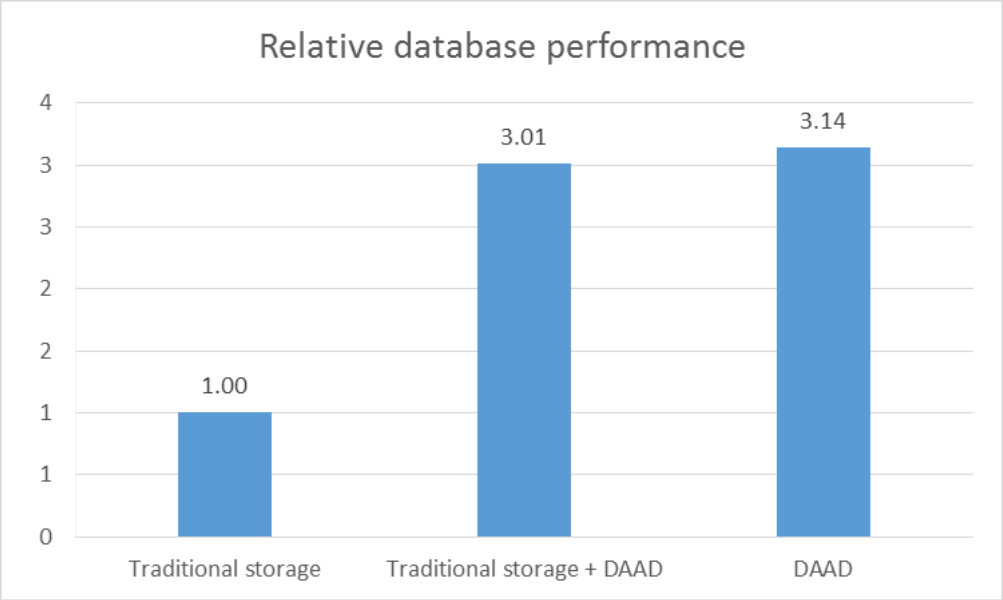

Figure 1: The highly available Dell Acceleration Appliance for Databases solution delivered up to 3.14 times the database performance of an environment using traditional storage.

While the Dell PowerEdge R720 with Intel Xeon processors E5-2650 v2 performed admirably with a traditional storage appliance, adding a highly available DAAD to the existing environment provided a significant performance boost, and replacing the traditional storage with the DAAD also yielded big performance advantages. By using a high availability DAAD to improve database performance, you can avoid having to make major changes to your database infrastructure as your user demands increase with your growing business.

The ability of the DAAD solution to maintain high availability through redundant pairs works to ensure that your databases keep going, even if a piece of hardware fails. This is critical to a successful database infrastructure, because any downtime where your hardware isn't working can be very costly as business stops.

CONCLUSION

If your business is expanding and you need to support more users accessing your databases, it's time to act. Upgrading your database infrastructure with a flash storage-based solution is a smart way to improve performance without adding more servers or taking up very much rack space, which comes at a premium. The Dell Acceleration Appliance for Databases addresses this by providing strong performance when combined with your existing infrastructure or on its own.

We found that adding a highly available DAAD solution to our database application provided up to 3.01 times the Oracle Database 12*c* performance, which can make a big difference to your bottom line. Additionally, the DAAD delivered 3.14 times the database performance when replacing traditional storage completely, which could enable your infrastructure to keep up with your growing business' needs.

APPENDIX A – ABOUT THE COMPONENTS

About the Dell Acceleration Appliance for Databases 2.0 hardware

The 2U Dell PowerEdge R730 rack servers used in the highly available DAAD configuration are each powered by two Intel Xeon processors E5-2667 v3, 384 GB of DDR4 RAM, and QLogic QLE2662 16Gb Fibre Channel HBAs, providing functional flexibility in the datacenter. The Dell PowerEdge R730 rack servers are each powered by four 6400 GB SanDisk SX300 ioMemory PCIe SSDs to reduce storage bottlenecks.

With redundant power supply units, hot-swappable hardware, and Dual SDTM card for Failsafe Hypervisors, the Dell PowerEdge R730 supports hardware high availability. The PowerEdge R730 comes standard with iDRAC8 with Lifecycle Controller and Dell OpenManage, which all work to streamline management. For more details on the Dell PowerEdge R730, visit http://www.dell.com/us/business/p/poweredge-r730/pd. For more details on the Intel Xeon processor E5-2600 v3 series, visit www.intel.com/content/dam/www/public/us/en/documents/product-briefs/xeon-e5-brief.pdf.

About the benchmark we used - HammerDB

HammerDB is an open-source benchmark tool that tests the database performance of many databases, including Oracle Database, Microsoft® SQL Server®, PostgreSQL®, MySQL™, and more. The benchmark includes built-in workloads derived from industry-standard benchmarks, such as a transactional (TPC-C-derived) workload and a data warehouse (TPC-H-derived) workload. For this study, we used the transactional workload. The TPC-C HammerDB workload is derived from TPC-C and as such is not comparable to published TPC-C results. For more information, visit www.hammerora.sourceforge.net.

APPENDIX B – SYSTEM CONFIGURATION INFORMATION

Figure 2 provides detailed configuration information for the test systems.

System	Dell PowerEdge R720	Dell Acceleration Appliance for Databases 2.0
Power supplies		
Total number	1	2 (HA pair)
Vendor and model number	Dell 05NF18X02	Dell 0G6W6KX02
Wattage of each (W)	750	750
General		
Number of processor packages	2	2
Number of cores per processor	8	8
Number of hardware threads per core	2	2
System power management policy	Performance	Performance
CPU		
Vendor	Intel	Intel
Name	Xeon	Xeon
Model number	E5-2650 v2	E5-2667 v3
Socket type	LGA 2011	FCLGA2011-3
Core frequency (GHz)	2.6	3.2
Bus frequency	7.2 GT/s	9.6 GT/s
L1 cache	32 + 32 KB (per core)	32 + 32 KB (per core)
L2 cache	256 KB (per core)	256 KB (per core)
L3 cache	20 MB	20 MB
Platform		
Vendor and model number	Dell PowerEdge R720	Dell PowerEdge R730
Motherboard model number	0M1GCR	0599V5
BIOS name and version	2.4.3	1.1.4
BIOS settings	Defaults	Defaults
Memory module(s)		
Total RAM in system (GB)	128	384
Vendor and model number	Hynix HMT31GR7BFR4A-H9	Hynix HMA42GR7MFR4N-TFT1
Type	PC3-10600	PC4-17000
Speed (MHz)	1,333	2,133
Speed running in the system (MHz)	1,333	2,133
Size (GB)	8	16
Number of RAM module(s)	16	24
Chip organization	Double-sided	Double-sided
Rank	2Rx4	2Rx4
Operating system		
Name	Red Hat® Enterprise Linux®	DAAD ION Accelerator
Build number	6.6	2.5.1-413

System	Dell PowerEdge R720	Dell Acceleration Appliance for Databases 2.0
File system	ext4	btrfs
Kernel	2.6.32-504.el6.x86_64	3.0.101-0.15.1.6651.0.PTF-default (x86_64)
Language	English	English
RAID controller		
Vendor and model number	Dell PERC H710P Mini	Dell PERC H730P Mini
Firmware version	21.3.0-0009	25.2.1.0037
Cache size (GB)	1	2
RAID configuration	1 × RAID50	1 × RAID 1
Hard disk types		
Hard disks (OS)		
Vendor and model number	Dell ST9146852SS	Dell ST300MM0006
Number of disks	16	2
Size (GB)	146	300
RPM	15K	10K
Type	SAS	SAS
PCI-e SSDs		
Vendor and model number	N/A	Fusion ioMemory SX300
Number of disks	N/A	4
Size (GB)	N/A	6400
Type	N/A	PCI-e
Ethernet adapters		
Vendor and model number	Intel I350 Quad-port 1Gb Daughter card	Broadcom® NetXtreme® BCM5720 Quad-port 1 GbE
Firmware	14.5.9	7.10.18
Type	rNDC	rNDC
Fibre Channel adapters		
Vendor and model number	QLogic® QLE2672 16Gb Fibre Channel Adapter	QLogic QLE2672 16Gb Fibre Channel adapter
Firmware	03.11.09	03.11.09
Type	PCI-e	PCI-e
USB ports		
Number	4	4
Type	USB 2.0	USB 2.0

Figure 2: System configuration information for the test systems.

APPENDIX C – CONFIGURING THE ORACLE DATABASE INFRASTRUCTURE

Configuration and test details

Traditional solution

For the traditional storage setup, we configured a Dell PowerEdge R720 server with Red Hat Enterprise Linux 6.6. We then configured a standalone Oracle Database 12*c* Enterprise Edition installation with a 5,000-warehouse database leveraging Oracle Automatic Storage Management, storing our data, logs, and backups on LUNs created on the traditional storage appliance. For our storage area network (SAN), we attached the R720's Fibre Channel HBA as well as the traditional solution's Fibre Channel HBAs to a Brocade SW6505 Fibre Channel switch. See Figure 3 for the layout of this configuration.

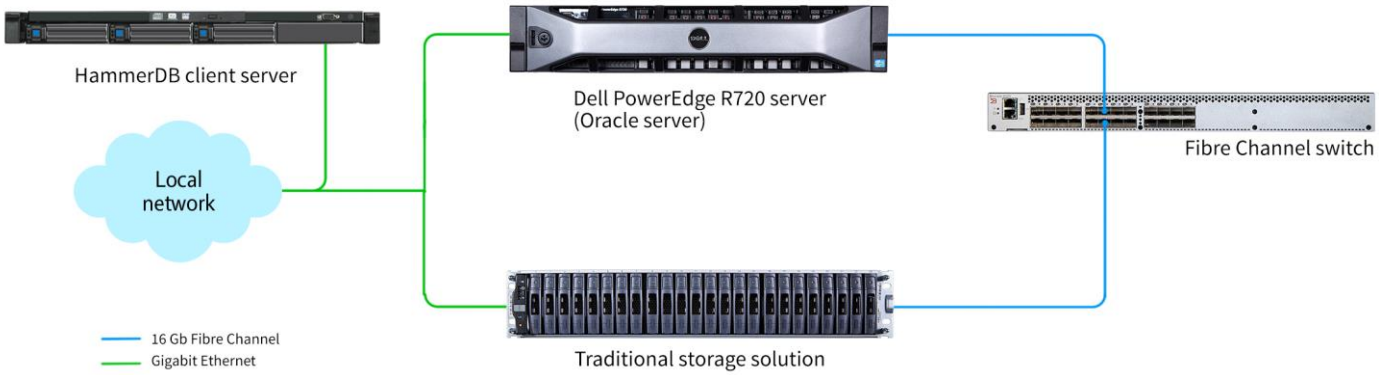

Figure 3: The traditional storage configuration we tested.

Combined DAAD/traditional storage solution

One of the benefits of the DAAD solution is being able to add it to your existing database infrastructure. We combined the highly available DAAD with our traditional storage solution, putting them into the same Fibre Channel SAN. We moved our backup and data LUNs to the DAADs, and the logs LUNs to the traditional storage for this configuration. See Figure 4 for the layout of this configuration.

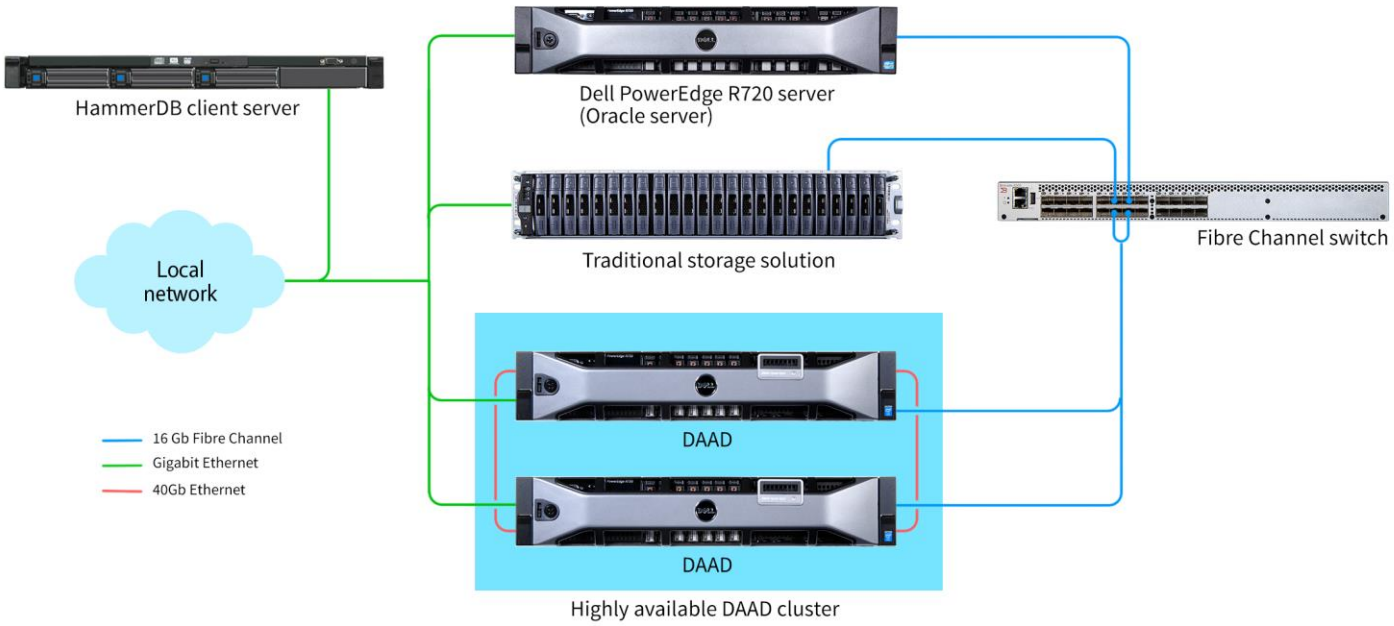

Figure 4: The combined traditional storage/highly available DAAD configuration we tested.

Dell Acceleration Appliance for Databases solution

For the Dell Acceleration Appliance for Databases solution, we replaced the traditional storage appliance with a highly available DAAD, storing our data, logs, and backups on LUNs created on the DAADs. We then attached the DAAD's 16Gb Fibre Channel HBAs to our infrastructure's Brocade Fibre Channel switch. See Figure 5 for the layout of this configuration.

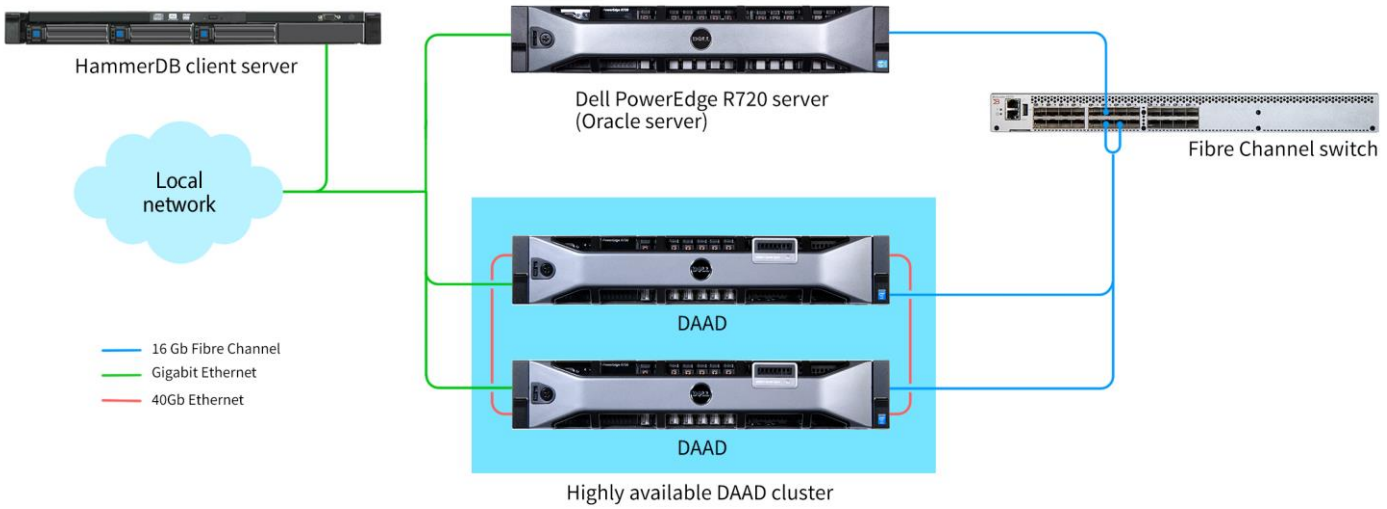

Figure 5: The highly available DAAD configuration we tested.

Setting up storage on the DAAD

1. In a terminal, ssh into one of the DAAD nodes with the admin credentials.
2. Enter the following commands to create mirrored volumes across the DAAD nodes, an initiator group, and then 8 LUNs to be presented to the database server:

```
profile:create direct
volume:create -n ion01 -n ion02 fcion_v_a1 50% jbod_pool-1
volume:create -n ion02 -n ion01 fcion_v_a2 100% jbod_pool-1
volume:create -n ion01 -n ion02 fcion_v_b1 50% jbod_pool-2
volume:create -n ion02 -n ion01 fcion_v_b2 100% jbod_pool-2
volume:create -n ion01 -n ion02 fcion_v_c1 50% jbod_pool-3
volume:create -n ion02 -n ion01 fcion_v_c2 100% jbod_pool-3
volume:create -n ion01 -n ion02 fcion_v_d1 50% jbod_pool-4
volume:create -n ion02 -n ion01 fcion_v_d2 100% jbod_pool-4
inigroup:create ig_all 20:01:00:0e:1e:09:d6:1c 20:01:00:0e:1e:09:d6:1d
lun:create fcion_v_a1 ig_all -b 512 –a
lun:create fcion_v_a2 ig_all -b 512 -a
lun:create fcion_v_b1 ig_all -b 512 -a
lun:create fcion_v_b2 ig_all -b 512 -a
lun:create fcion_v_c1 ig_all -b 512 -a
lun:create fcion_v_c2 ig_all -b 512 -a
lun:create fcion_v_d1 ig_all -b 512 –a
lun:create fcion_v_d2 ig_all -b 512 -a
```

Setting up storage on the traditional storage solution

We created four 300GB data LUNs, four 100GB logs LUNs, and two 1,000GB backup LUNs, and configured them for use over Fibre Channel.

Configuring Red Hat Enterprise Linux and Oracle Database 12*c*

We installed Red Hat Enterprise Linux on the Dell PowerEdge R720 server and configured settings as we specify below.

Installing Red Hat Enterprise Linux

1. Insert the Red Hat Enterprise Linux 6.6 DVD into the server, and boot to it.
2. Select Install or upgrade an existing system.
3. If you are unsure of the fidelity of your installation disk, select OK to test the installation media; otherwise, select Skip.
4. In the opening splash screen, select Next.
5. Choose the language you wish to use, and click Next.
6. Select the keyboard layout, and click Next.
7. Select Basic Storage Devices, and click Next.
8. Click Yes, discard any data at the Storage Device Warning.
9. Insert your hostname, and click Next.
10. Select the nearest city in your time zone, and click Next.
11. Enter your root password, and click Next.

12. Select Create Custom Layout, and click Next.
13. Select the install drive and click Create. (Create the following volumes and size: Root = 300GB, Home = 500GB, Boot = 200MB, SWAP = 20GB)
14. Click Next.
15. Click Write changes to disk.
16. Select the appropriate Data Store Devices and select where the Bootloader will go, and click Next.
17. Select Software Basic Server, and click Next. Linux installation begins.
18. When the installation completes, select Reboot to restart the server.

Performing initial configuration tasks

Complete the following steps to provide the base functionality that Oracle Database requires. We performed all of these tasks as root.

1. Disable SELINUX.
```
vi /etc/selinux/config
SELINUX=disabled
```
2. Set CPU Governor type.
```
vi /etc/sysconfig/cpuspeed
GOVERNOR=performance
```
3. Disable the firewall for IPv4 and IPv6.
```
chkconfig iptables off
chkconfig ip6tables off
```
4. To update the operating system packages, type the following:
```
yum update -y
```
5. To install additional packages, type the following commands:
```
yum install -y acpid cpuspeed wget vim nfs-utils openssh-clients man
lsscsi unzip smartmontools numactl ipmitool OpenIPMI
```
6. Reboot the server.
```
reboot
```
7. Install additional packages with the following commands:
```
yum install -y \
binutils \
compat-libcap1  \
compat-libstdc++-33  \
compat-libstdc++-33.i686  \
device-mapper-multipath \
gcc  \
gcc-c++  \
glibc  \
glibc.i686  \
glibc-devel  \
glibc-devel.i686  \
ksh  \
```

```
libgcc  \
libgcc.i686  \
libstdc++  \
libstdc++.i686  \
libstdc++-devel  \
libstdc++-devel.i686  \
libaio  \
libaio.i686  \
libaio-devel  \
libaio-devel.i686  \
libXext  \
libXext.i686  \
libXtst  \
libXtst.i686  \
libX11  \
libX11.i686  \
libXau  \
libXau.i686  \
libxcb  \
libxcb.i686  \
libXi  \
libXi.i686  \
make  \
sysstat  \
unixODBC  \
unixODBC-devel \
xorg-x11-xauth \
xorg-x11-utils
```

8. Edit the sysctl file.

```
vim /etc/sysctl.conf
fs.file-max = 6815744
kernel.sem = 250 32000 100 128
kernel.shmmni = 4096
kernel.shmall = 1073741824
kernel.shmmax = 4398046511104
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
fs.aio-max-nr = 1048576
net.ipv4.ip_local_port_range = 9000 65500
vm.nr_hugepages = 52248
vm.hugetlb_shm_group = 54321
```

9. Apply the changes with the following command:

```
sysctl -p
```

10. Edit the security limits configuration.

```
vim /etc/security/limits.conf
oracle  soft    nofile  1024
oracle  hard    nofile  65536
oracle  soft    nproc   2047
oracle  hard    nproc   16384
oracle  soft    stack   10240
oracle  hard    stack   32768
grid    -  nofile  65536
grid    -  nproc   16384
grid    -  stack   32768
oracle  soft    memlock 536870912
oracle  hard    memlock 536870912
```

11. Add the necessary groups and users.

```
groupadd -g 1001 oinstall
groupadd -g 1002 dba
groupadd -g 1003 asmadmin
groupadd -g 1004 asmdba
useradd -m -u 1002 -g oinstall -G dba,asmadmin,asmdba oracle
useradd -m -u 1003 -g oinstall -G dba,asmadmin,asmdba grid
```

12. Add the following lines to the .bash_profile file for the oracle user:

```
export TMP=/tmp
export TMPDIR=$TMP
export ORACLE_HOSTNAME=hostname
export ORACLE_UNQNAME=ORCL
export ORACLE_BASE=/u01/app/oracle
export GRID_HOME=/u01/app/12.1.0/grid
export DB_HOME=$ORACLE_BASE/product/12.1.0/dbhome_1
export ORACLE_HOME=$DB_HOME
export ORACLE_SID=orcl
export ORACLE_TERM=xterm
export BASE_PATH=/usr/sbin:$PATH
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export
CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib
```

13. Add the following lines to the .bash_profile file for the grid user:

```
export TMP=/tmp
export TMPDIR=$TMP
export ORACLE_HOSTNAME=hostname
export ORACLE_BASE=/u01/app/grid
export GRID_HOME=/u01/app/12.1.0/grid
export ORACLE_HOME=$GRID_HOME
export ORACLE_SID=+ASM1
```

```
export ORACLE_TERM=xterm
export BASE_PATH=/usr/sbin:$PATH
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export
CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib
```

14. Create the following directories, and assign the following permissions.

```
mkdir -p /u01/app/oracle
mkdir /u01/app/grid
chown -R oracle:oinstall /u01/app
chmod -R g+w /u01/app
```

15. Create passwords for the oracle and grid accounts with `passwd`.
16. Edit the hosts file.

```
vim /etc/hosts
127.0.0.1       R720 R720.localhost.localdomain localhost
localhost.localdomain localhost4 localhost4.localdomain4
::1             R720 R720.localhost.localdomain localhost
localhost.localdomain localhost6 localhost6.localdomain6
```

17. Edit the /etc/fstab file.

```
vim /etc/fstab
```

Adding this line:

```
tmpf /dev/shm     tmpfs   defaults,size=89G        0 0
```

18. Edit the 90-nproc.conf file.

```
vim /etc/security/limits.d/90-nproc.conf
```

Modifying this line:

```
*          soft    nproc    1024
```

To reflect this change:

```
* - nproc 16384
```

19. Download ION Tuner from Fusion-Io and install with `rpm -i iontuner_1.1.0_1.el6.noarch.rpm`. This is available to download from www.fusionio.com/files/ion-optimization-scripts.
20. Edit the /boot/grub/menu.lst file:

```
vim /boot/grub/menu.lst
```

Append this to the `kernel` line:

```
intel_idle.max_cstate=0 processor.max_cstate=0
```

21. Enable the multipath configuration file and start the multipath daemon:

```
mpathconf --enable --with_multipathd y
```

22. For the DAAD configurations, edit /etc/multipath.conf to add the following:

```
vim /etc/multipath.conf
defaults {
   user_friendly_names     yes
   queue_without_daemon    no
}
devices {
      device {
```

```
        vendor                  "FUSIONIO"
           features                "3 queue_if_no_path pg_init_retries 50"
           hardware_handler        "1 alua"            path_grouping_policy
           group_by_prio            path_selector           "queue-length 0"
           failback                immediate
           path_checker            tur
           prio                    alua

           fast_io_fail_tmo       15
           dev_loss_tmo           60
        }
    }
```

23. Perform a Loop Initialization Protocol and rescan the Fibre Channel interconnects with `echo "1" > /sys/class/fc_host/hostX/issue_lip`, replacing X with each of the hosts being used on the system.
24. Restart the multipath service with `service multipathd restart`.
25. Verify that multipath is working with `multipath -ll`.
26. Create a primary partition with the fdisk utility on each multipath disk with the following:
    a. `fdisk /dev/mapper/mpathX`, where X is the identifier of the disk in question.
    b. `n`
    c. `p`
    d. `1`
    e. `1`
    f. Press Enter.
    g. `w`
27. Reboot the server to make the kernel use the new table.
28. Edit the scsi_id file.
    `echo "options=-g" > /etc/scsi_id.config`

Providing LUNs to ASM

We used the steps below to configure the multipathed storage LUNs prior to setting up ASM.

1. Edit the 99-oracle-asmdevices rules file.

```
vim /etc/udev/rules.d/99-oracle-asmdevices.rules
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-26230373534643030",
SYMLINK+="oracleasm/DAADdata1", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-26565316265346439",
SYMLINK+="oracleasm/DAADdata2", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-23261306136633365",
SYMLINK+="oracleasm/DAADdata3", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-23436316433663237",
SYMLINK+="oracleasm/DAADdata4", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-26235623336316364",
SYMLINK+="oracleasm/DAADdata5", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-26364393665633631",
SYMLINK+="oracleasm/DAADdata6", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-26539376463663337",
SYMLINK+="oracleasm/DAADlogs1", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-23736666233353565",
SYMLINK+="oracleasm/DAADlogs2", OWNER="oracle", GROUP="dba", MODE="0660"
```

```
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-23534343539323132",
SYMLINK+="oracleasm/TRADITIONALdata1", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e6939",
SYMLINK+="oracleasm/TRADITIONALdata2", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e6941",
SYMLINK+="oracleasm/TRADITIONALdata3", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e6943",
SYMLINK+="oracleasm/TRADITIONALdata4", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e6945",
SYMLINK+="oracleasm/TRADITIONALlogs1", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e6947",
SYMLINK+="oracleasm/TRADITIONALlogs2", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e6949",
SYMLINK+="oracleasm/TRADITIONALlogs3", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e694b",
SYMLINK+="oracleasm/TRADITIONALlogs4", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e694d",
SYMLINK+="oracleasm/TRADITIONALbackup1", OWNER="oracle", GROUP="dba",
MODE="0660"
KERNEL=="dm-*", ENV{DM_UUID}=="mpath-360a980004431455a345d4733696e694f",
SYMLINK+="oracleasm/TRADITIONALbackup2", OWNER="oracle", GROUP="dba",
MODE="0660"
```

2. Execute udevadm and start udev.

```
udevadm control --reload-rules
start_udev
```

3. List the ASM devices.

```
ls -l /dev/oracleasm/
```

```
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADdata1 -> ../dm-17
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADdata2 -> ../dm-18
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADdata3 -> ../dm-13
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADdata4 -> ../dm-19
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADdata5 -> ../dm-14
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADdata6 -> ../dm-20
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADlogs1 -> ../dm-15
lrwxrwxrwx 1 root root 8 Mar 20 11:29 DAADlogs2 -> ../dm-16
lrwxrwxrwx 1 root root 8 Mar 20 11:29 TRADITIONALbackup1 -> ../dm-10
lrwxrwxrwx 1 root root 8 Mar 20 11:29 TRADITIONALbackup2 -> ../dm-11
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALdata1 -> ../dm-3
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALdata2 -> ../dm-4
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALdata3 -> ../dm-2
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALdata4 -> ../dm-5
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALlogs1 -> ../dm-7
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALlogs2 -> ../dm-6
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALlogs3 -> ../dm-8
lrwxrwxrwx 1 root root 7 Mar 20 11:29 TRADITIONALlogs4 -> ../dm-9
```

Installing Oracle Grid Infrastructure for Standalone Server 12*c*

Prior to starting the steps below, we downloaded the Oracle 12*c* Grid installation and extracted it to the /grid directory.

1. Run the GUI installer for Oracle Database using the following commands:
   ```
   ssh –Y grid@R720_IP_address
   cd /grid
   ./runInstaller
   ```
2. Launch the Oracle Grid Infrastructure installation wizard.
3. In Software Updates, select Skip software updates, and click Next.
4. In Installation Options, select Install and Configure Oracle Grid Infrastructure for a Standalone Server, and click Next.
5. In Product Languages, select English, and click the right-pointing arrow between the two selection panels to add English to the Selected languages panel. Click Next.
6. In Create ASM Disk Group, set the Disk group name to DATA.
7. Click Change Discovery Path.
8. Enter `/dev/oracleasm` for the Disk Discovery Path, and click OK.
9. Check the boxes for the data disks, and Click Next.
   a. For the DAAD-only configuration, select only the DAADs' data disks.
   b. For the traditional storage configuration, select only the traditional storage's data disks.
10. In ASM Password, select Use same passwords for these accounts. Enter and confirm the password, and click Next.
11. In Operating System Groups, set Oracle ASM Administrator Group to asmadmin, and Oracle ASM DBA Group to asmdba, and click Next.
12. In Installation Location, accept the default locations provided, and click Next.
13. In Create Inventory, accept the defaults, and click Next.
14. In Root Script Execution, check the box for Automatically run configuration scripts.
15. Select Use "root" user credential, and provide the root password. Click Next.
16. In Summary, review the information, and click Install to begin installation.
17. Click Yes to confirm using the privileged user for the installer.
18. In Finish, click Close to exit the installer.

Installing Oracle Database 12*c*

Prior to starting the steps below, we downloaded the Oracle Database 12*c* installation and extracted it to the /database directory.

1. Run the GUI installer for Oracle Database using the following commands:
   ```
   ssh –Y oracle@R720_IP_address
   cd /database
   ./runInstaller
   ```
2. Launch the Oracle Database 12*c* Release 1 Installer.
3. In Configure Security Updates, uncheck the I wish to receive security updates via My Oracle Support checkbox. Click Next.
4. Click Yes to confirm no email provided, and continue.
5. In Software Updates, select Skip software updates, and click Next.
6. In Installation Options, select Install database software only, and click Next.

7. In Grid Installation Options, select Single instance database installation, and click Next.
8. In Product Languages, select English and click the right-pointing arrow located between the two selection panels to add English to the Selected languages panel. Click Next.
9. In Database Edition, select Enterprise Edition, and click Next.
10. In Installation Location, accept the default locations provided, and click Next.
11. In Operating System Groups, accept the defaults, and click Next.
12. In Summary, review the information, and click Install to begin installation.
13. When prompted, follow the instructions to execute the scripts. Click OK when the scripts have completed.
14. In Finish, click Close to exit the installer.

Creating Oracle ASM disk groups for the database

1. Log into the system as the grid user.
2. Start the ASM configuration assistant, asmca.
3. On the Disk Groups tab, click Create.
4. On the Create Disk Group pop-up screen, enter LOGS for the Disk Group Name.
5. Select External (None) for Redundancy.
6. Select /dev/oracleasm/ for the Disk Path.
7. Select the logs disks.
    a. To start with the DAAD-only configuration, select only the DAADs' logs disks.
    b. When switching to another configuration later, launch asmca again, and perform the Drop Disks and Add Disks operations as necessary to achieve the correct arrangements.
        i. For the traditional storage configuration, the DATA disk group should contain the traditional storage's data disks, the LOGS disk group should contain the traditional storage's logs disks, and the BACKUP disk group should contain the traditional storage's backup disks.
        ii. For the mixed configuration, the DATA disk group should contain the DAAD disks, the LOGS disk group should contain the traditional storage's disks, and the BACKUP disk group should still contain the traditional storage's backup disks.
        iii. For the DAAD-only configuration, the DATA disk group should contain six of the DAAD disks, the LOGS disk group should contain the other two DAAD disks, and the BACKUP disk group should still contain the traditional storage's backup disks.
8. Click OK to create the LOGS disk group, and click OK on the completed-task pop-up screen.
9. Repeat steps 4-8 to create a BACKUP disk group with the remaining BACKUP disks.
10. Click Exit to close the ASM configuration assistant.

Creating the Oracle Database (using DBCA)

1. Log into the system as the oracle user.
2. Launch the Database Configuration Assistant (DBCA).
3. In Database Operations, select Create Database, and click Next.
4. In Creation Mode, select Advanced Mode, and click Next.
5. In Database Template, select the Template for General Purpose or Transaction Processing, and click Next.

6. If prompted about continuing without upgrade, click Yes.
7. In Database Identification, type `orcl` for the Global Database Name.
8. Type `orcl` for the SID. Click Next.
9. In Management Options, select Configure Enterprise Manager (EM) Database Express. Click Next.
10. In Database Credentials, select Use the Same Administrative Password for All Accounts.
11. Enter and confirm the administrative password, and click Next.
12. In Network Configuration, check the boxes for all listeners, and click Next.
13. In Storage Locations, select ASM for Database Storage Type.
14. Select User Common Location for All Database Files, and type +DATA into the Database Files Location field.
15. Select ASM for Recovery files Storage Type.
16. Specify Fast Recovery Area. Type +BACKUP in the Fast Recovery Area field.
17. Set the Fast Recovery Area size to 1850 GB, and click Next.
18. In Database Options, accept the defaults, and click Next.
19. In Initialization Parameters and under typical settings, set the Memory Size to 40%, and click Next.
20. In Creation Options, select Create Database.
21. Click Next.
22. Review the Summary. To complete the database creation, click Finish.
23. Review the information on the screen, and click Exit.
24. To exit the DBCA, click Close.

Configuring the Oracle Tablespace and redo logs

Alter the tablespace and redo logs as follows:

```
ALTER DATABASE ADD LOGFILE GROUP 4 ( '/tmp/temp1.log' ) SIZE 50M;
ALTER DATABASE ADD LOGFILE GROUP 5 ( '/tmp/temp2.log' ) SIZE 50M;

ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM CHECKPOINT;

ALTER DATABASE DROP LOGFILE GROUP 1;
ALTER DATABASE DROP LOGFILE GROUP 2;
ALTER DATABASE DROP LOGFILE GROUP 3;

ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM CHECKPOINT;

ALTER DATABASE DROP LOGFILE GROUP 1;
ALTER DATABASE DROP LOGFILE GROUP 2;
ALTER DATABASE DROP LOGFILE GROUP 3;

-- DELETE LOGS IN ASM HERE --

ALTER DATABASE ADD LOGFILE GROUP 1 ( '+LOGS/ORCL/ONLINELOG/redo01.log' ) SIZE 20G;
ALTER DATABASE ADD LOGFILE GROUP 2 ( '+LOGS/ORCL/ONLINELOG/redo02.log' ) SIZE 20G;
ALTER DATABASE ADD LOGFILE GROUP 3 ( '+LOGS/ORCL/ONLINELOG/redo03.log' ) SIZE 20G;
```

```
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM CHECKPOINT;

ALTER DATABASE DROP LOGFILE GROUP 4;
ALTER DATABASE DROP LOGFILE GROUP 5;

ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM CHECKPOINT;

ALTER DATABASE DROP LOGFILE GROUP 4;
ALTER DATABASE DROP LOGFILE GROUP 5;

HOST rm -f /tmp/temp*.log

CREATE BIGFILE TABLESPACE "TPCC"
 DATAFILE '+DATA/orcl/tpcc.dbf' SIZE 400G AUTOEXTEND ON NEXT 1G
 BLOCKSIZE 8K
 EXTENT MANAGEMENT LOCAL AUTOALLOCATE
 SEGMENT SPACE MANAGEMENT AUTO;

CREATE BIGFILE TABLESPACE "TPCC_OL"
 DATAFILE '+DATA/orcl/tpcc_ol.dbf' SIZE 150G AUTOEXTEND ON NEXT 1G
 BLOCKSIZE 16K
 EXTENT MANAGEMENT LOCAL AUTOALLOCATE
 SEGMENT SPACE MANAGEMENT AUTO;

ALTER DATABASE DATAFILE '+DATA/ORCL/DATAFILE/UNDOTBS1.260.873465895' RESIZE
32760M;
```

Configuring the Oracle pfile

Alter the Oracle pfile as follows:

```
orcl.__oracle_base='/u01/app/oracle'#ORACLE_BASE set from environment
_enable_NUMA_support=TRUE
_kgl_hot_object_copies=2
_shared_io_pool_size=512m
audit_file_dest='/u01/app/oracle/admin/orcl/adump'
audit_trail='NONE'
compatible='12.1.0.2.0'
control_files='+DATA/ORCL/CONTROLFILE/current.261.873465981','+BACKUP/ORCL/CONTRO
LFILE/current.256.873465981'#Restore Controlfile
db_16k_cache_size=16g
db_block_size=8192
db_cache_size=64g
db_create_file_dest='+DATA'
db_domain=''
db_name='orcl'
db_recovery_file_dest_size=1850g
db_recovery_file_dest='+BACKUP'
diagnostic_dest='/u01/app/oracle'
disk_asynch_io=TRUE
```

```
dispatchers='(PROTOCOL=TCP) (SERVICE=orclXDB)'
dml_locks=500
fast_start_mttr_target=300
java_pool_size=4g
job_queue_processes=0
large_pool_size=1g
local_listener='LISTENER_ORCL'
lock_sga=TRUE
log_buffer=536870912#log buffer update
log_checkpoint_interval=0
log_checkpoint_timeout=0
log_checkpoints_to_alert=TRUE
open_cursors=3000
parallel_max_servers=0
pga_aggregate_target=6g
recovery_parallelism=4
remote_login_passwordfile='EXCLUSIVE'
resource_manager_plan='FORCE'
shared_pool_size=3g
undo_management='AUTO'
undo_retention=1
undo_tablespace='UNDOTBS1'
use_large_pages='ONLY'
```

Setting up the HammerDB client

We used a dual-socket server running Red Hat Enterprise Linux 6.6 for the HammerDB client. We followed the installation steps at the beginning of this appendix to install Red Hat Enterprise Linux, and also installed the GUI. We then installed the HammerDB client software.

Installing HammerDB

Download and install version 2.16 on the Red Hat client. We downloaded HammerDB from the following location: hammerora.sourceforge.net/download.html. We installed HammerDB according to the installation guide (hammerora.sourceforge.net/hammerdb_install_guide.pdf).

Installing HammerDB Oracle libraries

Complete the following steps on the application server.

1. Launch the Oracle Client Installer.
2. In Select Installation Type, select Administrator (1.8 GB) as the installation type, and click Next.
3. In Software Updates, select Skip software updates, and click Next.
4. In Select Product Languages, select English and click the right-pointing arrow located between the two selection panels to add English to the Selected languages panel. Click Next.
5. In Specify Installation Location, accept the default locations provided, and click Next.
6. In Create Inventory, accept the defaults, and click Next.
7. In Summary, review the information, and click Install to begin installation.
8. In Install Product, follow the instructions to execute the scripts. Click OK when the scripts have completed.
9. In Finish, click Close to exit the installer.

Configuring the database

We used the TPC-C build schema build options for Oracle inside HammerDB to build the database. We set the following options in the build schema:

```
Oracle Service Name = R720_IP_address/orcl
System User = SYSTEM
System User Password = Password1
TPC-C User = tpcc
TPC-C User Password = tpcc
TPC-C Default Tablespace = tpcc
Order Line Tablespace = tpcc_ol
TPC-C Temporary Tablespace = temp
TimesTen Database Compatible = unchecked
Partition Order Line Table = checked
Number of Warehouses = 5000
Virtual Users to Build Schema = 20
Use PL/SQL Server Side Load = unchecked
Server Side Log Directory = /tmp
```

Running HammerDB

We ran HammerDB by filling in the appropriate information for the driver options. We tested with 10,000,000 transactions per user, a 20-minute ramp up time, and a 20-minute test duration. We used 101 virtual users with 0-ms user delay and repeat delay. We used rman to back up the database before testing, and we ran restores between runs.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.