

Platform Communications

Title: Operating System, RDBMS & Additional Component Patches Required for Installation on PT 8.49

Topic: CC - Platforms

Date:

Operating System, RDBMS & Additional Component Patches Required for Installation PeopleTools 8.49

**NEW FOR PeopleTools 8.49:
New minimums: JRE 1.5.0, Tuxedo 9.1, Weblogic 9.2,
WebSphere 6.1.0.3
TRU64 is no longer supported**

▼ Operating System Patches

PLEASE READ THE FOLLOWING IMPORTANT NOTES BELOW, which apply to all Unix and NT installations:

Note #1: PATCHES: For Unix and NT users, we list the "minimum patch version required". In many cases a patch becomes superseded and is no longer available. Our policy is to build against the "lowest common denominator" meaning that we do not automatically test new or superseding patches. If you come across a patch that has been superseded, please feel free to use the new one, as they generally incorporate the fix originally required. We only certify new operating system releases, not patches. When we discover a patch that WE require, we post it in the tables contained in this document. Your RDBMS or Third Party vendors may require additional or different patches; please verify this with them.

Note #2: COMPILERS: We do not require a C or C++ compiler for any of the Unix or NT platforms, when installing or operating our Batch or Application Servers in a standard fashion. We test and certify our products using standard C and C++ runtime libraries; however, do not confuse this with a requirement for a compiler. If for some reason a platform is missing a C or C++ runtime libraries that we are dependent on, then you will need to load them. (For example, for Solaris you may download the compiler suite, but you do not need to purchase a license for it, as we need the libraries but not the compiler itself). **We do require a COBOL compiler for those of our applications that use COBOL.** Additionally, some 3rd-party apps may require re-linking. All of the Unix platforms come standard with linkers and loaders. **If you purchase a product or develop a product in-house that requires you to write to the standard API's we provide, then you will require a compiler.** The type of compiler needed will be based upon whether you are writing C or C++ source code.

Note #3: JRE: You must get the **LATEST patches for your Operating System** from your Operating System vendor's website for the Java Runtime Engine (JRE) you have installed. (Please see the JAVA

section for patch information for JRE itself.)

- ▶ AIX
- ▶ HP-UX PA Risc
- ▶ HP-UX IPF
- ▼ Linux

Operating System	Bundle/Patch Level	Symptom
SuSE 9.0	2.6.5-7 (on both x86-32 and s390x zseries) SP2 on x86-32 and SP3 on s390x	Minimum Release Level
RedHat Linux Advanced Server 4.0 and Oracle Enterprise Linux 4 Update 4. (EL4.4)	2.6.9-11 (on both x86-32 and s390x zseries)	Minimum Release Level

- ▶ Solaris
- ▶ Windows
- ▶ z/OS

▼ RDBMS Patches

- ▼ Oracle

- ▼ Oracle 11g:

- ▼ **Note #1: Oracle 11g - Oracle shared library name renamed to libclntsh.so.11.1 (Linux/UNIX)**

A specific PeopleTools release is compatible with the Oracle database version it was released on, as well as any subsequent Oracle database version supported by PeopleTools. Out of the box, PeopleTools will work with Oracle 9i. However, for Oracle 11g, a symlink is required.

In Oracle 11g, the Oracle shared library name is libclntsh.xx.11.1, while PeopleTools looks for libclntsh.xx.9.0.

In order for Application Server to boot up and PSRUN.MAK to run successfully, a Unix symlink must be defined from libclntsh.xx.11.1 to libclntsh.xx.9.0 (Note: The '.xx' suffix is platform specific. Solaris, Linux, and HP-Itanium use '.so'; HP-PARISC uses '.sl'). To accomplish this:

- 1) Log in as the Oracle owner.
- 2) `cd $ORACLE_HOME/lib`
- 3) Depending on your platform, create the symlink:
 - a) For Oracle11g on Solaris, RH 4.0, SuSE 10, or HP-Itanium, execute the following command to establish the symlink:
`ln -s $ORACLE_HOME/lib/libclntsh.so.11.1 libclntsh.so.9.0`
 - b) For Oracle11g on HP-PARISC, execute the following command to establish the symlink:
`ln -s $ORACLE_HOME/lib/libclntsh.sl.11.1 libclntsh.sl.9.0`

After the symlink is established, you may rerun PSRUN.MAK or boot up the Servers.

RDBMS	Patch Level and Comments	Symptoms
Oracle 11gR1	<p>For customers using Oracle 11.1.0.6, the following patches are required in addition to the base release 11.1.0.6.0:</p> <p>UNIX and Linux platforms - apply Patch #6145570.</p> <p>Windows x86 platforms - Patch #6145570 will be incorporated into the next quarterly Windows bundle for 11g.</p>	Patch #6145570 addresses 'LIKE CLAUSE RETURNS WRONG RESULTS WITH VIEW - UNION ALL' errors found while testing FSCM 8.9/9.0 during off-cycle certification.

▼ **Oracle 10g:**

▼ **Note #1: Oracle 10g, Oracle shared library name renamed to libclntsh.so.10.1 (Linux/Unix)**

A specific PeopleTools release is compatible with the Oracle database version it was released on, as well as any subsequent Oracle database version supported by PeopleTools. Out of the box, PeopleTools will work with Oracle 9i. However, for Oracle 10g, a symlink is required (Note: This is not required for IBM AIX).

In Oracle 10g the Oracle shared library name is libclntsh.xx.10.1, while PeopleTools looks for libclntsh.xx.9.0.

In order for Application Server to boot up and PSRUN.MAK to run successfully, a Unix symlink must be defined from libclntsh.xx.10.1 to libclntsh.xx.9.0 (Note: The '.xx' suffix is platform specific. Solaris, Linux, and HP-Itanium use '.so'; HP-PARISC uses '.sl'). To accomplish this:

- 1) Log in as the Oracle owner.
- 2) `cd $ORACLE_HOME/lib`
- 3) Depending on your platform, create the symlink:

a) For Oracle10g on Solaris, RH 4.0, SuSE 9.0, or HP-Itanium, execute the following command to establish the symlink:

`ln -s $ORACLE_HOME/lib/libclntsh.so.10.1 libclntsh.so.9.0`

b) For Oracle10g on HP-PARISC, execute the following command to establish the symlink:

`ln -s $ORACLE_HOME/lib/libclntsh.sl.10.1 libclntsh.sl.9.0`

After the symlink is established, you may rerun PSRUN.MAK or boot up the Servers.

▼ **Note #2: Required INIT.ORA Parameters for Oracle 10g R1.**

▼ **Change #1 - `_optimizer_cost_based_transformation=off`**

During the course of HCM 8.9 Global Payroll testing we discovered a new ORA-00600 bug. This problem occurred while executing specific SELECT statements within Global Payroll. Although Bug# 4182639 had been filed, the Oracle database team recommended that PeopleSoft customers set an init.ora parameter (described below) for Oracle 10g versions 10.1.0.3.0 and 10.1.0.4.0 as an immediate workaround to this bug. Although the Oracle database team indicates this problem no longer exists in Oracle 10.1.0.5, we've found some regressions and have filed additional bugs (Bug#s 4939157 and 5092688) for the regressions found. .

BUG 5092688 - SQL EXECUTION REGRESSION RETURNING INCORRECT# OF ROWS ON 10.2.0.1.0

BUG 4939157 - ORA-07445 [EVAOPN2()+209] WITH A FUNCTIONAL-INDEX

The following init.ora parameter still needs to be in place for Oracle 10.1.0.5. The bugs have only been fixed for Oracle 10.2.0.3 (and above, assuming no regressions occur).

`_optimizer_cost_based_transformation=off`

According to Oracle database support: "This parameter controls whether query transformations (sub-query unnesting, complex view merging and push join predicate) are performed using an algorithm that relies only on optimizer cost (new in 10g) or an algorithm that uses a combination of optimizer cost and heuristics (pre-10g). "

▼ **Note #3: Workaround for performance issue with 'ignore dup' processing**

This problem has been documented by bug#s 5103126 and 4618615. The clearest description of the problem is from Bug#4618615.

Excerpt from bug# 4618615

PROBLEM:

When you insert against a unique / primary key constraint / index and the system returns ORA-00001, the following query is issued against the dictionary:

```
select c.name, u.name
from con$ c, cdef$ cd, user$ u
where c.con# = cd.con#
and cd.enabled = :1
and c.owner# = u.user#
```

The response time of this query is too long.

DIAGNOSTIC ANALYSIS:

The access plan the query uses is:

```
-----  
--  
| Id | Operation | Name | Rows | Bytes | Cost (%CPU)|  
-----  
--  
| 0 | SELECT STATEMENT | | 6062 | 278K | 38 (6)|  
|* 1 | HASH JOIN | | 6062 | 278K | 38 (6)|  
| 2 | TABLE ACCESS FULL | USER$ | 52 | 780 | 2 (0)|  
|* 3 | HASH JOIN | | 6062 | 189K | 35 (3)|  
|* 4 | TABLE ACCESS FULL | CDEF$ | 6062 | 48496 | 26 (0)|  
| 5 | TABLE ACCESS FULL | CON$ | 7722 | 180K | 8 (0)|  
-----  
--
```

Due to the full table scan on cdef\$, which can have more than 100000 Rows we have a performance problem.

The workaround recommended from Oracle support is to delete the statistics on the following three catalog tables :

**analyze table sys.CDEF\$ delete statistics;
analyze table sys.CON\$ delete statistics;
analyze table sys.USER\$ delete statistics;**

This bug has been fixed in 10.1.0.6.0 forward, 10.2.0.2.0 forward and 11g. If you are on these versions or higher this workaround is no longer necessary.

▼ **Note #4: Recommended INIT.ORA Parameters for Oracle 10g (R1 or R2) for all platforms:**

Problem:

Oracle 10g and "Select Distinct" SQL:

For any query with Distinct and with no ORDER BY Clause, the Oracle 10g optimizer now does HASH UNIQUE instead of SORT (Unique) as before. Rows that are returned from "Select Distinct" statements are now sorted by this internal hash key if there is no "order by" clause (as opposed to by the record key). This is not an Oracle bug, but a change in behaviour by design.

This changes in behaviour does not impact the integrity of the result set - the correct rows are still returned - although they may not be in the order one would expect if you're reviewing the list visually as compared to 9i. Note, this doesn't always appear - whether the symptom is noticeable or not seems to vary by the number of rows returned.

There are quite a few statements in PeopleSoft applications which will have functional impact because of this optimizer change.

Solution:

Oracle has never stated that DISTINCT will return results in SORTED order by Default.

In order to have results in sorted order, applications should have always be coded using ORDER BY when using DISTINCT clause. The solution to this issue is that development can change application and Tools generated statements (eg, like those "Select Distinct" SQL's generated for drop down lists on PIA pages) to add the ORDER BY clause where necessary.

This problem has been identified and the Tools changes necessary will be available with the PT8.49 release and backported to the PT8.48 release. The applications will be adding the ORDER BY clause where necessary, in subsequent patch releases.

Workaround:

Until the coding changes are made the customer may use the workaround below to get the same functionality they experienced with 9i when using SELECT distincts with no order by clause.

In the init.ora change set the following flag and bounce the SID.

[_gby_hash_aggregation_enabled=false](#)

A complete description of the problem can be referenced on Metalink by note# 341838.1

Problem:

Poor performance with SQL STATEMENTS utilizing nested subqueries.

Workaround:

[_unnest_subquery = false](#)

RDBMS	Patch Level and Comments	Symptoms
Oracle 10gR1 Version 10.1.0.5	<p>There are no known required patches for 10.1.0.5 at this time.</p>	
Oracle 10gR2 Version 10.2.0.3.0	<p>For customers using Oracle 10gR2 (10.2.0.3.0)*, the following patches are required in addition patchset 10.2.0.3.0:</p> <p>Patch# 5128368 addresses ORA-3113 errors found while testing HCM 9.0.</p> <p>Solaris 10 x86-64</p>	<p>Patch# 5128368 addresses ORA-3113 errors found while testing HCM 9.0.</p> <p>Patch# 5574966 addresses ORA-3113 errors found while testing FSCM 9.0.when selecting from PS_AUC_AWD_BIDNAME in FSCM9 application</p>

customers: Please download Patch# **6061749** instead of Patch# **5128368**.

Patch# 5574966 addresses ORA-3113 errors found while testing FSCM 9.0 when selecting from PS_AUC_AWD_BIDNAME in FSCM9 application

Patch# 5683652 - Addresses Oracle 10gR2 Client performance regression with freeing OCIHandles. This problem has been identified with causing up to 100% regression in PeopleTools.

Windows 32bit platform specific versions of these patches are available as:

Patch# 5946186 - Platform specific patch for Win32 for bug# 5128368

Patch# 6012742 - Platform specific patch for Win32 for bug# 5574966

Patch# 6215383 - Platform specific patch for Win32 for bug# 5683652 (OCIHandle performance regression)

Note: Apply patch to both client and server installs.

The fix for bug# **5683652** is available for all other Windows OS platforms(x86 -64bit, etc.) as Windows rollup Patch 10.2.0.3.9

All of the aforementioned

Patch# 5683652 - Addresses Oracle 10gR2 Client performance regression with freeing OCIHandles. This problem has been identified with causing up to 100% regression in PeopleTools.
Note: Apply patch to both client and server installs.

patches will be rolled into the next available patchset 10.2.0.4.0.

* Patches mentioned are to be applied by all 10.2.0.3 customers. Pre-Apps 9.0 customers using 10.2.0.3 will also need to apply these patches.

▶ Oracle 9i:

- ▶ DB2 UDB for OS/390
- ▶ DB2 UDB for Linux, UNIX and Windows
- ▶ Microsoft SQL Server
- ▶ Sybase
- ▶ Informix

▼ **Additional Component Patches**

- ▶ Crystal Reports
- ▶ BEA Tuxedo
- ▶ Oracle Application Server
- ▶ Micro Focus COBOL
- ▶ JRE
- ▶ Oracle Portal
- ▶ WebLogic