

Oracle® Parallel Server

Addendum

Release 3 (8.1.7) for Alpha OpenVMS

December 2000

Release 3 (8.1.7)

The OpenVMS release of Oracle 8i Release 3 (8.1.7) includes the Oracle Parallel Server option. Oracle Parallel Server allows multiple database instances to share the same database files. For complete information about configuring and using Oracle Parallel Server, please see the *Oracle Parallel Server Concepts and Administration manual* for Release 3 (8.1.7).

Oracle8i Release 3 (8.1.7) now incorporates the use of the Oracle Group Membership Services within the Oracle Parallel Server. As of this release and future releases there is no need to start up the OGMS daemon. This process is now automatically started from within the Oracle Parallel Server.

Note: There still is a need to have one OGMS home and install area per system.

More information about starting, configuring, and using the OGMS daemon are given in the *Oracle Parallel Server Concepts and Administration manual*.

Oracle Parallel Server for Oracle8i incorporates the Oracle Distributed Lock Manager. However, a limited number of native VMS locks are also used in the clustering interface. The usage of VMS locks has been greatly reduced from that of Parallel Server for Oracle 7.3.

Use of OpenVMS 7.2

In order to use Oracle Parallel Server for OpenVMS version 8.1.7, your system **MUST** be running Alpha OpenVMS version 7.2 or above. Also, depending on your OpenVMS version, you need have applied one of the Compaq ECO's as of October 5, 2000:

ORACLE®

Copyright © 2000, Oracle Corporation.
All Rights Reserved.

Oracle, the Oracle logo, Net8, Pro*COBOL, Pro*FORTRAN, SQL*Module, and SQL*Plus are registered trademarks of Oracle Corporation. Advanced Networking Option, Advanced Replication Option, Oracle Applications, Oracle Call Interface, Oracle Enterprise Manager, Oracle Installer, Oracle Names, Oracle Parallel Server, Oracle Server Manager, Oracle WebServer, Oracle7 Server, Oracle8, PL/SQL, and Pro*C/C++ are trademarks of Oracle Corporation.
All other product or company names mentioned are used for identification purposes only, and may be trademarks of their respective owners.

The Compaq ECO details are as follows:

For OpenVMS 7.2-1:

- VMS721_SYS-V0600 Alpha V7.2-1 System ECO
- OpenVMS VMS721_ACRTL-V0200 Alpha V7.2-1 Compaq C RTL ECO
- TCP/IP TCPIPALP_E01A50 TCP/IP for OpenVMS V5.0A

For OpenVMS 7.21-H1:

- OpenVMS VMS721H1_UPDATE-V0300 Alpha V7.2-1H1 UPDATE ECO
- OpenVMS VMS721H1_ACRTL-V0100 Alpha V7.2-1H1 Compaq C RTL ECO
- TCP/IP TCPIPALP_E01A50 TCP/IP for OpenVMS V5.0A ECO

These ECOs can be downloaded from the Compaq Services website. Oracle8i Release 3 (8.1.7) Parallel Server makes use TCP/IP services. For more information about configuring and using the TCP/IP feature, please see your OpenVMS system or OpenVMS TCP/IP documentation.

Changes from Oracle7

Note that Oracle8i Parallel Server has many architectural changes from the version supplied with Oracle 7. Users are urged to review the new features section of the *Oracle Parallel Server Concepts and Administration manual* for a complete discussion of these changes.

The keywords PARALLEL and EXCLUSIVE have been obsoleted from the STARTUP command. In Oracle7, this was reflected in the creation of STARTUP_EXCLUSIVE_<dbname>.COM, STARTUP_EXCLUSIVE_<dbname>.SQL, STARTUP_PARALLEL_<dbname>.COM, and STARTUP_PARALLEL_<dbname>.SQL files when you set up your Parallel Server instance. Since these keywords have been removed from Oracle 8, now only STARTUP_<dbname>.COM and STARTUP_<dbname>.SQL files are created when you create a new database.

Oracle8i Parallel Server requires the installation and operation of the Oracle Group Membership Services daemon. The Daemon must be installed and run on every system where you want to run Parallel Server.

The keywords used inside of the `init.ora` parameter file for Parallel Server have been changed. Users should not simply copy a pre-existing options file into a new installation. Please refer to the *Oracle Parallel Server Concepts and Administration manual* for complete information on these changes.

Because of these changes, a `psinit.ora` file is no longer created when you create a new database under Oracle8i.

Oracle Group Membership Services

In the previous versions of Oracle8, only one Oracle Group Membership Services (OGMS) Daemon could be run per system. Since the Oracle Group Membership Services have been incorporated into Oracle8i Parallel Server Instance, this restriction still applies. Only one Oracle8i Parallel Server Instance can be run per system. In addition, only one Oracle8 OGMS daemon, or Oracle8i Parallel Server Instance can be running per system. Each instance or OGMS daemon that runs on a system must be able to access the home and installation directories of the OGMS product area.

The OGMS Daemon is a quasi-separate product under Oracle8i. Installation of the Daemon is selected by an option under the Oracle PSOPT Build Options menu of the OpenVMS Oracle installer.

The installation procedure for Oracle Group Membership Services runs as an optional part of the Parallel Server Option (PSOPT) installation procedure. The OGMS installation creates a new directory that you specify, and places within it a directory called OGMS.DIR, which is the directory into which OGMS will be installed. In addition, separate directories for each node using this OGMS installation are created as a subdirectory of the OGMS.DIR directory. These directories contain the logs and other files of the OGMS daemons running on the different nodes.

You must have the appropriate privilege on your system to create the OGMS installation directories. The directories created will have the ORA_DBA access right granted to them during installation. This allows any account granted the ORA_DBA access right to administer the OGMS Daemon.

The installation procedure creates a `define_ogms.com` file in the OGMS installation directory. This procedure makes the appropriate system-wide definitions necessary to run the OGMS Daemon.

Two system wide logical names are defined: OGMS_INSTALL and OGMS_HOME. The OGMS_INSTALL logical name points at the OGMS.DIR directory created during installation. OGMS_HOME points at one of the per-node subdirectories of the OGMS.DIR directory. For example, if you have a system called AP8VMS, and use the defaults under the PSOPT menu, the following directories and definitions would be made:

```
OGMS_INSTALL
```

```
as
```

```
SYS$COMMON: [SYSEXE.Oracle.OGMS]
```

```
OGMS_HOME
```

```
as
```

```
SYS$COMMON: [SYSEXE.Oracle.OGMS.AP8VMS]
```

The `define_ogms.com` file must be run before launching the OGMS Daemon. A reference to this file can be put into your system startup files if you want the definitions to be in effect after a system restart.

The OGMS Daemon is controlled by the `ogmsctl` command. You can start and stop, query the basic status of the Daemon, and get a list of the membership of the OPS cluster with `ogmsctl`. For more information about the commands available to the OGMS Control program, please see the *Oracle Parallel Server Concepts and Administration* manual. For Oracle8i the `ogmsctl` command is not used and should not be invoked.

In Oracle8i there is no separate OGMS Daemon process. For previous Oracle8 installations when launched, the OGMS Daemon ran as process `ORA_GMS`. It was possible to effect the quotas given to this process by defining the appropriate logical name. This has no effect in Oracle8i. Thus, for Oracle8 only, to change the quota of the OGMS process, we define the following system wide logical name:

```
ORA_GMS_PQL$_<quota>    <new quota value>
```

For example, to boost BYTLM to 1 million, you would define:

```
$ define/system OGA_GMS_PQL$_BYTLM 1000000
```

Installation of the Oracle Parallel Server option

The Parallel Server Option is installed by selecting the PSOPT menu choice for the Oracle Installer Build and Load Options menus.

Note: You must link the Parallel Server Option into the RDBMS Server before Parallel Server operation is possible. This is done by selecting YES to the question on whether to include the Parallel Server Option in the RDBMS Server Build Options menu.

The PSOPT menu choice on the Build Options menu allows you to select a protection mode for the Parallel Server directories, as well as whether you want to install OGMS. The default menu for PSOPT is reproduced below:

Note: For Oracle Parallel Server 8.1.7.0 running on Alpha OpenVMS, there is a new default option included in the PSOPT build configuration screen. Item number 5 below is the inclusion of the parallel server communication option that will be included in your server.

PSOPT Configuration Options

OptionCurrent	Value
1. System or Group installation? [S/G]	S
2. Install Oracle Group Membership Services? [Y/N]	Y
3. Directory for OGMS install?	PTST05:[MJNOEL_816_DROP4A]
4. Override previous OGMS installation? [Y/N]	Y
5. Parallel Server Interprocess Communication package (TCP)?	TCP

Enter (A)LL to select all options.

Enter (E)XIT to exit this menu with selected options.

Enter (Q)UIT to quit this menu with no action.

You may select to have the installer override a current installation, if there are currently definitions of OGMS_INSTALL and OGMS_HOME. If you do not select YES for option 4, then the OGMS installation will fail if these logical names already exist.

If you have an OGMS Daemon running already on your system, the installer will create the directories you have indicated, if override is also selected, but it will not override the current definitions of OGMS_INSTALL and OGMS_HOME, nor will it shut down your currently running OGMS Daemon.

Please be cautioned that shutting down the OGMS Daemon will terminate any OPS instances running on that system.

Given below is the OGMS Installation procedure output for a default installation.

```
- Running ORA_PSOPT_BLD.COM.  
Building Oracle Group Membership Services.
```

The Oracle Group Membership Services will be installed into an OGMS directory which will be created in a directory that you have specified. Per node directories for each node which uses this installation will be created as OGMS Daemons are launched. Note that this directory must be accessible to all Parallel Server database instances that you wish to run on this node.

Also note that all database instances which participate in an OPS cluster must be able to access an OGMS installation. One OGMS Daemon can be run at a time on any cluster node.

Creating OGMS directory in SYS\$COMMON:[SYSEXE].

```
Creating OGMS install directory SYS$COMMON:[SYSEXE.ORACLE.OGMS].
Creating define_ogms.com.
Running define_ogms.com.
Copying files to SYS$COMMON:[SYSEXE.ORACLE.OGMS].
Linking OGMS 64 bit daemon
- Linking OGMS.EXE
Linking OGMS 64 bit control program
- Linking OGMSCTL.EXE
Setting protection on OGMS files.
Purging OGMS files.

OGMS installation complete.
```

Reproduced below is the output of the OGMS Installation procedure if the user has selected to override a current installation.

```
-----
- Running ORA_PSOPT_BLD.COM.
Building Oracle Group Membership Services.
```

An OGMS installation is already present on system AP8VMS.

OGMS_INSTALL is currently defined to be

```
SYS$COMMON:[SYSEXE.ORACLE.OGMS].
```

Overriding previous installation

The Oracle Group Membership Services will be installed into an OGMS directory which will be created in a directory that you have specified. Per node directories for each node which uses this installation will be created as

OGMS Daemons are launched. Note that this directory must be accessible to all Parallel Server database instances that you wish to run on this node.

Also note that all database instances which participate in an OPS cluster must be able to access an OGMS installation. Only one OGMS Daemon can be run at a time on any cluster node.

```
Creating OGMS directory in SYS$COMMON:[SYSEXE].
The OGMS directory already exists.
Creating define_ogms.com.
Running define_ogms.com.
%DCL-I-SUPERSEDE, previous value of OGMS_INSTALL has been superseded
%DCL-I-SUPERSEDE, previous value of OGMS_HOME has been superseded
Copying files to SYS$COMMON:[SYSEXE.ORACLE.OGMS].
Linking OGMS 64 bit daemon
- Linking OGMS.EXE
Linking OGMS 64 bit control program
- Linking OGMSCTL.EXE
Setting protection on OGMS files.
Purging OGMS files.
```

OGMS installation complete.

Addition to Createcat Procedure

Before operating your database in Parallel Server mode, be sure to add the following line to the end of the ora_db:createcat_<dbname>.sql file in your database (ora_db) directory:

```
@ora_rdbms_admin:catparr.sql
```

If you have already created the database before adding the reference to catparr, simply do the following:

```
connect internal as sysdba
@ora_rdbms_admin:catparr.sql
```

Quota Considerations for Oracle Parallel Server

Oracle8i Parallel Server Option is a completely rearchitected feature from the Parallel Server Option for the Oracle7 Server. Oracle8i now includes a portable Distributed Lock Manager (DLM) internal to the database server.

Because of this, the usage of ENQLM quota by the Parallel Server Option has been greatly reduced. Large ENQLM boosts for background processes are no longer necessary in the Oracle8i environment.

If you have an Oracle8 Parallel Server installation that uses the new OpenVMS Intra-Cluster Communications (ICC) feature, you should note that ICC buffers are stored in non-paged pool associated with each process. This storage is charged against the process pooled BYTLM quota.

In addition, some ICC data is stored in the P1 pool of a process. This storage is allocated by the setting of the CTLPAGES parameter in SYSGEN. The default setting of 256 pages is adequate for many Parallel Server environments. This is not a dynamic parameter, so a change here will require a system reboot.

The following formula may be used to estimate the BYTLM and P1 pool requirements for ICC consumption of each process:

NCon Number of connections

$NCon * 2048 * 6$ bytes of BYTLM quota

$Ncom * 256$ bytes of P1 Pool

Each database process connects with the OGMS Daemon. If there are N instances, then N would be the value for NCon.

However, Parallel Query Slaves connect in a mesh with all other Parallel Query Slaves. If there are 50 PQS per instance, and four instances, then each PQS will make connections with $200-1 = 199$ neighbors.

In addition, the Lock Manager Process, LMD0, connects with all other database backgrounds in its instance, as well as with all other LMD0s in the other instances. If there are 4 instances, and each instance has 10 normal backgrounds and 50 Parallel Query Slaves, then the Lock Manger connects with $10 + 50 + 4 + 1 = 65$ different processes.

Please note that these quota estimates are in addition to quota required by non-Parallel Server functions of the Oracle server.

The Process Quota estimation functions of the Oracle Server attempt to make reasonable calculations for the various backgrounds launched during server operations. You can change the quota allocated to a background process by defining certain system logical names before the instance is started.

To effect the process quota given to all backgrounds in an instance, define the following:

```
$ define/system ORA_<sid>_PQL$_<quota> <quota_value>
```

If you have a SID named TOM1, and you wish to boost the BYTLM to 2 megabytes, you would use the following:

```
$ define/system ORA_TOM1_PQL$_BYTLM 2000000
```

Or, you could effect the process quota of a particular background by defining the following:

```
$ define/system ORA_<sid>_<process>_PQL$_<quota> <quota_value>
```

If you wanted the LMD0 process of the TOM1 SID to get 3 megabytes of BYTLM, you would use the following:

```
$ define/system ORA_TOM1_LMD0_PQL$_BYTLM 3000000
```

Analyzing OGMS Log files

The Oracle GMS Daemon creates several files in the OGMS_HOME directory, including Daemon and Control file traces. These files may be useful in debugging OGMS and Parallel Server problems.

The following files are created in the OGMS_HOME directory when the OGMS control program starts the OGMS Daemon:

ogms_llnmap.dat	Local node map This is a binary format file which must not be altered.
gms0000.dat	Daemon contact information This is a binary format file which must not be altered.
cntl_<proc-id>.trc	OGMS control program trace file Traces execution of OGMSCTL command.
daem_<proc-id>.trc	OGMS daemon trace file Traces execution of the OGMS daemon.

Note: All status codes listed in these trace files by the ICC driver are printed in HEXADECIMAL, not decimal. These values are indicated by the prepending of a "0x" to the data output.

There is currently a problem in the ICC support of VMS 7.2 and 7.2-1 such that the ICC system calls can report a status of 25fc(hex), 9724(dec) or:

%SYSTEM-F-INSFP1POOL, insufficient process dynamic memory

when in reality the error is a shortage of BYTLM in the affected process. If this status code is listed in a message showing a status of 294(hex), 660(dec), or:

`%SYSTEM-F-REJECT, connect to network object rejected`

with a remote status of 25fc(hex), then the affected process is the remote side of the command in question.

The BYTLM of the affected process can be increased by defining the proper logical name. Please see the above section "Quota Considerations for Oracle Parallel Server" for more information.

It is also possible for the 25fc status to correctly indicate a shortage of P1 pool, in which case you must increase the SYSGEN parameter CTLPAGES and reboot the machines involved.