Oracle Parallel Server™ Administrator's Guide for Silicon Graphics® Systems

Document Number 007-2763-001

CONTRIBUTORS

Written by Carolyn Curtis
Illustrated by Dan Young and Carolyn Curtis
Production by Derrald Vogt
Engineering contributions by Sharad Srivastava, Ravi Chhabria, Kireeti Kompella, Ewan McKissock, Raj Das, and Kathy O'Connell
Cover design and illustration by Rob Aguilar, Rikk Carey, Dean Hodgkinson, Erik Lindholm, and Kay Maitz

Oracle Parallel Server™ Administrator's Guide for Silicon Graphics® Systems
Document Number 007-2763-001

Contents

List of Figures

List of Tables

About This Guide

The Oracle Parallel Server™ (OPS) represents a significant development in RDBMS capability for ORACLE® users. The ability to cluster two computer systems and coordinate their access to a single shared Oracle database provides for increased throughput and capacity with higher levels of application and database availability.

The Silicon Graphics® OPS consists of the following hardware components:

- two CHALLENGE™ servers
- optionally, for each CHALLENGE DM, L, or XL server, an external peripheral enclosure for SCSI storage devices: CHALLENGE Vault M, CHALLENGE Vault L, or CHALLENGE Vault XL, respectively
- one IRISconsole™, consisting of a 24-bit Indy™ workstation with a minimum of 32 MB of memory, running IRIX™ 5.3, XFS, an IRISconsole ST-1600™ serial port multiplexer, cables, and software
- one CHALLENGE RAID deskside storage system with two storage-control processors (SPs) and at least five disk modules; all disk modules can be either 2 GB or 4.3 GB
- required hardware upgrades and cables

The software for OPS consists of IRIX 5.3 with XFS, IRIX patches, OPS software for Silicon Graphics systems, IRISconsole software, FDDI software, and software for the component systems, such as the Indy workstation and the CHALLENGE RAID storage system. Optional software includes Performance Co-Pilot™ (PCP), IRIXPro™, Database Accelerator (DBA), and IRIX NetWorker™ 4.1.1. Software components of OPS that must be obtained from Oracle Corporation are ORACLE RDBMS and the Parallel Server Option.

Audience

This guide is written for the person who administers the OPS system. The OPS administrator is familiar with Oracle RDBMS in general and the specific Oracle database instances running on the CHALLENGE servers. The OPS administrator is also familiar with the operation of the Indy workstation and the CHALLENGE servers, as well as the optional Vault storage systems and CHALLENGE RAID storage system, if they are used in the OPS configuration. The OPS administrator has acquired familiarity with the Oracle Parallel Server system and uses IRISconsole to control the OPS hosts.

Structure of This Document

This guide contains the following chapters and appendix:

Chapter 1	"Features of the Oracle Parallel Server" explains capabilities of the Oracle Parallel Server and explains how system components work together.
Chapter 2	"Using the OPS System" explains how to configure and start the OPS system software and how to halt the OPS system.
Appendix A	"Troubleshooting" lists possible problem situations and suggests solutions.

Conventions

These type conventions and symbols are used in this guide:

Helvetica Bold	Hardware labels
Italics	Executable names, filenames, IRIX commands, manual or book titles, new terms, program variables, tools, utilities, variable command line arguments, variable coordinates, and variables to be supplied by the user in examples, code, and syntax statements
`Fixed-width type`	Error messages, prompts, and onscreen text

Bold fixed-width type
User input, including keyboard keys (printing and nonprinting); literals supplied by the user in examples, code, and syntax statements (*see also* <>)

“” (Double quotation marks) Onscreen menu items and references in text to document section titles

[] (Brackets) Surrounding optional syntax statement arguments

<> (Angle brackets) Surrounding nonprinting keyboard keys, for example, **<Esc>**, **<Ctrl-D>**

Chapter 1

Features of the Oracle Parallel Server

This chapter introduces the features and capabilities of the Oracle Parallel Server (OPS). It explains

- OPS configuration
- OPS instances and domains
- OPS architectural features
- OPS components and how they work together
- OPS performance and optional Silicon Graphics software

OPS Configuration

OPS is a collection of Oracle instances running on separate CHALLENGE servers, providing simultaneous access to the same physical database. The physical database is the same as that for an ordinary non-OPS (nonparallel) Oracle RDBMS, except that it has separate redo logs and rollback segments for each instance. The redo log file is a compressed record of changes that a transaction has made.

For the Silicon Graphics platform, OPS is available in a dual-host configuration; each node can access the same shared disk storage. The hosts can be two CHALLENGE DM systems, two CHALLENGE L systems, or two CHALLENGE XL systems. A third host, a 24-bit Indy workstation running the IRISconsole software, functions as the OPS node controller and as a single point of administration for OPS. Figure 1-1 diagrams Silicon Graphics OPS hardware configuration.

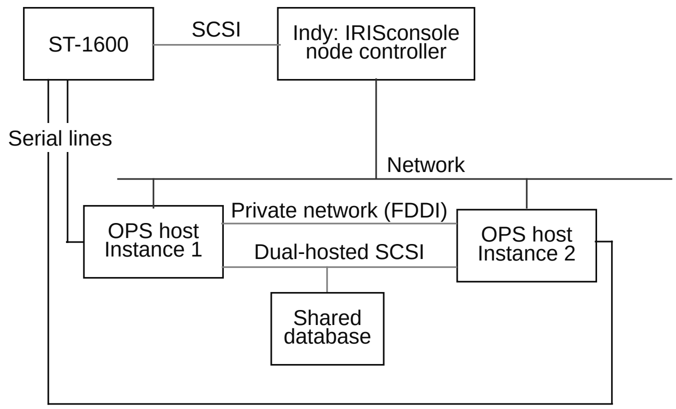

Figure 1-1 OPS Hardware Configuration

Besides the IRISconsole software, the node controller also runs the OPS node control software, *opsnc*. The OPS node controller *opsnc* implements a fail-stop mechanism: in the event of a private network partition (a private network failure that results in the OPS instances being isolated from each other), only one OPS instance is permitted to continue providing service. The other instance is forced to crash and must be restarted by the system administrator. See "Starting OPS" in Chapter 2 for instructions.

Each OPS instance consists of the following software components:

- Oracle RDBMS processes, including PMON, SMON, DBWR, LGWR
- OPS Distributed Lock Manager processes: *dlmmon* / *dlmd*

  OPS allows multiple hosts access to the same shared physical database. The Distributed Lock Manager (DLM) has the functionality that enables sharing. The Distributed Lock Manager program consists of two processes, *dlmd* and *dlmmon*.
- OPS Connection Manager process: *opscm*

  The OPS Connection Manager *opscm* implements a heartbeat protocol across both hosts to detect host and private network failures, monitors local DLM processes to detect lock manager failure, and provides a sync service to coordinate recovery for host failure and reintegration.

Figure 1-2 diagrams OPS software configuration.

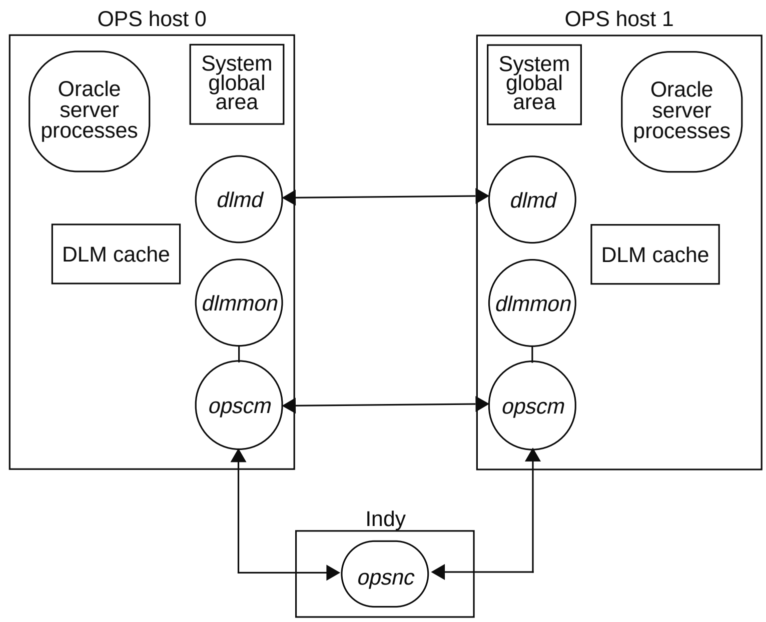

Figure 1-2 OPS Software Configuration

OPS Instances and Domains

An Oracle instance consists of a *system global area* (SGA) and a set of server processes that access the physical database located on disk. The SGA is a section of shared memory accessed by each of the OPS server processes in an instance. In an OPS system, multiple Oracle instances, each with its own SGA, server processes, redo log files, rollback segments, the Distributed Lock Manager (DLM) processes, and the Connection Manager (CM) process access the same physical database. These constitute one *domain*.

DLM domains are numbered starting with 0; DLM instances are numbered 0 and 1. Figure 1-3 diagrams an OPS configuration with a single domain.

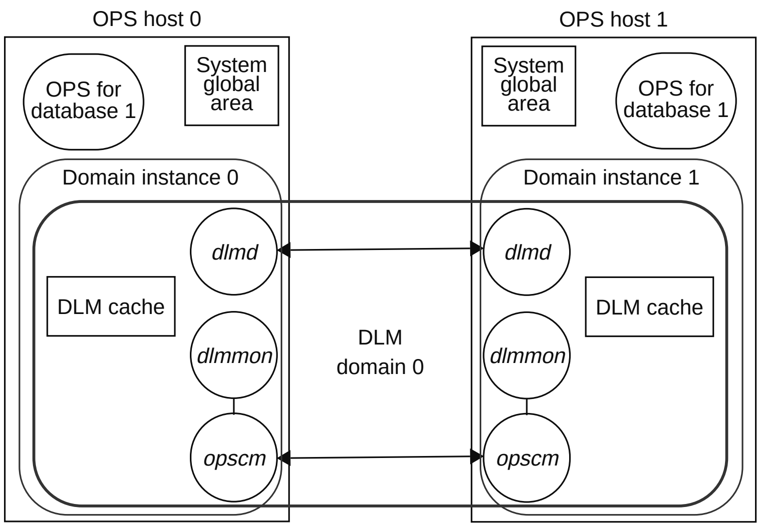

Figure 1-3 One Domain on an OPS Configuration

In the configuration shown in Figure 1-3, the DLM domains are

- 0,0 (DLM domain 0, domain instance 0)
- 0,1 (DLM domain 0, domain instance 1)

Figure 1-4 diagrams an OPS configuration with two domains. In this configuration, the DLM domains are

- 0,0 (DLM domain 0, domain instance 0)
- 0,1 (DLM domain 0, domain instance 1)
- 1,0 (DLM domain 1, domain instance 0)
- 1,1 (DLM domain 1, domain instance 1)

In any of these cases, all Oracle instances that mount the same database must use the same domain.

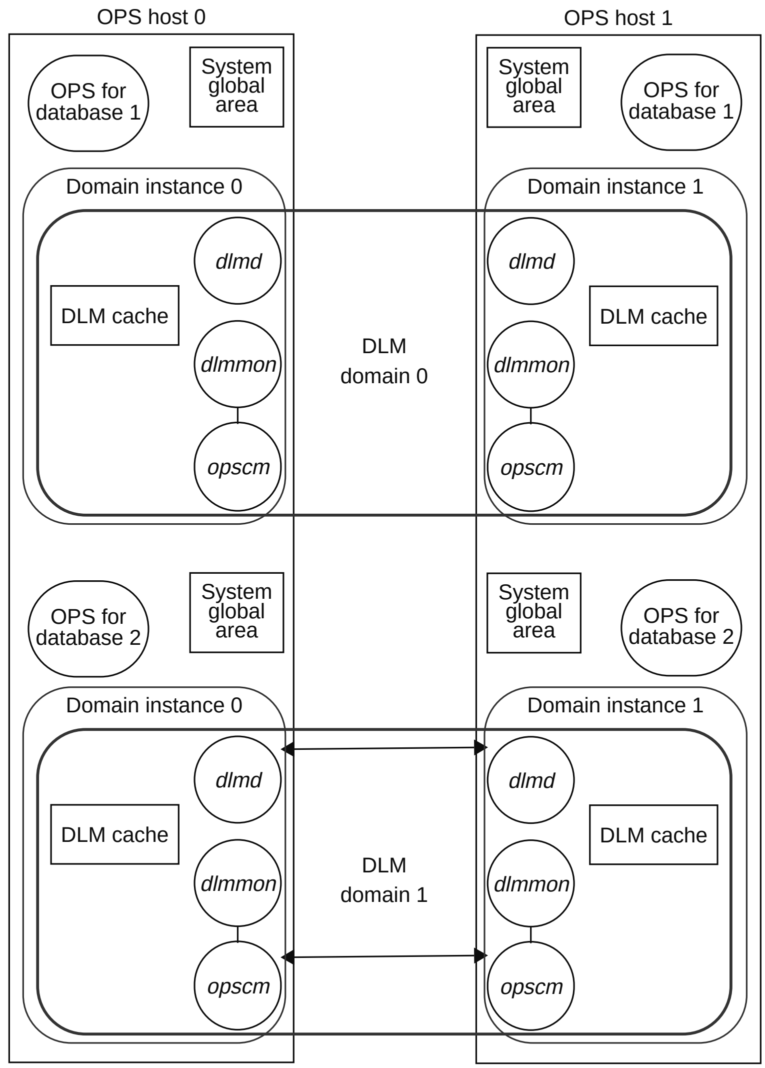

Figure 1-4 OPS Configuration With Two Domains

OPS Architectural Features

Major architectural features of OPS are

- high availability

  High availability is provided at multiple levels:

  - If a node fails, the database is still accessible from the surviving node.
  - If the CHALLENGE RAID storage system is used, RAID-5 provides tolerance to any single point of failure within the RAID.
  - Each redo log file can be mirrored, so that an instance can survive failure of a log file.
- consolidation of database administration, using the Indy workstation as node controller
- high performance

  OPS utilizes the full power of CHALLENGE system memory and its high-speed system bus performance. Operating system enhancements include changes to virtual memory for more efficient multiprocessing, raw I/Os, multiprocess networking, and process scheduling. Besides these enhancements, IRIX already supports real-time scheduling, CPU affinity, and, for the 64-bit OS, CPU partitioning (the ability to steer interrupts to specific CPUs), which are critical for DBMS performance.
- distributed locks

  Row-level locking, the finest level of locking granularity, minimizes the amount of data contention between transactions and maximizes concurrency. Oracle Parallel Server extends this feature by allowing multiple transactions on different nodes to lock and update different rows of any table in the database.

  Row-level locking is independent of the parallel cache manager's use of distributed lock, which is used to keep the SGAs consistent with each other. Row-level locking is achieved by Oracle's internal concurrency control architecture. For distributed locking, the parallel cache manager uses a special background process, the LCNO process, which requests locks from the Distributed Lock Manager. The DLM is not used for row-level locking; thus its use is minimized and performance is enhanced.

For more information on Oracle RDBMS and OPS operation, consult its documentation.

OPS Components and How They Work Together

OPS allows Oracle7 instances running on the two nodes to access a common Oracle database. This design allows users on multiple systems seamless access to common data, so that more computing resources are available to all applications that access the same database.

OPS is designed to allow any node in the cluster to be brought down, either voluntarily or involuntarily, without interrupting access to the database from the other nodes. If an Oracle instance or a system node fails, users from the failed node can migrate to the other running node and reconnect to the database.

Table 1-1 outlines the required and optional software for OPS.

Table 1-1 OPS Hardware and Software

Hardware	Required Software	Optional Software
Two CHALLENGE DM, L, or XL servers (OPS hosts)	IRIX 5.3 with XFS OPS software for Silicon Graphics CHALLENGE RAID software Oracle Corp. RDBMS Oracle Corp. OPS software	Performance Co-Pilot (PCP) IRIXpro IRIX NetWorker 4.1 Database Accelerator (DBA)
CHALLENGE RAID storage system with two storage-control processors (SPs) and at least five disk modules		
IRISconsole: Indy workstation, with Silicon Graphics ST-1600 serial port multiplexer, standard Silicon Graphics CD-ROM, and cables	IRIX 5.3 with XFS Silicon Graphics OPS software *opsnc* IRISconsole software (GUI)	
Optional Vault storage enclosures (Vault M, L, or XL, respectively)		

Figure 1-5 diagrams an example OPS installation with storage systems.

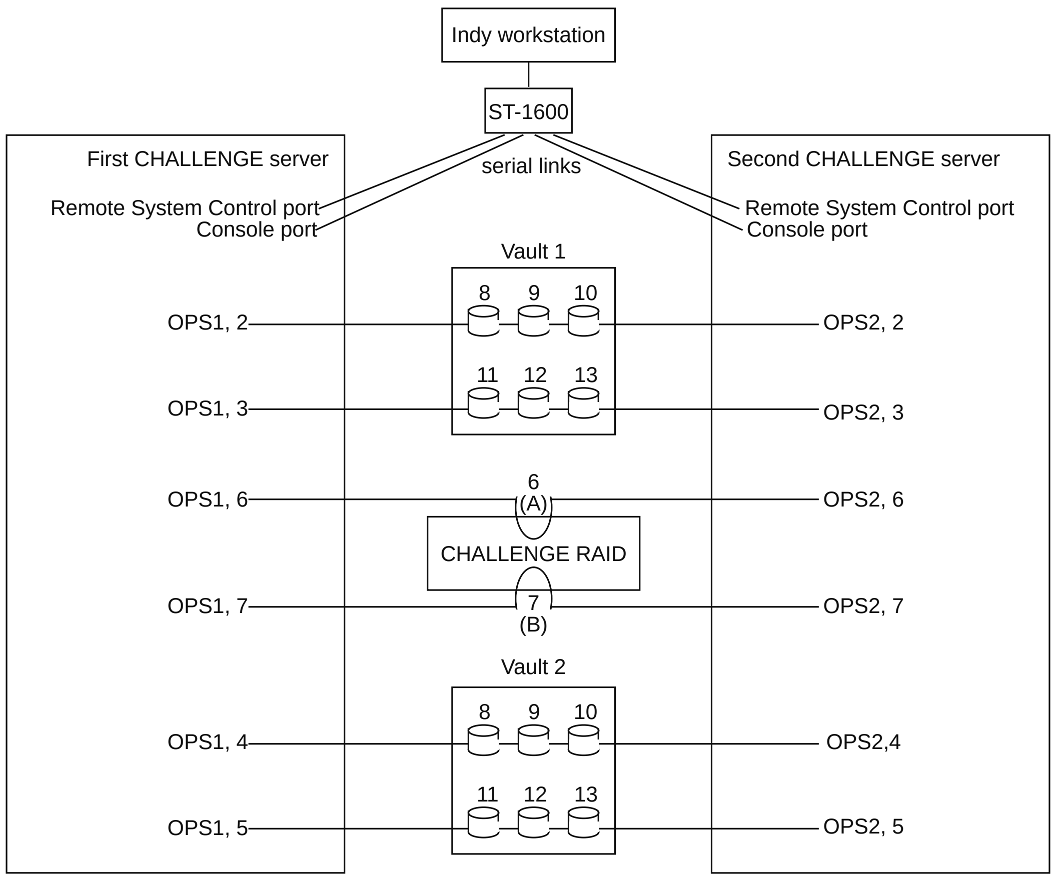

Figure 1-5 Example OPS Site With Two Vaults and One CHALLENGE RAID Storage System

The rest of this section describes specific components of OPS:

- IRISconsole
- CHALLENGE RAID storage system
- XFS filesystem
- Database Accelerator (DBA)

IRISconsole

IRISconsole runs on a 24-bit Indy workstation with a minimum of 32 MB of memory and 20-inch display. For OPS, it is made up of the following:

- IRISconsole software, including a graphical user interface, running under IRIX 5.3 with XFS
- an IRISconsole ST-1600 multiplexer, including cabling connecting the Indy workstation to the ST-1600
- a pair of serial cables included in the IRISconsole package, plus one additional pair, for connecting the two OPS hosts to the ST-1600

The IRISconsole software monitors each OPS host (node) through the host's Remote System Control and console ports via serial connection to the ST-1600 serial port server (multiplexer). If a node fails, IRISconsole can automatically start procedures defined by the OPS system administrator in addition to the failover procedures provided for in the OPS software.

Note: For full OPS and IRISconsole functionality, the Remote System Control and System Console ports on the CHALLENGE DM, L, or XL must be cabled to ports on the ST-1600.

The IRISconsole software enables the administrator to

- display, view, or take control of the console of an OPS host (or other attached system)
- view real-time graphs of hardware operating statistics of an OPS host, such as voltage, operating temperature, and blower speeds; save the graphs as files and display them
- set a threshold for operating statistics so that an alarm is activated when the threshold is reached and various activities can be triggered
- view console activity logs and other system reports

For complete information on the IRISconsole, see the documentation:

- *IRISconsole Administrator's Guide* (007-2872-00x)
- *IRISconsole ST-1600 Multiplexer Installation Guide* (007-2839-00x)

CHALLENGE RAID Storage System

The CHALLENGE RAID (Redundant Array of Inexpensive Disks) storage system provides a compact, high-capacity, high-availability source of disk storage for OPS in the form of multiple disk drive modules that you can replace when the storage system is powered on (hot-replaceable). Each CHALLENGE RAID storage system supports from five to twenty disk modules in groups of five.

The CHALLENGE RAID storage system supports RAID level 5: a group of five disk modules is bound together into a logical unit (LUN). A RAID-5 group maintains parity data that lets the disk group survive a disk module failure without losing data. In addition, in a CHALLENGE RAID storage system configured for OPS, the RAID-5 group can survive a single SCSI-2 internal bus failure, because each disk module in the group is bound on an independent SCSI-2 internal bus.

Through the storage-control processors (SPs), the SCSI-2 bus is split into five internal fast / narrow SCSI buses—A, B, C, D, and E—that connect the slots for the disk modules. For example, internal bus A connects the modules in slots A0, A1, A2, and A3, in that order. Figure 1-6 diagrams this configuration.

For OPS, the CHALLENGE RAID storage system must have two SPs. Each SP controls disk modules in groups of five. The second processor provides a second path to the disk modules as part of the failover strategy of OPS; see Figure 1-6. Each LUN is controlled by one of the SPs. The non-controlling SP takes over a LUN if its controlling SP fails.

In addition, both SPs are required for storage system caching to work: each processor temporarily stores modified data in its memory and writes the data to disk at the most expedient time.

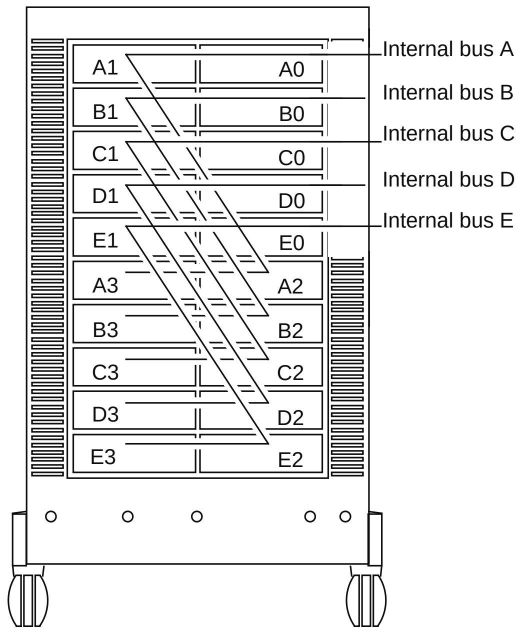

Figure 1-6 SCSI-2 Bus and Internal Buses (Front View)

For complete information on the CHALLENGE RAID storage system, see the *CHALLENGE RAID-5 Owner's Guide* (007-2532-00x).

XFS Filesystem

XFS is a journaled filesystem that allows for extremely fast recovery time of filesystem structures during reboot. Recovery of XFS filesystems is independent of filesystem size. For this reason, XFS is particularly useful for OPS operation.

On a traditional UNIX® filesystem, a full filesystem check takes an amount of time proportional to the size of the filesystem. On XFS, the recovery time is in the seconds, because it is dependent upon system activity level, rather than filesystem size. Using XFS speeds up the time required to bring a failed node back online.

For complete information on the XFS filesystem, see the *Getting Started With xFS Filesystems* (007-2549-00x). This document is viewable in InSight.

DBA (Database Accelerator)

The Database Accelerator (DBA) consists of kernel enhancements designed to boost performance specifically for Oracle. These kernel enhancements can help double the performance of write-intensive benchmarks, such as TPC-[AB], or building very large indexes for real-life applications. The kernel enhancements are as follows:

- *Postwait driver*, a kernel software driver, provides very fast multithreaded synchronization mechanism for Oracle processes. It replaces the standard SVR4 mechanism of semaphore, which is too slow for the high TPS rate.
- *Kernel list I/O*, an OS enhancement, allows the Oracle database writer to flush modified buffers to disks efficiently: a single Oracle database writer can flush at least 2000 buffers per second to disk drives. With only one system call, the database writer can initiate multiple writes to all disk drives in the system. Without this functionality, Oracle database writer performance would have to use shadow processes, thus incurring the overhead of process synchronization; another limitation would be the single-threaded nature of making one system call per disk write.

OPS Performance and Optional Silicon Graphics Software

This section briefly explains how the following optional Silicon Graphics software products can enhance OPS performance:

- PCP (Performance Co-Pilot)
- IRIXpro
- IRIX NetWorker 4.1.1

PCP (Performance Co-Pilot)

The Performance Co Pilot (PCP) provides a suite of tools for performance monitoring and performance management services across the spectrum of performance domains—hardware platforms, the operating systems, the DBMS, and the applications.

PCP runs in a client/server configuration: PCP agents (clients) monitor domains and send information to the PCP server, which graphically displays the information on the workstation. PCP can be used to monitor Oracle and system activity on both nodes in the OPS cluster.

IRIXpro

IRIXpro is a suite of tools for the professional systems administrator. Applications included in IRIXpro are

- Propel: software environment distribution and file management
- Provision: distributed system monitoring
- Problema: request desk
- Proclaim: network configuration server

In particular, Propel can be used to transfer software from one OPS host to the other.

IRIX NetWorker 4.1.1

IRIX NetWorker reliably protects files against loss across an entire network of systems. NetWorker saves valuable administrator time by speeding and simplifying daily backup operations. As NetWorker backs up data, it creates a database of the saved data, making it easy to locate a file for recovery. Furthermore, as the network and number of files expand, NetWorker has the capacity and performance to handle the load.

The IRIX NetWorker 4.1.1 OS includes extended support for autochangers (jukeboxes and tape libraries), and archiving and retrieval capability. Its ability to back up raw files particularly makes it particularly suitable for use with OPS, since all Oracle files are XLV raw devices.

For complete information on NetWorker 4.1.1, see the documentation:

- *IRIX NetWorker Administrator's Guide* (007-1458-0x0)
- *IRIX NetWorker User's Guide* (007-2304-00x)

These documents are viewable in Insight.

Chapter 2

Using the OPS System

This chapter explains

- configuring OPS
- starting OPS
- halting the OPS system

Configuring OPS

OPS configuration files that are edited or created when your OPS system is set up are

- */usr/opscm/conf/opsconf* on each server
- */usr/opscm/conf/sidconf* on each server
- */usr/opsnc/conf/ncconf* on the Indy

Each process is explained separately in this section.

Editing opsconf

To check the contents of the *opsconf* configuration file, follow these steps:

1. Specify the Indy workstation used for IRISconsole: open */usr/opscm/conf/opsconf* on each server. Check to see that the line

   `CLUSTER 1 test 2` *ichostname*

   has the correct system name of the Indy workstation used for IRISconsole.

2. Define the DLM domain(s) and DLM instances, as explained "OPS Instances and Domains" in Chapter 1. Check the lines under

```
#NODE dom inst ndname ndaddress      cmsvc   apsvc   wt
```

have accurate information on your OPS servers. For example:

```
NODE  0   0    host1  150.166.42.37  opscm   opsdlm  1
NODE  0   1    host2  150.166.42.38  opscm   opsdlm  1
```

3. Save and close */usr/opscm/conf/opsconf*.
4. Edit the file */etc/services*; add three services:

```
opscm       newnumber1/tcp
opsdlm      newnumber2/tcp
opsnc       newnumber3/tcp
```

where *newnumbern* is a number not used elsewhere at this site. For example:

```
opscm       7018/tcp
opsdlm      7019/tcp
opsnc       7020/tcp
```

5. Save and exit the file.

Editing sidconf

The *sidconf* file maps each Oracle instance (sid, or Oracle system ID) to a DLM domain-instance pair. To check the *sidconf* configuration file, follow these steps:

1. Determine the sid of the Oracle database for each instance.
2. In */usr/opscm/conf/sidconf*, check that the lines

`MAP` *sid0 domainnumber instancenumber*
`MAP` *sid1 domainnumber instancenumber*

contain accurate information for the servers. For example:

```
MAP finance1 0 0
MAP finance2 0 1
```

3. Save and close */usr/opscm/conf/sidconf*.
4. If necessary, change permissions on this file so that it can be read by all.

Editing ncconf

To edit the *ncconf* configuration file, follow these steps:

1. Note the ports on the ST-1600 to which the Remote System Control server ports are cabled.
2. In an Indy node controller IRIX window, open the */usr/opsnc/conf/ncconf* file.
3. Check that the entries under

   ```
   #nodename   ttyname
   ```

   contain accurate information for the servers. For example:

   ```
   ops1   /dev/ttyf031
   ops2   /dev/ttyf033
   ```

   The last digit in each *tty* entry should be the ST-1600 port into which the Remote System Control for each server is cabled. The numbers 1 through 16 on the ST-1600 correspond to 0 through f in the tty entries. In the example above, the Remote System Control ports are cabled to ports 2 and 4 on the ST-1600.
4. Save and close */usr/opsnc/conf/ncconf*.

Starting OPS

This section explains

- starting the node controller software (*opscm*) on the Indy workstation
- starting the Connection Manager software (opscm) on the CHALLENGE servers
- starting OPS automatically
- starting a node for single-host operation

Starting the Node Controller Software on the Indy Workstation

To restart the OPS software for normal dual-host operation on both hosts and the Indy node controller, follow these steps:

1. As root on the Indy workstation, start the OPS node controller software by typing

   /usr/opsnc/bin/opsnc

2. Start IRISconsole by typing

   /usr/sbin/ic

3. In the IRISconsole site window, click the *Get Console* button for each OPS host.

Starting the Connection Manager Software on the CHALLENGE Servers

Follow these steps:

1. Check for the presence of the CM lock file in */tmp*. This filename has the format *.nn*, where each number stands for the DLM domain and instance, for example, *.00*. If this file exists, delete it.

2. Check to see if the CM is already running by typing

   ps -ef | grep opscm

   If it is running, type the following to kill it:

   killall -TERM opscm

3. Check to see if the DLM is already running by typing

   ps -ef | grep dlm

   If it is running, type the following to kill its processes:

   ```
   killall -TERM dlmmon
   killall -TERM dlmd
   ```

4. Run *ipcs* to determine the shared memory and semaphores used on the host. The following is an example output:

```
IPC status from /dev/kmem as of Thu May 18 14:31:22 1995
T     ID     KEY        MODE        OWNER    GROUP
Message Queues:
Shared Memory:
m      0 0x53637444 --rw-r--r--     root      sys
m    301 0x000022bb --rw-rw----   oracle      dba
m   2202 0x0c33b7c9 --rw-r-----   oracle      dba
Semaphores:
s   2200 0x00000000 --ra-r-----   oracle      dba
```

5. If Oracle or DLM is using any shared memory segments or semaphores, save them to another location if you need them for debugging a DLM or Oracle crash; otherwise, delete them with *ipcrm*. In the example in step 4, you would use

```
ipcrm -m 301 -m 2202 -s 2200
```

6. For each host, create a startup script containing the following lines:

```
#!/sbin/sh

ORACLE_HOME=/usr/people/oracle
ORACLE_SID=sidname
LKDOM=0
LKINST=0
USER=oracle
GROUP=dba

export ORACLE_HOME ORACLE_SID LKDOM LKINST

/usr/opscm/bin/opscm
$ORACLE_HOME/bin/lkmgrd -u $USER -g $GROUP
```

   In each script, make sure that the values for `LKDOM=` and `LKINST=` are accurate for the domain and instance on that host. Theses values must match those in */usr/opscm/conf/sidconf*, as explained in “Editing sidconf,” earlier in this chapter.

7. As root, tun the startup script on each host to bring up OPS.

8. In each host console window, start the Oracle database.

Starting OPS Automatically

To enable *opscm* and the DLM to start automatically at boot time, follow these steps:

1. Edit the */etc/init.d/opscmgr* script. This script is similar to the startup script created in "Starting the Connection Manager Software on the CHALLENGE Servers," earlier in this section.
2. Run this command as root:

   `chkconfig -f opscm on`

Starting a Node for Single-Host Operation

To start a node for single-host operation, run *opscm* with the -F option:

`opscm -F`

Caution: Do not use the -F option for normal OPS dual-host operation.

Halting the OPS System

To halt the OPS system, follow these steps:

1. Back up Oracle database information as needed.
2. On each OPS host, halt Oracle database operation.
3. Type the following to terminate CM gracefully:

   `killall -TERM opscm`

The log and control files that are relevant for OPS failures are

- node control and CM log information are in *syslog*
- DLM log: stored in */var/tmp/dlm/*

Appendix A

Troubleshooting

This appendix explains

- checking OPS host and CHALLENGE RAID SCSI IDs
- checking disk partitions
- replacing a CHALLENGE RAID disk module

Checking OPS Host and CHALLENGE RAID SCSI IDs

This section explains how to make sure that the two OPS hosts (or other SCSI devices on the same bus) do not have the same SCSI ID number. Follow these steps:

1. On the first OPS host, enter the command *nvram*. This command's output should contain the line

   `scsihostid=`

   This output means that the SCSI host ID for this OPS host is 0.

2. Repeat step 1 in the second window, for the second OPS host. The output should contain the line

   `scsihostid=2`

   This output means that the SCSI host ID for the second OPS host is 2.

3. If the host SCSI IDs are not 0 and 2, respectively, reboot and start the System Maintenance menu; choose item 5, the Command Monitor.

4. In the Command Monitor, set the SCSI ID for this system:

   `setenv scsihostid 2`

5. To verify that the SCSI ID was set, enter

   `printenv`

The output should include the line `scsihostid=2`.

6. Exit the System Maintenance menu and restart the system.
7. Enter the command *hinv* on each server and compare the output.

   Caution: No SCSI device on any bus on the second server (OPS2) should have SCSI ID 2.

Table A-1 charts SCSI ID switch settings for the CHALLENGE RAID storage system. Note that the settings do not conform to frequently used numbering schemes.

Table A-1 CHALLENGE RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number				Comment
	ID 0	ID 1	ID 2	ID 3	
1	Off	On	On	On	
2	On	Off	On	On	Do not use
3	Off	Off	On	On	
4	On	On	Off	On	
5	Off	On	Off	On	
6	On	Off	Off	On	Shown in Figure A-1 as A
7	Off	Off	Off	On	Shown in Figure A-1 as B
8	On	On	On	Off	
9	Off	On	On	Off	
10	On	Off	On	Off	
11	Off	Off	On	Off	
12	On	On	Off	Off	
13	Off	On	Off	Off	
14	On	Off	Off	Off	
15	Off	Off	Off	Off	

Figure A-1 shows a typical OPS cabling scheme.

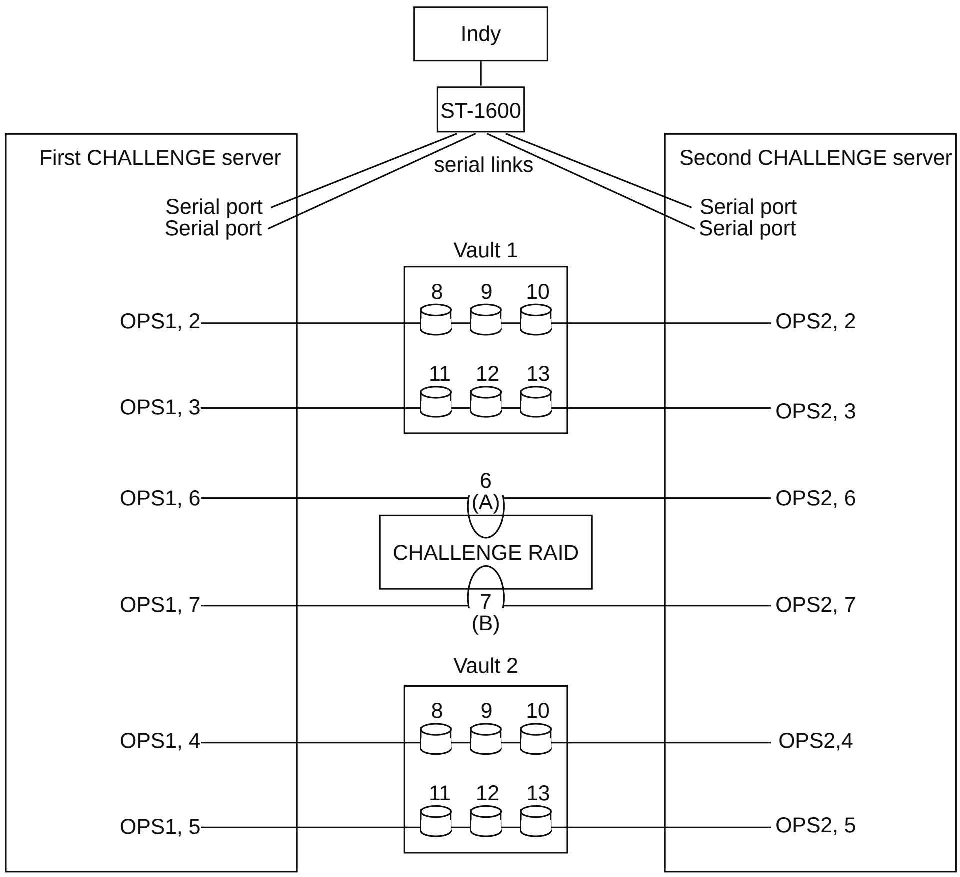

Figure A-1 Example OPS Cabling Scheme

Checking Disk Partitions

Never create disk partitions 8, 9, and 10 manually. If partition 9 on the drive has been put on the drive at Silicon Graphics or during the installation process, do not reuse it or the sectors that it spans. Thus, 13 usable partitions per LUN remain: 0, 1, 2, 3, 4, 5, 6, 7, 11, 12, 13, 14, and 15.

Replacing a CHALLENGE RAID Disk Module

To replace a failed disk module in the CHALLENGE RAID, follow instructions in the *CHALLENGE RAID-5 Owner's Guide*.

Use only CHALLENGE RAID disk modules as replacements; only they contain the correct device firmware. The replacement 2 GB drive part number is 9410113; the replacement 4.3 GB drive part number is 9410114. Other disk modules, even those from other Silicon Graphics equipment, will not work. Do not mix disk modules of different capacities within one array.

After replacing a failed disk module, update the firmware on the CHALLENGE RAID storage-control processor. Type as root:

raid5 -d *device* **firmware /usr/raid5/flare7.67.0.prom.7.99.bin**

Caution: You must use this command every time you replace a disk module.

The image in the file given in the command contains microcode that runs on the storage-control processor and possibly also a microcode image destined for the storage-control processor PROM, which runs the power-on diagnostics.

This command has no output.

We'd Like to Hear From You

As a user of Silicon Graphics documentation, your comments are important to us. They help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics to comment on:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please include the title and part number of the document you are commenting on. The part number for this document is 007-2763-001.

Thank you!

Three Ways to Reach Us

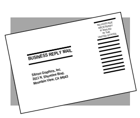

The **postcard** opposite this page has space for your comments. Write your comments on the postage-paid card for your country, then detach and mail it. If your country is not listed, either use the international card and apply the necessary postage or use electronic mail or FAX for your reply.

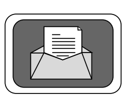

If **electronic mail** is available to you, write your comments in an e-mail message and mail it to either of these addresses:

- If you are on the Internet, use this address: techpubs@sgi.com
- For UUCP mail, use this address through any backbone site: *[your_site]*!sgi!techpubs

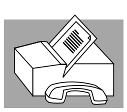

You can forward your comments (or annotated copies of manual pages) to Technical Publications at this **FAX** number:

415 965-0964