Technical Report: HP-UX NFS Performance with Oracle Database 10*g* Using NetApp Storage

Dave Olker
Technical Consultant
HP

John Elliott
Database Performance Engineer
NetApp, Inc.

November 2008 | TR-3557

Abstract

This technical report contains performance results and tuning recommendations for Oracle® Database 10g running on HP-UX 11i using NetApp® storage over NFS. The focus of this paper is technical and the reader should be experienced with HP-UX system administration, Oracle Database 10g administration, network connectivity, and NetApp storage system administration.

Table of Contents

1. Introduction 3

2. Hardware Configuration 4

2.1. HP-UX Server Hardware 4

2.2. NetApp Storage System Hardware 4

3. Database Workload Description and Test Results 5

4. NetApp Storage System Tuning 6

5. Oracle Configuration on HP-UX Systems 7

5.1. HP-UX Kernel Parameter Tuning for Oracle Workloads 7

5.2. Oracle Configuration Parameter Tuning 11

5.3 Oracle NFS-Mounted File Systems 12

6. HP-UX 11i v2 / 11i v3 Common Recommendations for Network Performance (Including NFS) 12

6.1. TCP Window Sizes 13

6.2. TCP Segmentation Offload and Checksum Offload 13

6.3. NFS Mount Options 14

7. HP-UX 11i v2 OS-Specific Recommendations 15

7.1. HP-UX 11i v2 Kernel Parameter Recommendations 15

7.2. NFS Block I/O Daemons (BIOD Processes) 16

7.3. HP-UX 11i v2 Patch Recommendations 17

8. HP-UX 11i v3 OS-Specific Recommendations 22

8.1. HP-UX 11i v3 Kernel Parameter Recommendations 22

8.2. HP-UX 11i v3 Patch Recommendations 22

9. Conclusions 24

10. Acknowledgements 25

1. Introduction

Oracle Databases running on HP-UX servers with NetApp storage systems provide a very powerful environment for enterprise-class Oracle workloads. Together they offer stability, high performance, scalability, and a rich set of tools for serving, monitoring, managing, and protecting mission critical data.

This technical report contains performance data and tuning recommendations for Oracle Database 10*g* running on HP-UX 11i servers, using NetApp storage systems with NFS. NetApp storage systems and Hewlett-Packard servers (and software) are designed to implement the latest computing technologies and to provide the highest level of performance for a wide range of business and scientific applications.

Those applications can be quite different in the ways they manage and utilize resources. As a result, it becomes imperative that those resources be tuned to meet the requirements of the applications they host (including Oracle Databases). The information presented in this document will make that very clear. In fact, we were able to increase the throughput of a test database running on HP-UX 11i v2 by around 640% by simply applying a few changes to the default configuration through a series of tuning and test iterations. Those changes included the implementation of direct I/O, HP-UX kernel tuning, storage tuning, and network tuning.

The first part of this document describes the HP and NetApp hardware used during this exercise. The next section includes a description of our test workload along with actual test results obtained using the hardware described in the previous section and the tuning recommendations that follow. The next part of this document contains a detailed list of recommendations for configuring the Oracle application on HP-UX systems with NetApp storage. After that we'll focus on tuning recommendations common to both HP-UX 11i v2 and 11i v3 systems, followed by a listing of OS-specific recommendations applicable to 11i v2 and v3 systems, respectively.

2. Hardware Configuration

2.1. HP-UX Server Hardware

For this exercise we used a Hewlett-Packard Integrity server configured as described in Table 1.

Table 1) HP server hardware

Component	Details
Operating system	HP-UX 11i v2 mission-critical operating environment component
Version	B.11.23.0603 U ia64
System type	rx2600
Database server	Oracle 10.2.0.2.0 64-bit
Total physical RAM	6GB
Processor	2 Intel® Itanium-2 CPUs 1.4 GHz
Storage network	2 X HP-A6825 GbE NICs for NFS, MTU Size 9000 (Jumbo Frames)

2.2. NetApp Storage System Hardware

We used a NetApp FAS3050 storage system configured as described in Table 2. The NFS client and server systems were interconnected as displayed in Figure 1.

Table 2) NetApp storage system configuration

Component	Details
Operating system	Data ONTAP® 7.2.1
Storage interconnect	2X1Gb Ethernet for NFS
Disks	8 DS14s of 72GB, 10K RPM disks
Storage controller	Single FAS3050
Disk shelf to storage controller connection	2 backside FCAL with 4 shelves per FCAL loop
Storage switches	Database server and storage system were direct connected with crossover cables

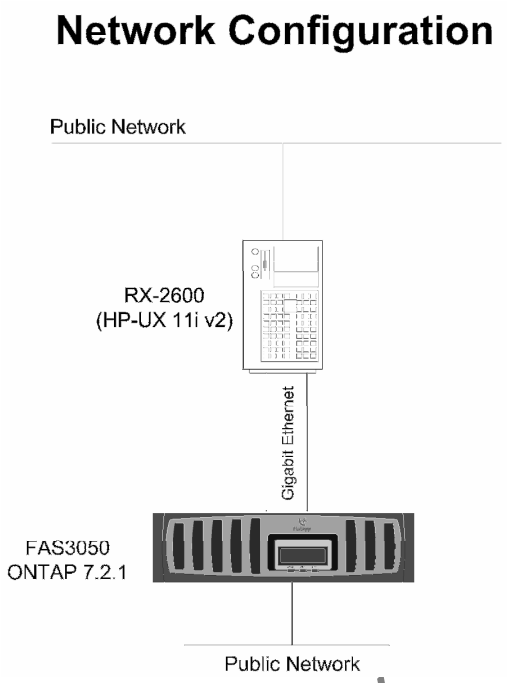

Figure 1) Schematic of HP-UX server and NetApp storage for the NFS test environment

3. Database Workload Description and Test Results

The test environment was designed to stress the database server software and hardware with the goal of utilizing all database server CPU resources for each test run.

The database used for testing can best be described as OLTP in nature with a physical size of approximately 350GB. After creating our database, we generated an OLTP-type workload consisting of a steady stream of small, random read and write operations (approximately 2:1 reads to writes). This workload was designed to emulate the real-life activities of a wholesale supplier order processing system in which inventory is spread across several regional warehouses.

The database utilized both primary and secondary keys for data access. In terms of measured database throughput, the metric of interest was defined as the number of “order entry transactions per minute,” or simply OETs.

The test was run in *server-only* mode. This means that all the user processes and the database engine were running on the rx2600 server. Server-only mode also implies that the user processes are running without any think times. This means that the users continually submit transactions without simulating any delay between those transactions.

Test results in “order entry transactions per minute,” before and after tuning, are recorded in Figure 2.

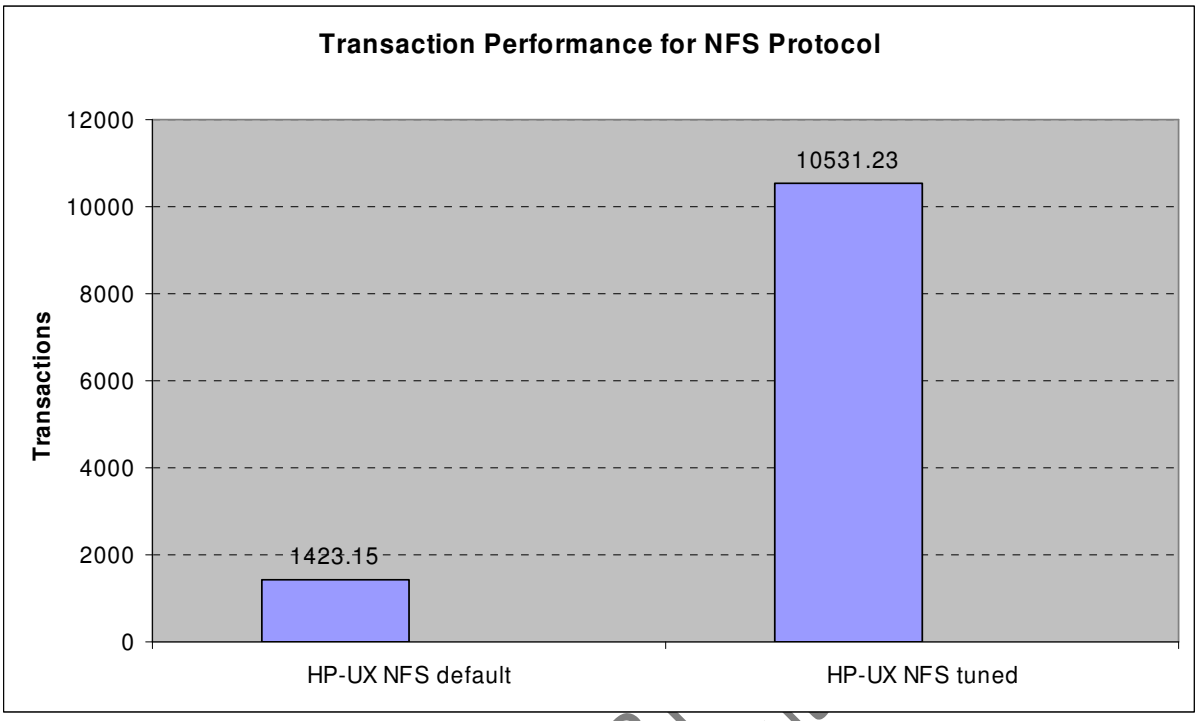

Figure 2) Order entry transactions per minute for untuned and tuned configurations

To achieve these results, we conducted a series of testing and tuning iterations in which several Oracle, HP-UX, Data ONTAP, and network configuration parameters were modified and the resulting performance improvement measured. As you can see in Figure 2, we saw performance improve approximately 640%, going from 1,423 OETs using default configuration parameters to 10,531 OETs using our tuned configuration. That 640% includes about 482% performance improvement obtained by simply enabling direct I/O. The following sections provide more information about the NetApp storage, HP-UX, and Oracle settings used in the tuned configuration.

4. NetApp Storage System Tuning

Table 3 lists the recommended Data ONTAP settings for best database performance.

Table 3) Recommended Data ONTAP settings for optimal database performance

Data ONTAP Setting	Recommended Value	Description
nfs.tcp.recvwindowsize	262144	Maximize the TCP window size
nfs.tcp.xfersize	65536	Maximize the NFS transfer size

Data ONTAP Setting	Recommended Value	Description
no_atime_update	On (for data file, log, and control file volumes)	Prevent access time update for inodes

Storage Provisioning Details

All the disks from eight shelves of 72GB 10K RPM drives were placed in two aggregates. The aggregates were created using RAID-DP® with 16 disks in a RAID group. One FlexVol® root volume, vol0, was created inside the aggregate aggr0, and two FlexVol volumes, data1 and log1, were created inside the aggregate aggr1:

aggr0 – 3 disks, RAID-DP, RAID size 16
    /vol/vol0 flexvol, root volume
aggr1 – 104 disks, RAID-DP, RAID size 16
    /vol/oradata1 flexvol, 1600GB, used for Oracle datafiles
    /vol/oralog1 flexvol, 200GB, used for Oracle redo logs and control files

5. Oracle Configuration on HP-UX Systems

5.1. HP-UX Kernel Parameter Tuning for Oracle Workloads

HP-UX 11i provides numerous kernel tunable parameters. These parameters are designed to allow administrators to adjust the system's behavior and performance to meet the needs of the application workload of a given system. This means the list of tunable values appropriate for a system running Oracle will likely differ from a system running a different application. Proper tuning of the HP-UX kernel is essential to achieve optimal application performance and efficient use of system resources.

The complete list of kernel tunables on HP-UX 11i systems is quite large (well over 300 tunables, when both public and private tunables are considered). The vast majority of these tunables can be left at their default values for most application workloads to run optimally. However, in certain cases these tunables need to be adjusted to support a specific application, such as Oracle.

HP-UX Kernel Tuning Using the kctune(1M) Command

HP offers many different tools to manage kernel tunables, including the System Administration Manager (SAM) tool on HP-UX 11i v2 and the System Management Homepage (SMH) tool on HP-UX 11i v3. However, neither of these tools can view or modify the many private kernel tunable parameters provided on HP-UX systems. For this reason, HP recommends using the kctune(1M) command to view and modify kernel tunable parameters. Kctune is described in detail in the "HP-UX System Administrator's Guide: Configuration Management" paper available at http://docs.hp.com/en/5992-4607/5992-4607.pdf.

The Tune-N-Tools Bundle and the Server-Tunables Product

HP recently released HP-UX 11i v3 Update 3, and one of the improvements included with this operating system is a product called Server-Tunables. This product includes a new tool called tuneserver(1M), which may be used to easily configure an HP-UX 11i v3 system to optimally run Oracle workloads. HP has released a white paper called “Server Tuning on HP-UX” that discusses this tool:

http://docs.hp.com/en/5992-4222ENW/5992-4222ENW.pdf.

The Server-Tunables product is included as part of the new Tune-N-Tools bundle, which may be downloaded free of charge from http://software.hp.com/portal/swdepot/displayProductInfo.do?productNumber=Tune-N-Tools.

Recommended Common HP-UX Kernel Parameter Settings for Oracle Workloads

Table 4 lists the set of kernel tunable parameter values recommended for an HP-UX 11i v2 or 11i v3 system running an Oracle workload. Most of these values were generated by the tuneserver(1M) command, while others were taken from Oracle installation guides.

Table 4) HP-UX recommended kernel parameter settings for Oracle workloads

Kernel Parameter	Recommended Value	Default Value	Details
max_async_ports	27000	50 (11i v2) 4096 (11i v3)	Maximum number of asynchronous disk ports that can be open at any time
max_thread_proc	3000	256	Maximum number of concurrent threads allowed per process
maxdsiz	3221225472 (3GB)	1GB	Maximum size (in bytes) of the data segment for any 32-bit user process
maxdsiz_64bit	274877906944 (256GB)	4GB	Maximum size (in bytes) of the data segment for any 64-bit user process
maxfiles	8192	2048	Initial (soft) maximum number of file descriptors per process
maxfiles_lim	8192	4096	Hard maximum number of file descriptors per process
maxssiz	134217728 (128MB)	8MB	Maximum size (in bytes) of the stack for any 32-bit user process

Kernel Parameter	Recommended Value	Default Value	Details
maxssiz_64bit	1073741824 (1GB)	256MB	Maximum size (in bytes) of the stack for any 64-bit user process
maxtsiz	1073741824 (1GB)	96MB	Maximum size (in bytes) of the text segment for any 32-bit user process
maxtsiz_64bit	8589934592 (8GB)	1GB	Maximum size (in bytes) of the text segment for any 64-bit user process
maxuprc	27000	256	Limits the maximum number of concurrent user processes per user
msgmnb	65536	16384	Maximum number of bytes on a single System V IPC message queue
msgmni	4096	512	Maximum number of system-wide System V IPC message queues allowed
msgtql	4096	1024	Maximum number of System V IPC messages in the system at any time
ncsize	36672	8976	Number of Directory Name Lookup Cache (DNLC) entries
nflocks	4200	4096	Maximum number of file/record locks allowed on the system at any time
nkthread	250000	((nproc*2)+16)	Limits the number of threads allowed to run simultaneously
nproc	30000	4200	Limits the number of processes allowed to exist simultaneously
npty	200	60	Maximum number of BSD pseudo terminals (ptys)
nstrpty	200	60	Maximum number of STREAMS-based pseudo terminals (pts)

Kernel Parameter	Recommended Value	Default Value	Details
o_sync_is_o_dsync	1	0	Enables translation of O_SYNC to O_DSYNC in open(2) / fcntl(2) calls
semmni	8192	2048	Number of System V IPC system-wide semaphore identifiers
semmns	60000	4096	Number of System V IPC system-wide semaphores
semmnu	27000	256	Maximum number of System V IPC undo structures for processes
semmsl	128	2048	Maximum number of System V IPC semaphores per identifier
semume	512	100	Maximum number of System V IPC undo entries per process
shmmax	4398046511104 (4096GB)	1GB	Maximum size (in bytes) for a System V shared memory segment
shmmni	4096	400	Number of System V shared memory segment identifiers in the system
shmseg	512	300	Maximum number of System V shared memory segments per process
swchunk	65536	2048	Swap chunk size in 1 KB blocks
vps_ceiling	64	16	Maximum (in kilobytes) of system-selectable page size
vps_chatr_ceiling	4194304 (4GB)	1GB	Maximum (in kilobytes) of user selectable page size

File System Data Cache Tuning

Another important area of HP-UX kernel tuning applicable to Oracle workloads is the size of the cache used to hold file system data. On 11i v2 systems this is known as the Dynamic Buffer Cache and is

controlled by the tunables *dbc_min_pct* and *dbc_max_pct*. On 11i v3 this cache is called the Unified File Cache and is controlled by the tunables *filecache_min* and *filecache_max*. These tunables control the percentage of physical memory in the system dedicated to file data caching. Their default values are 5% (dbc_min_pct or filecache_min) and 50% (dbc_max_pct or filecache_max).

On most Oracle systems the use of file data caching is discouraged since the cache resources are managed directly using the Oracle SGA. In cases such as these, the system administrator should strongly consider decreasing the default size of the file data cache by tuning dbc_max_pct or filecache_max tunable (11i v2 or 11i v3, respectively). On most Oracle systems this value can be tuned down to 25% of physical memory or less.

5.2. Oracle Configuration Parameter Tuning

Table 5 contains a list of key Oracle initialization parameters along with guidelines for setting them.

Table 5) Oracle Database 10*g* initialization parameter settings

Oracle Parameter	Recommended Value	Guidelines
cpu_count	Workload dependent	A good aggressive starting value is twice the number of CPUs in the database host.
db_writer_processes	Workload dependent	In many cases where asynchronous I/O is in use, a value of “1” is sufficient. However, multiple db writers may be beneficial even with asynchronous I/O. A good aggressive starting value is twice the number of host CPUs.
disk_asynch_io	Workload dependent	See section 3, “HP-UX NFS configuration,” for guidelines on when to use async I/O vs. direct I/O. HP-UX 11i v3 and previous releases do not support asynchronous direct I/O. disk_asynch_io should be set to “true” if async I/O is to be used and false for directio.
filesystemio_options	setall	NetApp best practices documentation recommends a setting of “setall.”

Oracle Parameter	Recommended Value	Guidelines
_kgl_large_heap_warning_threshhold	4194304	Per Metalink Note: 330239.1 to avoid KGL heap size warning messages with Oracle Database 10*g* R2.
hpux_sched_noage	178	Per Metalink Note: 217990.1
db_cache_size, log_buffer, shared_pool_size, etc	Workload dependent	There are no guidelines specific to HP-UX.

5.3 Oracle NFS-Mounted File Systems

Table 6 lists the NFS file systems used during this exercise and their usage by the Oracle application.

Table 6) NFS mount and Oracle usage

NFS Mount	Oracle Usage
/mnt/oradata1	Oracle data files
/mnt/oradata2	Oracle data files
/mnt/oralog1	Oracle redo log and control files
/mnt/oralog2	Oracle redo log and control files

6. HP-UX 11i v2 / 11i v3 Common Recommendations for Network Performance (Including NFS)

This section discusses the many networking features common to both HP-UX 11i v2 and 11i v3 that can have a significant effect on NFS performance, including TCP window sizes, TCP segmentation offload, checksum offload, and NFS mount options. For more information on the many ways to tune the HP-UX networking subsystem for optimal performance, see the “HP-UX TCP/IP Performance” white paper available at http://docs.hp.com/en/11890/perf-whitepaper-tcpip-v1_1.pdf. Later sections will describe the performance recommendations specific to each version of HP-UX 11i.

6.1. TCP Window Sizes

To increase the TCP send and receive window sizes to 16MB on HP-UX systems, add the following lines to the /etc/rc.config.d/nddconf file and issue the command “ndd –c” to activate the changes:

```
TRANSPORT_NAME[0]=tcp
NDD_NAME[0]=tcp_recv_hiwater_def
NDD_VALUE[0]=16777216

TRANSPORT_NAME[1]=tcp
NDD_NAME[0]=tcp_xmit_hiwater_def
NDD_VALUE[0]=16777216
```

6.2. TCP Segmentation Offload and Checksum Offload

The TCP Segmentation Offload (TSO) and Checksum Offload (CKO) features are supported on most Gigabit and 10 Gigabit Ethernet adapters available on HP-UX systems. To enable TSO and CKO on a specific Gigabit or 10 Gigabit Ethernet adapter, first determine the type of interface installed in your system. The ioscan(1M) command, shown below, lists the type of networking interfaces installed in the system and the driver used by each interface:

```
# ioscan -fk -C lan
Class     I  H/W Path       Driver   Description
==========================================================================
lan       0  0/0/8/1/0/4/0  igelan   HP A6794-60001 PCI 1000Base-T
lan       1  0/0/12/1/0     iether   HP A7012-60001 PCI/PCI-X 1000Base-T Dual-port Adapter
lan       2  0/0/12/1/1     iether   HP A7012-60001 PCI/PCI-X 1000Base-T Dual-port Adapter
lan       3  0/0/14/1/0     iether   HP A7012-60001 PCI/PCI-X 1000Base-T Dual-port Adapter
lan       4  0/0/14/1/1     iether   HP A7012-60001 PCI/PCI-X 1000Base-T Dual-port Adapter
```


igelan Devices

For igelan Gigabit Ethernet interfaces, modify the /etc/rc.config.d/hpigelanconf file as follows:

```
HP_IGELAN_INTERFACE_NAME[n]=lann
HP_IGELAN_VMTU[n]=32160
HP_IGELAN_SEND_CKO[n]=ON
HP_IGELAN_RECV_CKO[n]=ON
```

Either reboot the HP-UX server or issue the following commands to restart the driver and thereby activate the change:

```
/sbin/init.d/hpigelan stop
/sbin/init.d/hpigelan start
```

iether Devices

For iether Gigabit Ethernet interfaces, modify the /etc/rc.config.d/hpietherconf file as follows:

HP_IETHER_INTERFACE_NAME[*n*]=lan*n*
HP_IETHER_VMTU[*n*]=32160
HP_IETHER_SEND_CKO[*n*]=ON
HP_IETHER_RECV_CKO[*n*]=ON

Either reboot the HP-UX server or issue the following commands to restart the driver and thereby activate the change:

/sbin/init.d/hpiether stop
/sbin/init.d/hpiether start

ixgbe Devices
For ixgbe 10 Gigabit Ethernet interfaces, modify the /etc/rc.config.d/hpixgbeconf file as follows:

HP_IXGBE_INTERFACE_NAME[*n*]=lan*n*
HP_IXGBE_VMTU[*n*]=32160
HP_IXGBE_TX_CHECKSUM_OFFLOAD[*n*]=ON
HP_IXGBE_RX_CHECKSUM_OFFLOAD[*n*]=ON

Either reboot the HP-UX server or issue the following commands to restart the driver and thereby activate the change:

/sbin/init.d/hpixgbe stop
/sbin/init.d/hpixgbe start

Warning: Before stopping and restarting network drivers such as igelan, iether, and ixgbe, verify the following:

- All affected Oracle Database instances have been shut down
- Oracle clusterware (RAC environments only) using the affected networks has been shut down
- Dependent NFS-mounted volumes have been unmounted
- Non-Oracle clusterware such as HP Service Guard has been shut down

If the above prerequisites cannot be met, you should activate changes made to the respective configuration files by rebooting the HP-UX server(s) during the next maintenance window instead of manually stopping and restarting the network drivers.

6.3. NFS Mount Options

The following mount options were used for the database volumes in our HP-UX 11i test environment:

rw,bg,hard,nointr,forcedirectio,proto=tcp,vers=3

As you can see, we're using the "forcedirectio" option, the TCP transport, and NFS version 3.

Direct I/O Semantics Using the "forcedirectio" Mount Option

Database applications usually give better performance if the database is allowed to manage its own cache. The use of direct I/O causes NFS requests to bypass NFS buffer cache on the server, allowing

the database to do just that. With direct I/O, an Oracle Database is able to manage its own caches using the Oracle SGA, thereby avoiding the overhead of double caching. Direct I/O can be invoked by mounting the NFS volume with the “forcedirectio” mount option.

Implementation in the Oracle instance is accomplished by setting the Oracle initialization parameter “filesystemio_options” to either “directio” or “setall.” A value of “setall” enables Oracle to use both direct I/O and async I/O. To implement either setting, you must shut down and restart the Oracle instance. NetApp best practices document kb7518 recommends a value of “setall” for this parameter: http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb7518.

Our database server had enough physical memory to provide for a sufficiently large SGA for our workload, making direct I/O the better choice. That being the case, our test database throughput improved by 482% with direct I/O enabled.

Local Locking Semantics Enabled Using the “llock” Mount Option

In some cases bypassing file system cache can be bad for database performance. Below are a few examples where that might be true:

- Limited system memory is being shared by multiple database instances.
- Limited system memory is being shared by both applications and database instances.
- The maximum Oracle SGA size is limited by the use of 32-bit Oracle RDBMS software.

In those cases, file system cache might actually improve database performance. Local locking using the “llock” mount option should be used if direct I/O is not used. “llock” modifies NFS behavior in a couple of ways. It causes file locks to be maintained locally by the HP-UX kernel instead of by the NFS server, resulting in fewer NFS operations and less network traffic between the HP-UX host and storage. It also improves file system cache performance by avoiding the aggressive cache purging that occurs without “llock.” For this reason, we strongly recommend using *either* direct I/O or local locking.

7. HP-UX 11i v2 OS-Specific Recommendations

This section contains tuning recommendations specific to the HP-UX 11i v2 operating system, including OS-specific kernel parameter settings and patches.

7.1. HP-UX 11i v2 Kernel Parameter Recommendations

HP recently introduced a number of kernel parameters for HP-UX 11i v2 systems that directly influence NFS client performance and behavior. Table 7 lists these recommended HP-UX 11i v2 kernel parameter settings. This list of tunables is in addition to the parameters listed in Table 4.

Table 7) HP-UX 11i v2 NFS kernel parameter recommendations

Kernel Parameter	Recommended Value	Details
nfs_async_read_avoidance_enabled	1	Controls whether NFS clients are able to issue NFS READ calls when all the biod daemons are busy.
nfs_fine_grain_fs_lock	2	Reduces the occurrence of semaphore contention in NFSv3 file operations.
nfs_new_lock_code	1	Controls the performance of asynchronous I/O against locked files on NFS file systems.
nfs_new_rnode_lock_code	1	Allows threads that are waiting for an rnode lock to sleep at an interruptible state.
nfs_wakeup_one	2	Eliminates "thundering herd" conditions in the rnode lock code and the main read/write code paths of the NFS client biod daemons.
nfs3_max_transfer_size	65536	Controls both the maximum size of the NFS data client may request and the size of NFS data the server may return in a single request.
nfs3_max_transfer_size_cots	65536	Controls the maximum request size for requests sent over connection-oriented transports, such as TCP.
nfs3_new_acache	1	Enables the use of a new hashed algorithm used to control the NFS client access cache.

For a complete listing of all HP-UX 11i v2 NFS and RPC kernel parameters, see the "Managing NFS and KRPC Kernel Configurations in HP-UX 11i v2" white paper available at http://docs.hp.com/en/14765/NFSTunablesWP11iv2_101608.pdf.

7.2. NFS Block I/O Daemons (BIOD Processes)

HP-UX 11i v2 uses biod daemons on the NFS client to improve NFS I/O performance by asynchronously scheduling read-ahead and write-behind requests. The number of biod processes is determined by the value of the NUM_NFSIOD parameter in the /etc/rc.config.d/nfsconf file. The default value in HP-UX 11i v2 is 16, which is sufficient for most database workloads. A change in the NUM_NFSIOD value can be activated by either stopping and restarting the NFS client service (that is, /sbin/init.d/nfs.client stop/start) or by rebooting the HP-UX system. In HP-UX 11i v3 the biod daemons are replaced by pools of kernel threads, and thus the NUM_NFSIOD parameter is not applicable to HP-UX 11i v3.

7.3. HP-UX 11i v2 Patch Recommendations

This section lists the HP-UX 11i v2 patches recommended for optimal NFS and Oracle Database 10*g* performance. Table 8 lists the recommended patch bundles for HP-UX 11i v2 systems. Table 9 lists the current individual NFS patches available for 11i v2 systems. This table also lists all dependent patches and patches for critical subsystems, like ARPA transport, STREAMS, Gigabit Ethernet, and so on, which NFS relies on for proper behavior and performance. Table 10 lists the current patch recommendations for optimal Oracle Database 10*g* performance.

If you install the patch bundles shown in Table 8, there will likely be some duplication between the patches included in those bundles and those listed in Table 9 and Table 10. However, as some administrators prefer to patch their systems using patch bundles and others prefer individual patches, we decided to list both options.

All of the patches and bundles listed in this section are available for download from HP's IT Resource Center (ITRC) site: http://itrc.hp.com. The patches and bundles listed in these tables were current as of the writing of this paper. Be sure to frequently check the ITRC for the latest information about patch bundles, quality packs, and the current NFS patches available for HP-UX 11i v2 systems.

Table 8) Recommended patch bundles for HP-UX 11i v2 systems

Patch Bundle Name	Patch Bundle Description
BUNDLE11i(B.11.23.0409.3)	Required Bundle for HP-UX 11i v2 (B.11.23), September 2004
HWEnable11i(B.11.23.0712.070)	Hardware Enablement Patches for HP-UX 11i v2, December 2007
FEATURE11i(B.11.23.0803.070a)	Feature Enablement Patches for HP-UX 11i v2, March 2008
QPK1123(B.11.23.0806.072)	Quality Pack Depot for 11i v2, June 2008

Table 9) Recommended individual NFS patches (and dependencies) for HP-UX 11i v2 systems

Patch Name	Patch Description
PHKL_31500	Sept04 base patch
PHCO_36744	LVM commands patch
PHCO_32523	quota(1) on an NFS client;uid enhancement
PHKL_32713	nfs client hang; cache FS umount

Patch Name	Patch Description
PHKL_33990	VxFS 3.5 : Quota metadata corruption
PHKL_35240	rwsleep locks
PHKL_35242	JFS3.5 bmap performance improvement
PHKL_35243	VxFS 3.5 VFS destacking support;DIO hang;bdf
PHKL_35244	vx_vget
PHKL_35479	Security Enhancement,vfs_teardown_stack
PHKL_35503	NFS Stale Handle Fix
PHKL_36745	LVM Cumulative Patch
PHKL_38287	VxFS 3.5 mount: Quota;logiosize
PHKL_38714	VxFS 3.5 cumulative patch
PHKL_38902	JFS3.5 DIO performance; extent rollback
PHNE_32057	CacheFS cumulative patch
PHNE_33100	AutoFS cumulative patch
PHNE_33102	RPC commands and daemons cumulative patch
PHNE_33103	NIS/NIS+ cumulative patch
PHNE_34756	Core NFS cumulative patch
PHNE_35117	Kernel RPC cumulative patch
PHNE_35120	libnsl cumulative patch
PHNE_37489	Lock Manager cumulative patch
PHNE_38252	NFS cumulative patch
PHKL_36577	PM-PSTAT section 2 manpage changes
PHNE_36236	IETHER 1000Base-SX/T B.11.23.[01-0706] patch
PHNE_36575	Cumulative STREAMS Patch
PHNE_37324	IGELAN 1000Base-SX/T B.11.23.[01-0712] patch

Patch Name	Patch Description
PHNE_37670	cumulative ARPA Transport patch

HP-UX 11i v2 Patches for Optimal Oracle Database 10*g* Performance

Table 10 lists the current patches (and any patch dependencies) recommended for optimal Oracle Database 10*g* performance. The Oracle installation may not complete if any of the listed HP-UX patches are missing.

Table 10) HP-UX 11i v2 patches required for Oracle Database 10*g* R2 installation

Patch Name	Patch Description
PHCO_32426	reboot(1M) cumulative patch
PHCO_36744	LVM commands patch
PHCO_37940	pthread library cumulative patch
PHCO_38273	libc cumulative patch
PHKL_31500	Sept04 base patch
PHKL_32631	GPE cell number change for vPars
PHKL_32637	vPars CPU init; Reuse interrupts; olrad msgs
PHKL_33808	PCI Error Handling feature
PHKL_33813	PCI Error Handling feature
PHKL_33820	mmap64(2) size limit;mutex unlock PA 64 bit
PHKL_34546	VM Pageout Data Corruption Fix
PHKL_34901	_OSC ACPI method
PHKL_34902	Firmware interfaces for PCIe
PHKL_34909	CPU OLA/D Soft interrupts, sw_callback
PHKL_35174	Partition ID, PCI EH, PCIe
PHKL_34941	Oracle Clusterware, Setboot for PCI Express

Patch Name	Patch Description
PHKL_34942	kepd pseudo driver, Setboot for PCI Express
PHKL_35737	PCI OLAR Cumulative IPF Patch, IOTreeProvidr
PHKL_36103	wsio.h header file cumulative patch
PHKL_36105	GIO cumulative patch
PHKL_36233	vPars,core cell,maxcpus,PCIe,TPM,MMIOPORT
PHKL_36319	mlockall(2), shmget(2) cumulative patch
PHKL_36320	VM physio fix
PHKL_36668	mlock header file fix
PHKL_36669	mlock performance fix
PHKL_36745	LVM Cumulative Patch
PHKL_36750	PCI OLA/R, data corruption,PCI EH,PCI-X2.0
PHKL_37106	WSIO IO subsystem cumulative patch
PHKL_37121	ksleep kwakeup performance cumulative patch
PHKL_37122	pthread_condvar_prio_boost patch
PHKL_37619	VM physio vas fix
PHKL_37653	mmap(2),mlock[all](2),madvise(2),msync(2)
PHKL_37654	VM pstat
PHKL_37655	pstat(2) fix
PHKL_37656	pstat(2) pstat header fix
PHKL_37657	pstat(2) PSTAT_PROC_VM_LITE
PHKL_37777	mmap(2) Data Corruption Fix
PHKL_38151	mmap(2),mmap64(2),madvise(2),mlockall(2)
PHKL_38364	MSI,PCI EH,PCIe and cumulative defect fixes
PHSS_34445	milli cumulative patch

Patch Name	Patch Description
PHSS_36347	assembler patch
PHSS_37552	Aries cumulative patch
PHSS_38134	linker + fdp cumulative patch
PHSS_38136	Math Library Cumulative Patch
PHSS_38138	Integrity Unwind Library
PHSS_38140	aC++ Runtime (IA: A.06.20, PA: A.03.85)

8. HP-UX 11i v3 OS-Specific Recommendations

This section contains tuning recommendations specific to the HP-UX 11i v3 operating system, including OS-specific kernel parameter settings and patches.

8.1. HP-UX 11i v3 Kernel Parameter Recommendations

HP-UX 11i v3 shipped with a number of new kernel-tunable parameters that can directly influence NFS client performance and behavior. Table 11 lists the recommended HP-UX 11i v3 kernel parameter settings for optimal NFS client behavior and performance. This list of kernel tunables is in addition to those parameters listed earlier in Table 4.

Table 11) HP-UX 11i v3 NFS kernel parameter recommendations

Kernel Parameter	Recommended Value	Details
nfs3_bsize	1048576	Controls both the maximum size of the NFS data client may request and the size of NFS data the server may return in a single request.
nfs3_max_threads	16	Control the number of kernel threads that perform asynchronous I/O for the NFS version 3 client.
nfs3_nra	8	Control the number of read-ahead operations queued by the NFS version 3 client when sequentially accessing a file.
intr_strobe_ics_pct	100	System limit for % of time a processor is allowed to spend in interrupt context.

For a complete listing of all HP-UX 11i v3 NFS and RPC kernel parameters, see the “Managing NFS and KRPC Kernel Configurations in HP-UX 11i v3” white paper available at http://docs.hp.com/en/13579/NFSTunablesWP_11iv3.pdf.

8.2. HP-UX 11i v3 Patch Recommendations

This section lists the HP-UX 11i v3 patches recommended for optimal NFS and Oracle Database 10*g* performance. Table 12 lists the recommended patch bundles for HP-UX 11i v3 systems.

If you install the patch bundles shown in Table 12, there will likely be some duplication between the patches included in those bundles and those listed in Table 13. However, as some system

administrators prefer to patch their systems using patch bundles and others prefer individual patches, we decided to list both options.

The bundles listed in Table 12 are available for download from HP's IT Resource Center (ITRC) site: http://itrc.hp.com. The bundles listed were current as of the writing of this paper. Be sure to frequently check the ITRC for the latest information about patch bundles and quality packs for 11i v3 systems.

Table 12) Recommended patch bundles for HP-UX 11i v3 systems

Patch Bundle Name	Patch Bundle Description
HWEnable11i(B.11.31.0809.326a)	Hardware Enablement Patches for HP-UX 11i v3, September 2008
FEATURE11i(B.11.31.0809.326a)	Feature Enablement Patches for HP-UX 11i v3, September 2008
QPK1131(B.11.31.0809.326)	Quality Pack Depot for 11i v3, September 2008

HP-UX 11i v3 ONCplus ISU Updates

On HP-UX versions prior to 11i v3, Open Network Computing (ONC) products were delivered as a core product (called NFS). Core products cannot be released independent of the release of the operating system. Any changes (for example, defect fixes) to core products were only available to customers by installing HP-UX patches. Also, there are strict limitations about providing new features to core products using HP-UX patches. For these reasons, HP decided to release NFS on HP-UX 11i v3 as an Independent Software Unit (ISU).

The ONCplus ISU provides system administrators the most flexibility in fixing problems along with the opportunity to take advantage of new functionality. The ONCplus ISU bundles are available for download from the Software Depot Web site:

http://software.hp.com/portal/swdepot/displayProductInfo.do?productNumber=ONCplus.

The current bundle version is B.11.31.05. New ONCplus ISU bundles are released approximately every three months. These bundles contain NFS-specific defect fixes, new functionality, performance enhancements, stability improvements, and so on. System administrators are strongly encouraged to review the contents of the latest ONCplus ISU and consider updating to the most recent ONCplus ISU version.

HP-UX 11i v3 Patches for Optimal Oracle Database 10*g* Performance

Table 13 lists the current patches (and any patch dependencies) recommended for optimal Oracle Database 10*g* performance. The Oracle installation may not complete if any of the listed HP-UX patches are missing.

Table 13) HP-UX 11i v3 patches required for Oracle Database 10*g* R2 installation

Patch Name	Patch Description
PHCO_38050	pthread library cumulative patch
PHKL_36261	Integrity Virtual Machine Capabilities Interface
PHKL_37460	system call infrastructure cumulative patch
PHKL_37462	diagnostic memory driver cumulative patch
PHKL_37463	kernel tunables infrastructure cumulative patch
PHKL_38034	MCA/Fault when executing in kernel mode on vPars/IPF
PHKL_38035	svc cumulative patch
PHKL_38038	vm cumulative patch
PHKL_38040	Kernel debugger cumulative patch
PHKL_38080	vPars PSM component for vPars A.05.04 release
PHKL_38091	cdfs cumulative patch
PHKL_38398	sba cumulative patch
PHKL_38414	dump cumulative patch
PHSS_37553	Aries cumulative patch
PHSS_38135	linker + fdp cumulative patch

9. Conclusions

NetApp storage systems provide superior data protection, availability, administration, and backup and recovery through NetApp tools and features that include the following:

- Fast reliable backups using NetApp Snapshot™ technology
- Disk redundancy through RAID-DP
- Instant clones, requiring very little additional storage space, using NetApp FlexClone® technology
- Deduplication
- Disaster recovery using NetApp SyncMirror® technology

Integrated with HP-UX servers using NFS, you can achieve cost-effective, enterprise-class support for business applications using Oracle Databases. Together they deliver high performance, scalability, and a rich set of features for processing, managing, and protecting data. This paper has focused on performance, providing recommendations and guidelines for tuning the HP-UX server, NetApp storage system, Oracle Database instance, and storage network in order to achieve optimum database performance with NFS. The potential impact of following these tuning guidelines is clearly demonstrated using actual test results from NetApp performance labs.

As with any computing environment, performance tuning is an iterative process, and the results are often influenced by workload characteristics. This document provides guidelines and methodology for that process.

10. Acknowledgements

Special thanks to the following people for their contributions and support:

NetApp, Inc.

Lee Dorrier, Keith Griffin, Ricky Stout