



HP-UX Kernel Configuration for Oracle data bases and data warehouse

Last Update: March 2010
Oracle release: Oracle 11gR2/R1
Oracle 10gR1/R2
Oracle 9i R1/9i R2
Operating System: HP-UX 11iv3 (also called 11.31)
HP-UX 11i v2 (also called 11.23 and 11.23PI)
HP-UX 11i v1 (also called 11.11PA-Risc)
Author: Bernd Menth/HP Oracle CTC

Content:

- 1. Oracle release information**
- 2. HP-UX release information**
- 3. additional SW/patch information and links**
- 4. HP-UX Kernel Parameter**
- 5. Helpful Tools**
- 6. Appendix 1: db_block_size**
- 7. Appendix 2: Raw Devices + Asynchronous IO**
- 8. Appendix 3: Parallel Query Option**
- 9. Appendix 4: shmmax**
- 10. Appendix 5: Swapping/Paging**
- 11. Appendix 6: Privileges for SCHED_NOAGE and MLOCK**
- 12. Appendix 7: cell local memory vs. interleaved memory**





HP-UX Kernel Configuration for Oracle data bases and data warehouse

1. Oracle release information

1.1 Oracle 11gR2 (11.20.1) recommendations:

Memory

At least 4 GB of RAM

To determine the RAM size, enter the following command:

```
# /usr/contrib/bin/machinfo | grep -i Memory
```

Swap Space

Between 4 GB and 8 GB 2 times the size of RAM

Between 8 GB and 32 GB 1.5 times the size of RAM

More than 32 GB use 32 GB if you are short in memory

To determine the size of the configured swap space, enter the following command:

```
# /usr/sbin/swapinfo -a
```

tmp space

Temporary Disk Space: 1 GB of space in the /tmp

If there is not 1 GB available space in the /tmp directory,

then you can set the TMPDIR or TMP (used by Oracle) environment

variable to point to a directory that has at least 400 MB of available space.

To determine the amount of disk space available in the /tmp directory, enter the following command:

```
# bdf /tmp
```

disk space requirements for Oracle binaries

Enterprise Edition 6.89 GB

Custom (Enterprise Edition + custom install options) 6.89 GB

Optional : additional 4.5 GB for the grid/ASM infrastructure home

disk space requirements for Oracle data files (optional)

Enterprise Edition 1.69 GB

Custom (Enterprise Edition + custom install options) 1.81 GB

Additional disk space, either on a file system or on an Automatic Storage Management disk group is required for the fast recovery area if you choose to configure automated backups.





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Operating System Requirements

HP-UX 11i V3 patch Bundle Sep/ 2008 (B.11.31.0809.326a) or higher

To determine the distribution and version of HP-UX installed, enter the following command:

```
# uname -a
```

Verify that the system meets the minimum patch bundle requirements using the following command:

```
# /usr/sbin/swlist -l bundle |grep QPK
```

The QPK (Quality Pack) bundles have version numbers of the form

B.11.31.0809.326a (for the September 2008 release), B.11.31.0903.334a (for the March 2009 release), and so on.

Compiler Requirements

The following are the compiler requirements for HP-UX systems Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, and Oracle XML Developer's Kit (XDK) with Oracle Database 11g Release 2 (11.2):

A.06.20 (HP C/aC++ Swlist Bundle - C.11.31.04) - Sep 2008

More information regarding the different link options during the installation.

<http://h21007.www2.hp.com/portal/site/dspp/menuitem.863c3e4cbcdc3f3515b49c108973a801/?ciid=e408e467f5f02110e467f5f02110275d6e10RCRD>

Patch Requirement

The following, or later, patches and bundles are required for Oracle Database 11g Release 2 (11.2) for HP-UX systems.

http://download.oracle.com/docs/cd/E11882_01/install.112/e10813.pdf

All installations for HP-UX 11i V3 (11.31) requires the following patches (or newer versions: the current patch successor status March 2010)

■ PHCO_40381 11.31 Disk Owner Patch (as this a new patch, make sure this patch is installed)

■ PHKL_38038 VM patch - hot patching/Core file creation directory

-> replaced by [PHKL_40942](#)

■ PHKL_38938 11.31 SCSI cumulative I/O patch

-> replaced by [PHKL_39646](#)

■ PHKL_39351 Scheduler patch : post wait hang

-> replaced by [PHKL_40207](#)

■ PHSS_36354 11.31 assembler patch

-> replaced by [PHSS_40546](#)

■ PHSS_37042 11.31 hppac (packed decimal)

■ PHSS_37959 Libcl patch for alternate stack issue fix (QXCR1000818011)

-> replaced by [PHSS_40804](#)

■ PHSS_39094 11.31 linker + fdp cumulative patch

-> replaced by [PHSS_40538](#)

■ PHSS_39100 11.31 Math Library Cumulative Patch

-> replaced by [PHSS_40540](#)





HP-UX Kernel Configuration for Oracle data bases and data warehouse

■ PHSS_39102 11.31 Integrity Unwind Library

-> replaced by [PHSS_40542](#)

■ PHSS_38141 11.31 aC++ Runtime

-> replaced by [PHSS_40544](#)

Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)

Patch for HP-UX 11i V3 (11.31) on HP-UX Itanium:

PHSS_39824 - 11.31 HP C/aC++ Compiler (A.06.23) patch

➔ Replaced by [PHSS_40631](#)

VERITAS File System

PHKL_39773: 11.31 VRTS 5.0 GARP6 VRTSvxfs Kernel Patch

➔ Replaced by [PHKL_40650](#)

Note: The VERITAS file system is optional. This patch is required only if you want to use a VERITAS File System 5.0.

To ensure that the system meets these requirements:

- To determine whether a patch is installed, enter a command similar to the following:

```
# /usr/sbin/swlist -l patch | grep PHCO_40381
```

Alternatively, to list all installed patches, enter the following command:

```
# /usr/sbin/swlist -l patch | more
```

- If a required patch is not installed, then download it from the following Web site and install it:

<http://itresourcecenter.hp.com>

If the Web site displays a recent version of the patch, then download and install that version.

Installation Fixup Scripts

During installation, for certain prerequisite check failures, you can click Fix & Check Again to generate a fixup script (runfixup.sh). You can run this script as a root user to complete the required preinstallation steps.

The fixup script does the following:

■ Checks and sets kernel parameters to values required for successful installation, including:

– Shared memory parameters

– Open file descriptor and UDP send/receive parameters

Oracle recommends that you do not modify the contents of the generated fixup script but it also not guarantee that all necessary changes are done.





HP-UX Kernel Configuration for Oracle data bases and data warehouse

1.2 Oracle 11gR1 (11.1.0.7) recommendations:

- Check process during installation:

Temp 415 MB

Swap 250 MB

/opt/oracle 5.5 GB

Checking operating system requirements ...

Expected result: One of 11.23,11.31

Actual Result: 11.31

Check complete. The overall result of this check is: Passed

=====

Checking recommended operating system patches -> could be already replaced by a new version

Checking for PHKL_35936;

Checking for PHKL_35900;

Checking for PHKL_36248;

Checking for PHKL_36249;

=====

Checking kernel parameters

Checking for ksi_alloc_max=32768;

Checking for max_thread_proc=1024;

Checking for maxdsiz=1073741824;

Checking for maxdsiz_64bit=2147483648;

Checking for maxssiz=134217728;

Checking for maxssiz_64bit=1073741824;

Checking for maxuprc=3687;

Checking for msgmni=4096;

Checking for msgtql=4096;

Checking for nsize=35840;

Checking for nflocks=4096;

Checking for ninode=34816;

Checking for nkthread=7184;

Checking for nproc=4096;

Checking for semmni=4096;

Checking for semmns=8192;

Checking for semmnu=4092;

Checking for semvmx=32767;

Checking for shmmax=1073741824;

Checking for shmmni=4096;

Checking for shmseg=512;

Checking for vps_ceiling=64;

=====

Checking physical memory requirements ...

Expected result: 1024MB

Checking available swap space requirements ...

Validating ORACLE_BASE location (if set) ...





HP-UX Kernel Configuration for Oracle data bases and data warehouse

1.3 Oracle 10g recommendations:

- Minimum RAM = 1024 MB
- HP-UX 11i
- Swap Space: Minimum
 - o Between 1GB and 2GB RAM use 1.5 times the size of RAM for swap
 - o Between 2GB and 8 GB RAM use equal to the size of RAM
 - o More than 8 GB RAM 0.75 to equal size of RAMplease check Appendix 5 for large systems
- available Disk Space: up to 4 GB (sw only up to 3 GB)
this includes 1 GB from companion CD
- Temporary Disk Space: 400 MB of space in the /tmp directory
If there is not 400MB available space in the /tmp directory,
then you can set the TMPDIR or TMP (used by Oracle) environment
variable to point to a directory that has at least 400 MB of available space.
This directory must be writable by the user account performing the installation.)
- Disk space for DB files: 1 .2 GB
- Create OS UNIX groups for Oracle system:
 - o OSDBA group : Required
 - o OSOPER group: Optional
 - o ORAINVENTORY group: Optional
 - o APACHE group: Recommended
- Please check appendix 6 for SCHED_NOAGE usage and MLOCK

1.4 Oracle 9i recommendations:

- Minimum RAM = 256 MB (HP-UX 11.0/11.i v1)
- Minimum RAM = 512 MB (HP-UX 11.i v2 and v1.6)
- HP-UX 11.0 or HP-UX 11i v1 or HP-UX 110 v2.0
- Swap Space: Minimum 2 x RAM or min 400MB, whichever is greater,
please check Appendix 5 for large systems
- available Disk Space: 3 GB (HP-UX 11.0/11.i v1)
- available Disk Space: 3.5 GB (HP-UX 11.i v2 and v1.6)
- an additional 1 GB for a seed database
- Temporary Disk Space: up to 2GB of space in the /tmp directory
(If there is not 2 GB of available space in the /tmp directory,
then you can set the TMPDIR or TMP (used by Oracle) environment
variable to point to a directory that has at least 400 MB of available space.
This directory must be writable by the user account performing the installation.)
- Create OS UNIX groups for Oracle system:
 - o OSDBA group : Required
 - o OSOPER group: Optional
 - o ORAINVENTORY group: Optional
 - o APACHE group: Recommended
- Please check appendix 6 for SCHED_NOAGE usage and MLOCK





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Oracle® Database on Unix® Operating Systems Installation and Configuration
Requirements Quick Reference Doc ID: Note:169706.1 on <http://metalink.oracle.com>

Oracle 9i, 10g and 11gR requires some patches on the specific Operating System
These patches might be superseded.

HP-UX 11iv3 (11.31)

- **The following kernel parameters are obsolete on HP-UX 11iV3.**
 - maxswapchunks
 - msgmap
 - ncallout
 - semmap
 - vx_ncsize
 - dbc_max_pct replaced by filecache_max (50% default)
 - dbc_min_pct replaced by filecache_min (5% default)
 - the file caching with 11i v3 is different from v2
<http://docs.hp.com/en/oshpux11iv3.html>
- You need to install the following patches (or newer version) on top of HP-UX 11.31 for Oracle DB (status Feb 2009)
 - PHCO_40381 11.31 Disk Owner Patch (as this a new patch, make sure this patch is installed) only necessary for Oracle 11gR2
 - PHKL_38651 (replaces PHKL_35900) vm cumulative patch
 - dependencies
 - PHCO_38050 PHKL_36261 PHKL_37462 PHKL_37463 PHKL_38034 PHKL_38035
 - PHKL_38080 PHKL_38091 PHKL_38398 PHKL_38414 PHSS_37553 PHSS_37955
 - PHKL_38651 (replaces PHKL_35936) vm cumulative patch
 - PHKL_38604 (replaces PHKL_36248) esctl cumulative patch
 - PHKL_38605 (replaces PHKL_36249) esdisk cumulative patch
- We recommend to install the latest patch bundle,
f.e. the September 2007 version 11.31.0709
- For Oracle releases 11gR1, 10g and 9i, you need to use the additional flag –
ignoreSysPreReqs for the runInstaller, otherwise the installation
will not continue
 - runInstaller -ignoreSysPreReqs





HP-UX Kernel Configuration for Oracle data bases and data warehouse

HP-UX 11i v2 (11.23) (Itanium2 / PA-RISC):

- HP-UX B.11.23.0412 or later
- Quality Pack Bundle:
 - Latest patch bundle: Quality Pack Patches for HP-UX 11i v2
- HP-UX 11.23 Patches:
 - PHSS_31849: linker + fdp cumulative patch
 - PHSS_31852: aC++ Runtime (PA A.03.61)
 - PHSS_32511: HP aC++ Compiler (A.03.63)
 - PHSS_32512: ANSI C compiler B.11.11.12 cumulative patch
 - PHSS_32513: +O4/PBO Compiler B.11.11.12 cumulative match
 - PHCO_35524 (LVM patch date April 2007 necessary for 10203 with ASM)
- Serviceguard and OS Patches (optional, only if you want to use Serviceguard):
 - PHSS_32740: Serviceguard A.11.16.00 (replaces PHSS_31076)
 - PHSS_32659: Serviceguard Extension for RAC 11.16.00 (optional, only if you want to use SGeRAC) (replaces PHSS_31080)
 - PHSS_32741: COM B.03.00.00/COM B.03.00.01 (replaces PHSS_31078)
 - PHCO_31244: VxVM 3.5-IA.004 Command Patch 02 (only required for clusters using VxVM; required to support 8 node clusters with CVM)

HP-UX 11.11 (PA-RISC):

- Quality Pack Bundle:

HP-UX 11i Quality Pack (GOLDQPK11i)
- HP-UX 11.11 Patches:
 - PHNE_31097: ONC/NFS general release/performance patch
 - PHSS_31221: HP aC++ -AA runtime libraries (aCC A.03.60)
 - PHSS_32508: HP aC++ Compiler (A.03.63)
 - PHSS_32509: ANSI C compiler B.11.11.12 cumulative patch
 - PHSS_32510: +O4/PBO Compiler B.11.11.12 cumulative patch
- Patches for JDK on HP-UX 11.11:
 - PHSS_30970: ld(1) and linker tools cumulative patch
 - **Note:** Refer to the following Web site for information about additional patches that may be required by JDK 1.4.2:
<http://www.hp.com/products1/unix/java/patches/index.html>





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Creating Required Symbolic Links

To enable you to successfully relink Oracle products after installing this software, enter the following commands to create required X library symbolic links in the /usr/lib directory:

Note: This task is required only if the Motif 2.1 Development Environment package (X11MotifDevKit.MOTIF21-PRG) is not installed.

```
# cd /usr/lib
# ln -s libX11.3 libX11.sl
# ln -s libXIE.2 libXIE.sl
# ln -s libXext.3 libXext.sl
# ln -s libXhp11.3 libXhp11.sl
# ln -s libXi.3 libXi.sl
# ln -s libXm.4 libXm.sl
# ln -s libXp.2 libXp.sl
# ln -s libXt.3 libXt.sl
# ln -s libXtst.2 libXtst.sl
```

To ensure that the system meets these requirements, follow these steps:

- HP provides patch bundles at <http://www.itrc.hp.com/service/sum/home.do>
- To determine whether the HP-UX 11i Quality Pack is installed:
`# /usr/sbin/swlist -l bundle | grep GOLD`
- Individual patches can be downloaded from <http://itresourcecenter.hp.com/>
- To determine which operating system patches are installed, enter the following command:
`# /usr/sbin/swlist -l patch`
- To determine if a specific operating system patch has been installed, enter the following command:
`# /usr/sbin/swlist -l patch <patch_number>`
- To determine which operating system bundles are installed, enter the following command:
`# /usr/sbin/swlist -l bundle`

Required Executables

The make, ar, ld, and nm executables must be present in the /usr/ccs/bin directory

Additional SW:

To support Oracle Real Application Clusters 9i, the following operating system package and patch are required:

MC/ServiceGuard

MC/ServiceGuard extension for RAC

With Oracle 10g/11g RAC, you do have the choice between different cluster ware.





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Note: These are minimum kernel requirements for Oracle database. If you have previously tuned your kernel parameters to levels equal to or higher than these values, continue to use the higher values, except for file system cache parameters (tunable filecache_max , tunable filecache_min, dbc_max_pct, dbc_min_pct). The default parameter for files system cache is for an Oracle database most of the time to high (tunable filecache_max, dbc_max_pct = 50, means 50% of your memory is given to file system cache), it is better to give Oracle the memory to cache data in the SGA instead to have a large files system cache. A system restart is necessary for kernel changes to take effect.



HP-UX Kernel Configuration for Oracle data bases and data warehouse

Kernel Parameter Setting Purpose

HP-UX 11.x Kernel Parameter	Description, Default/Recommended Value
Bufpages	Number of static Buffer Pages of 4kB, enables dynamic buffer cache (see dbc_max/min_pct). Not used anymore >= HP-UX 11.23 default= 0 recommended: 0
dbc_max_pct (for 11iV3 use filecache_max)	Max dynamic Buffer Cache size in percent of system memory, when bufpages and nbuf are both set to 0. Do not set this value to high, it will use this percentage of the memory to file system cache and give it not to Oracle or any other application!!!! default=50. recommended: between 3 and 10 % of the memory It's more efficient to use system memory for Oracle SGA instead of file system buffer cache hpux 11i v1,v2
dbc_min_pct (for 11iV3 use filecache_min)	Min dynamic Buffer Cache size in percent of system memory (see dbc_max_pct), default=5. recommended: between 2 and 5 hpux 11i v1,v2
executable_stack	This tunable parameter controls whether program stacks are executable by default. It allows systems to be configured to have extra protection from stack buffer overflow attacks without sacrificing system performance. (HP-UX 11iV2 feature) Default=0 Recommended: >= Oracle 10g R2 = 0



HP-UX Kernel Configuration for Oracle data bases and data warehouse

filecache_max	<p>Max dynamic file system cache size in percent of system memory. Do not set this value to high, it will use this percentage of the memory to file system cache and give it not to Oracle or any other application!!!! default=50. recommended: between 3 and 10 % of the memory It's more efficient to use system memory for Oracle SGA instead of file system buffer cache >= hp-ux 11iv3 feature</p>
filecache_min	<p>Min dynamic file system cache size in percent of system memory (see dbc_max_pct), default=5. recommended: between 2 and 5 >= hp-ux 11iv3 feature</p>
fs_async	<p>see appendix 2 and appendix 6</p>
Ksi_alloc_max	<p>Defines the system wide limit of queued signal that can be allocated. Default: (NPROC * 8) ->Oracle 11gR2=32768</p>
lcpu_attr	<p>Enable/disable hyper threading online default=0 (off) recommend: dependent on the workload the use of hyper threading will increase your throughput. Single threaded process will probably not benefit from the use of hyper threading >= hp-ux 11iv3 feature enable hyperthreading</p> <ul style="list-style-type: none">• in firmware: HP-UX#setboot -m on• EFI command: EFI#cpuconfig threads on
max_async_ports	<p>Maximum number of asyncdsk ports that can be open at one time. If an error happens with this, it will not shown in the Oracle alert.log file or in trace files. default: 50 recommended: max. no of shadow processes + no of parallel query slaves (could go up to nproc)</p>





HP-UX Kernel Configuration for Oracle data bases and data warehouse

max_thread_proc	Max no. of threads allowed in each process default: 256 recommended: <= Oracle 8.1.6 = 64 Oracle 8.1.7 - Oracle10gR1= 256 >= Oracle 10gR2 = 1024
Maxfiles	Soft File Limit per Process, soft limit for number of files a process is allowed to have open simultaneously. Processes can increase their soft limit until they reach the hard limit maxfiles_lim. default=2048 Oracle11gR2= 1024
Maxfiles_lim	Hard File Limit per Process, default=4096 recommended=4096 (default). ->Oracle11gR2=63488
Maxdsiz	Refers to the maximum data segment size for 32-bit systems. Setting this value too low may cause the processes to run out of memory. Max Data Segment Size (Bytes) default: 1073741824 recommended: <= Oracle 8.1.7 = 256 MB >= Oracle 9i = 1073741824 bytes (1 GB) (0x40000000)
Maxdsiz_64bit	Refers to the maximum data segment size for 64-bit systems. Setting this value too low may cause the processes to run out of memory. Max Data Segment Size (Bytes) default: 1073741824 recommended: <= Oracle 8.1.7 = 1073741824 bytes (1 GB) (0x40000000) >= Oracle 9i >= 2147483648 bytes (2 GB) (0x80000000)
MAXSSIZ	Defines the maximum stack segment size in bytes for 32-bit systems. Recommended: >= Oracle 9i = 134217728 bytes (128 MB) (0x8000000)



HP-UX Kernel Configuration for Oracle data bases and data warehouse

MAXSSIZ_64BIT	Defines the maximum stack segment size in bytes for 64-bit systems. Recommended: >= Oracle 9i = 1073741824 bytes (1 GB) (0x40000000)
Maxswapchunks or swchunk	maximum number of swap chunks where SWCHUNK is the swap chunk size (1 KB blocks). (related to parameter swchunk) used with HP-UX 11, 11i, 11iv1.6 not used >= HP-UX 11iv2 Default: 2048 recommended: <= Oracle 8i = 4096 >= Oracle 9i = 16384
maxtsize	Max Text Segment size (Bytes) Recommended: 134217728 bytes (128 MB) (0x8000000) not used >= Oracle 10g
max_fcp_reqs	Maximum Number of Concurrent Fiber Channel Requests Per Adapter recommended: 512 >= HP-UX 11i v1.6 not used anymore
Maxuprc	Max Number of simultaneous user processes (per user-id!). default=75 Because all database processes often run with the Oracle user-id, it's recommended: <= Oracle 8.1.7 = maxusers * 5 = Oracle 9i = ((NPROC*9)/10) = Oracle 10g = (((NPROC*9)/10)+1) >=Oracle11gR2=3686
Maxusers	Value of MAXUSERS macro, limits the Ssstem resource allocation (not the actual number of users). Influences nproc, ninode, nfile. default=32. Not used anymore >= HP-UX 11.23 recommended: at least set to number of concurrent Oracle DB users + 64 not used >= Oracle 10g



HP-UX Kernel Configuration for Oracle data bases and data warehouse

Maxvgs	<p>Max. number of volume groups: Default: 10 Recommended: increase to the number of volume groups you would like to have on the system (maximum 256)</p>
MSGMAP	<p>Defines the maximum number of message map entries. Recommended: Oracle 9i = (NPROC + 2) Oracle 10g=(MSGTQL+2) Oracle 11gR2=4096</p> <p>Obsolete for >= HP-UX 11iv3 (11.31)</p>
MSGMNI	<p>Defines the number of message queue identifiers. Recommended: >= Oracle 9i = (NPROC) (at least 4096)</p>
MSGSEG	<p>Defines the number of segments available for messages. Recommended: >= Oracle 9i = (NPROC * 4) (at least 32767)</p>
MSGTQL	<p>Defines the number of message headers. Recommended: >= Oracle 9i = (NPROC) (at least 4096)</p>
NCALLOUT	<p>Defines the maximum number of pending timeouts. Not used anymore >= HP-UX 11.23 Recommended: >= Oracle 9i = (NPROC + 16) not used >= Oracle 10g</p>
NCSIZE	<p>Defines the Directory Name Lookup Cache (DNLC) space needed for inodes. Recommended: >= Oracle 9i = ((NINODE + VX_NCSIZE) VX_NCSIZE is by default 1024. >=Oracle11gR2=35840</p> <p>For >= HP-UX 11iv2 please use = (NINODE + 1024)</p> <p>If you have a large number of ninodes, than normally Oracle does not benefit from a large ncsiz.</p> <p>Ncsiz between 4096 and 8192 should be normally enough. Having ncsiz to large might decrease the performance because of serializing of processes.</p>





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Nfile	<p>Max Number of simultaneously Open files system-wide at any given time. Total number of slots in the file descriptor table, default=$16 * (nproc + 16 + maxusers) / 10 + 32 + 2 * (npty + nstrpty)$ recommended: <= Oracle 8.1.7 = to use default. >= Oracle 9i = $(15 * NPROC + 2048)$ >= Oracle 11gR2 = 63488</p> <p>For Oracle installations with a high number of data files this might be not enough, than use as an upper limit for nfile the formula <= $((\text{number of oracle processes}) * (\text{number of Oracle data files}) + 2048)$ (or a value between the standard formula and the formula for the high number of data files)</p>
Nflocks	<p>Max Number of File Locks available system-wide, (check with glance if you are getting close to the value you have set) default=200 or $200 + 10 * (\text{num_clients})$ recommended: <= Oracle 8.1.7 = at least 200 + sum off all db_files for all instances of Oracle >= Oracle 9i = (nproc) (at least 4096)</p>
NINODE	<p>Defines the maximum number of open inodes. Recommended: >= Oracle 9i = $(8 * NPROC + 2048)$ >= Oracle 11gR2 = 34816</p> <p>If you are limited in memory, having not a high number of files on the server and using hfs only for /stand than you could probably go with the number of nproc for ninode. However, if ninode is set low, be sure to tune ncsiz or vx_ncsiz on the system.</p>
NKTHREAD	<p>Defines the maximum number of kernel threads supported by the system. Recommended: >= Oracle 9i = $((NPROC * 7) / 4) + 16)$ >= Oracle 11g = 7184</p>





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Nproc	Max Number of Processes that can exist simultaneously in the system, default=(20+8*MAXUSERS), influences ninode, nfile. recommended: <= Oracle 8.1.7 = to use default >= Oracle 9i = (20+8*MAXUSERS) at least 4096
Npty	Number of ptys (pseudo ttys), default=60. recommended: if you are using rlogin/telnet connections from clients to server then increase npty up to the number of client users, otherwise use default. Note: ptys are not used when using Oracle SQL*Net
num_tachyon_adapters	HP-UX 11.0 only Number of Tachyon-based Fiber Channel Adapters in the System Recommended: minimum 5, set to real number of Fiber Channel Adapters HP-UX 11i – no support for tachyon adapters, the new is tachlite, which do not use this kernel parameter.
o_sync_is_o_dsync	Enable/Disable translation of O_SYNC to O_DSYNC in open()/fcntl() calls Not used >= HP-Ux 11iv2 default=0. recommended: < Oracle 7.3 = o_sync_is_o_dsync=1 or use Oracle Patch for Bug #310042. >= Oracle7.3 = use default for o_sync_is_o_dsync=0 (patch #310042 not needed).



HP-UX Kernel Configuration for Oracle data bases and data warehouse

scsi_max_qdepth	<p>Device Queue Depth is set to 8 by default on HP-UX. This value is per LUN. Queue depth should be increased if only a few LUNs are being accessed Default: 8 Recommended: Having a low number of LUNs like 8 with a high number of readers, increasing the scsi max queue depth f.e. to 32 or 64 will probably increase the throughput .</p> <p>With HP-UX 11iv3 you need to use scsimgr to set the max_q_depth per device, f.e. scsimgr get_attr -D /dev/rdisk/disk16 -> provide all the parameters for disk16 scsimgr get_attr -a max_q_depth -D /dev/rdisk/disk16 -> shows the max_q_depth for disk16 scsimgr set_attr -a max_q_depth=32 -D /dev/rdisk/disk16 -> set the max_q_depth to 32 for disk16 http://docs.hp.com/en/scsimgr/scsimgr_whp_AR0803.pdf?jumpid=reg_R1002_USEN</p>
SEMMAP	<p>Defines the maximum number of semaphore map entries. Recommended: >= Oracle 9i = (SEMMNI + 2)</p> <p>In HP-UX 11i v2, the memory allocation of semaphore sets previously controlled by semmap is done dynamically by the kernel. It is no longer necessary to tune the semmap kernel parameter</p>



HP-UX Kernel Configuration for Oracle data bases and data warehouse

Semmni	<p>Number of Semaphore Identifiers, specifies the maximum number of sets of semaphores that can exist simultaneously on the system, default=64. recommended: <= Oracle8i = use default or at least 10 per Oracle database >= Oracle9i = (nproc) >=Oracle11gR2= 4096 or (total semaphores sets) to semmns divided by semmsl, rounded up to the nearest multiple of 1024.</p>
Semmns	<p>Max Number of Semaphores, defines the system-wide maximum number of individual semaphores that can be allocated for users, default=128. recommended: <= Oracle 8.1.7 = 256 or at least 1 per Oracle process >= Oracle 9i = (semmni * 2) >= Oracle 11g = 8192 or 2 * sum (process parameters of all database instances on the system) + overhead for background processes + system and other application requirements</p>
SEMMSL	<p>Semmaphor for each set >= Oracle 11g=256</p>
SEMMNU	<p>Defines the number of semaphore undo structures. Recommended: >= Oracle 9i = (NPROC - 4) >=Oracle11g=4092</p>
SEMVMX	<p>Maximum value of a semaphore. Recommended: >= Oracle 9i = 32767</p>





HP-UX Kernel Configuration for Oracle data bases and data warehouse

shmmax (see appendix 4)	<p>maximum allowable size of one shared memory segment. The SHMMAX setting should be large enough to hold the entire SGA in one shared memory segment. A low setting can cause creation of multiple shared memory segments which may lead to performance degradation, but normally it is ok to use a small number of shared memory segments.</p> <p>default=0x04000000 (64MB) Oracle 7 0x40000000 (1GB = 1073741824 Bytes) Oracle 8, 8i, 9i, 10g,11g</p> <p>recommended: db_block_size * db_block_buffers + shared_pool_size + log_buffer + shared memory for the application or hold the entire SGA in one shared memory segment ! f.e. 0x100000000 = 4GB (for 64bit use up to size of free physical memory, but Oracle can also use multiple segments (please check appendix 4))</p>
Shmmni	<p>maximum number of shared memory segments in the entire system. default=200. recommended: <= Oracle 8.1.7 = use default. Oracle 9i, 10g = 512 >=Oracle11gR2=4096 (grid infrastructure requirement)</p>
Shmseg	<p>maximum number of shared memory segments one process can attach. default=12. recommended: <= Oracle 8.1.7 = 10 per Oracle database Oracle 9i = 32 per Oracle database Oracle10g = 120 >=Oracle11g=512</p>



HP-UX Kernel Configuration for Oracle data bases and data warehouse

Streampipes	Set this parameter to 0 during on the installation of Oracle DB will reduce the installation time. Default: 0 Recommended for installation: 0
swapmem_on (see appendix 5)	Allow Memory to Exceed Swap Space on disk. Will reserve some memory for pseudo swap (up to 20 - 25%). This memory will not be available than for Oracle SGA or other processes. For environments need to allocate more than 75 – 80% of memory the recommendation is not to use pseudo swap having the full physical swap space available on disks.. default: 1 (on) recommended: 1 (on), for environments which need to allocate more than 80% of memory to set to 0 (off)
Swchunk	Swap Chunk Size (1K Blocks) (related to kernel parameter maxswapchunks) default: 2048 recommended: 4096 (up to 65536 for large RAM)
VPS_CEILING	maximum System-Selected Page Size in kilobytes. Default: 64 Recommended: >= Oracle9i = 64 (up to 16384 = 16MB for large SGA)
VX_NCSIZE	the number of bytes to be reserved for the directory path-name cache used by the VxFS file system This tunable parameter, vx_ncsize, has been obsoleted from HP-UX 11i Version 2 and onwards; there is a change in the implementation of DNLC (Directory Name Lookup Cache) in VxFS/JFS 3.5, which uses it's own private DNLC. The kernel tunable ncsz has no impact on the size of the VxFS/JFS 3.5 DNLC.



HP-UX Kernel Configuration for Oracle data bases and data warehouse

Helpful Tools:

To find out which server you have

- model -> provide you a short information about the server
- machinfo -v -> provide you a detailed information about the server, cpu clock speed and cache, memory,

To list all kernel parameters on the system

- HP-UX11iv1/v2:
 - use SAM or
 - /usr/sbin/kmtune -l | more
- HP-UX 11iv3:
 - System Management Homepage (/usr/sbin/smh) or
 - /usr/sbin/kctune

To update a kernel parameter

- HP-UX11iv1/v2:
 - SAM (/usr/sbin/sam) or
 - /usr/sbin/kmtune -s <kernel parameter> {=|+}
- HP-UX 11iv3:
 - System Management Homepage (/usr/sbin/smh) or
 - /usr/sbin/kctune

To find out 32 or 64 bit HP-UX 11.x

- /bin/getconf KERNEL_BITS

Limitations on processes

- ulimit -a

The command sysdef analyzes the currently running system and reports on its tunable configuration parameters.

- sysdef | more

To find out 32 or 64 bit Oracle software version

```
$ cd $ORACLE_HOME/bin
```

```
$ file oracle
```

```
# if this is the answer, you have 32-bit Oracle
```

```
oracle: PA-RISC1.1 shared executable dynamically linked -not stripped
```

```
# if this is the answer, you have 64 bit Oracle
```

```
oracle: ELF-64 executable object file - PA-RISC 2.0 (LP64) # for PA-Risc systems
```

```
oracle: ELF-64 executable object file - IA64 # for Itanium systems
```





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Appendix 1: db_block_size

The default db_block_size on HP-UX is 2048. But we recommend to use the following values:

- db_block_size = 4096 to 8192 for OLTP Applications
- db_block_size = 8192 to 16384 for DSS/DW Applications

For RAC/OPS it makes still sense to use small db_block_size to reduce the amount of data to transfer between the nodes for cache fusion.

Appendix 2: Raw Devices + Asynchronous IO

We recommend to use raw devices instead of file system files for Oracle datafiles and to use asynchronous IO, both resulting in much better performance.

Use striping with a large number of disks to spread IO over multiple disks (high number is good). Define the number of disks in your stripe set also dependent of your future growth of your storage, otherwise you will get probably hot spots in the future. For the striping are 2 possibility exists. For a high number of disks we did not figured out big difference in performance.

Use the HP-UX Logical Volume Manager to stripe logical volumes across multiple disks.

Recommended stripe size: Quite often is 64KB still used

or

use HP-UX Extend-based-striping with 1 MB or even larger extents.

To implement asynchronous IO on HP-UX systems (raw devices are required)

- add the asyncdsk Driver (Asynchronous Disk Pseudo Driver) to the HP-UX Kernel (using SAM) and regenerate new kernel
- create the device file:
 - /sbin/mknod /dev/async c 101 0x100
 - chown oracle:dba /dev/async
 - chmod 660 /dev/async
 - HP-UX 11i (>=11.11)
 - To use asyncio: As root:
/usr/sbin/setprivgrp dba MLOCK
- set oracle init.ora parameters:
 - use_async_io=true (Oracle 7.3)
 - disk_async_io = true (>= Oracle 8i)
 - db_writers=1 (< Oracle 7.3)

To test the async IO

- check with fuser /dev/async





HP-UX Kernel Configuration for Oracle data bases and data warehouse

- or look with glance/gpm at the ora_dbwr.
 - look for open files:
 - /dev/async must be shown as opened file
- To check the privilege capabilities for a group, issue the command:
/usr/bin/getprivgrp
- If the output of getprivgrp(1) does not indicate that the group has the MLOCK privilege, it can be set by issuing the following command as root:
/usr/bin/setprivgrp dba MLOCK

When not using asynchronous IO

- set db_writers to number of disk (for Oracle 7)
- set db_writer_processes = 1 – 12 (>= Oracle 8i)
- set dbwr_io_slaves = 0 – 999 (>= Oracle 8i)
(Please test with your application, if the behaviour using multiple db_writer_processes is better than use multiple dbwr_io_slaves)

Please check also appendix 6 for Oracle (SCHED_NOAGE)

Using file system please have a look at least at these 3 parameters which might have an impact for Oracle DB performance from the HP-UX 11iv2, v3 vxtunefs parameter

```
max_diskq=1G
write_nstream=64
read_nstream=64
```

We have used these parameters with an EVA8000 with 112 disks. For other storage subsystems the parameter may be different.

To give you information about the VxFS 3.5 Defaults

```
read_pref_io = 65536
read_nstream = 1 (increase to number of LUNs)
read_unit_io = 65536
write_pref_io = 65536
write_nstream = 1 (increase to number of LUNs)
write_unit_io = 65536
pref_strength = 10
buf_breakup_size = 131072
discovered_direct_iosz = 262144
max_direct_iosz = 1048576
default_indir_size = 8192
qio_cache_enable = 0
write_throttle = 0
max_diskq = 1048576
initial_extent_size = 1
max_seqio_extent_size = 2048
max_buf_data_size = 8192
hsm_write_prealloc = 0
read_ahead = 1
```





HP-UX Kernel Configuration for Oracle data bases and data warehouse

Read Ahead

- JFS detects sequential pattern, prefetches data into buffer cache
- Read ahead size is calculated using $read_pref_io * read_nstream$
- Maintains 4 ranges of read ahead size
- Sequential read ahead affected by other processes or threads

Flush Behind

- Flush behind amount is calculated using $write_pref_io * write_nstream$
- When the number of dirty buffers for a file exceeds the flush behind amount, JFS will start to issue asynchronous writes to flush the dirty buffers.

Appendix 3: Parallel Query Option

When using Oracle Parallel Query Option, there are up to `parallel_max_servers` more Oracle Processes running on the system. It's necessary to increase the related HP-UX Parameters (`maxusers`, `maxuprc`, `nfile`, `nproc`)

Appendix 4: shmmax

If you got an Oracle error saying ORA-3113: end-of-file on communication channel than you have to adapt your `shmmax` to a usable value which means,

- for Oracle 32-bit max 1GB
- for Oracle 64-bit `shmmax` can be raised above 1 GB

Please remind, that you have normally more than one shared memory segments (`shmmni` (max number of shared memory segments allowed to exist simultaneously), `shmseg` (per process) are the related kernel parameters), which can be used by Oracle. If you have multiple Oracle instances on one system, it is normally better to decrease the size of shared memory and use multiple shared memory segments instead (which is done automatically).

Limitations:

- Oracle7
 - the buffer cache part can be composed of more than one segment, but the variable portion of the SGA (mainly the `shared_pool`) must fit in one shared memory segment





HP-UX Kernel Configuration for Oracle data bases and data warehouse

- Oracle8 / Oracle 9i / Oracle 10g / Oracle 11g
 - the buffer cache part can be composed of more than one segment, the variable portion of the SGA (mainly the shared_pool) can be composed of more than one shared memory segment, assuming that the shared memory segments are contiguous

Appendix 5: Swapping/Paging

- How much memory does the machine have?
 - `dmesg | grep Physical`
 - or `grep Physical /var/adm/syslog/syslog.log`
- What is the page size?
 - `dmesg | grep physical`
 - or `grep physical /var/adm/syslog/syslog.log`
- Is the machine swapping in the moment?
 - `vmstat 5 5`
 - `page.pi` indicated the number of pages moving from disk to memory
 - `page.po` indicated the number of pages moving from memory to disk
 - `memory free` indicates the amount of free memory in pages
- How much swap space is configured / used?
 - `swapinfo -a`
 - A swap area of `TYPE=dev` gives the `NAME` of the area of disk used for paging. `/etc/fstab` will show `NAME` as a swap device.
 - A swap area of `TYPE=reserve` indicates the amount of swap space that would be required if HP-UX needed to page out all in-memory pages for all processes that are currently running.
 - A swap area of `TYPE=memory` only appears if 'memory paging' is enabled, by setting the configurable kernel parameter `swapmem_on` to 1 (on). This default value for this parameter is 1. 'memory paging' is also known as 'pseudo-swap'. If present, the `AVAIL` figure is (total memory - memory used after bootup). This swap area is not actually used by HP-UX for paging (since it does not make sense to page from memory, and then back into memory).





HP-UX Kernel Configuration for Oracle data bases and data warehouse

- How much swap should I configure?
 - Oracle 9i
 - Recommendation is to have 2 x RAM size for swap at least 400 MB on a system running Oracle,
but on large systems 2*RAM size is not used anymore. You need at least the swap space that Oracle SGA fits in plus some additional swap space, even swapping out the SGA is not a good solution for a system.
(related kernel parameters for the swap are swchunk and maxswapchunks)
 - Oracle >= 10g
 - Recommendation is to have at least 1GB or 2 x size of RAM size for swap on a system running Oracle,
but on large systems with 2 GB or more 2* size of RAM size is not used anymore. You need at least the swap space that Oracle SGA fits in plus some additional swap space, even swapping out the SGA is not a good solution for a system.
(related kernel parameters for the swap are swchunk and maxswapchunks)

Appendix 6: Privileges for SCHED_NOAGE (>= Oracle 9i DB) and MLOCK

The SCHED_NOAGE policy gives processes holding a latch a fixed priority and makes them nonpreemptable during this time. This causes less latch waits and latch sleeps, hence higher throughput.

To enable the SCHED_NOAGE policy for Oracle, the following init.ora parameter needs to be added:

hpux_sched_noage=154 (for HP/UX 11.0)

hpux_sched_noage=178 (for >= HP/UX 11i)

In addition, the system privileges RTPRIO and RTSCHED need to be added to the DBA group as root in

/etc/privgroup:

dba RTPRIO RTSCHED

Complete root user set-up tasks:

The MLOCK privilege grants Oracle the ability to execute asynch I/Os through the HP asynch driver. Without this privilege, Oracle DB generates trace files with the following error message: Ioctl

ASYNCH_CONFIG error, errno = 1

As root, do the following:

_ If it does not already exist, create the /etc/privgroup file. Add the following line to the file:
dba MLOCK RTSCHED RTPRIO

_ Use the following command syntax to assign these privileges:

```
#setprivgrp groupname privileges
```

In the preceding command, *groupname* is the name of the group that receives the privileges, and *privileges* are the privileges that are granted to the group.





HP-UX Kernel Configuration for Oracle data bases and data warehouse

For example, to set the privileges for the OSDBA group, which is typically named dba, enter the following command:

```
#setprivgrp dba MLOCK RTSCHED RTPRIO
```

Appendix 7: ccNUMA Architecture: cell local memory vs. interleaved memory

ccNUMA (cache coherent Non-Uniform Memory Architecture) systems offer programmers and users the simplicity and flexibility of symmetric multiprocessing (SMP) with the memory scalability of clusters. In a ccNUMA system, processors, memory, and I/O are grouped together into cells. The latency and bandwidth characteristics of communication within a cell are faster, while going outside a cell is slower. Since the memory in ccNUMA systems is physically distributed but logically shared, these systems offer better performance to applications that are optimized to use their features. For non-optimized applications, they still offer better performance since the default behavior is designed to be benign—if not beneficial—and they still have access to much larger shared resources of memory, CPUs, and disk space. It's the basic technology that's used to build large machines like rx7640, rx8640 or Superdome.

Interleaved memory

Interleaved memory is memory for shared objects or data structures. A portion of memory is taken from cells of the system—typically all of the cells—and is mixed together in a round robin fashion of cache-line-size chunks. It has the characteristic that memory accesses take a uniform amount of time. In other words, it has uniform latency no matter which processor accesses it.

Using interleaved memory and not having cell local memory the following Oracle init.ora parameter needs to be set to switch off and cell based systems.

Customers can download and apply patch for Bug 8199533 to disable NUMA support. This is a database patch and should be applied to the Database home. This patch is available for common platforms on 10204 and 11.1.0.7 releases..

Oracle support does not recommended using init.ora parameter

"_enable_NUMA_optimization" to disable NUMA. Customers should apply fix for Patch 8199533 to disable NUMA . The patch is rolling upgradeable.

(The patch is setting the underscore parameter_enable_numa_optimization=FALSE and _db_block_numa=1)

For more information have a look at http://docs.hp.com/en/4913/ccNUMA_White_Paper.pdf

For additional information about Oracle on HP Itanium Server have a look at the following white paper <http://h20219.www2.hp.com/ERC/downloads/4AA2-0547ENW.pdf>

The white paper talks about Oracle on ccNUMA architecture and resource management with Oracle on HP-UX.

You also find more information at <http://metalink.oracle.com> at the metalink note called “Oracle Database ccNUMA support and dynamic partitioning on HP-UX” with the ID 761065.1

