

October 1999  
ECG517/1099

Prepared by Customer Relationship  
Management and Business  
Intelligence Group

Compaq Computer Corporation

**Contents**

<b>Introduction</b> .....	3
<b>TPC Benchmark C (TPC-C)</b> .....	3
<b>Oracle8i Server Overview</b> .....	3
<b>Key Features of Oracle8i Server</b> .....	10
<b>Tuning Goals</b> .....	11
<b>Basic Tuning</b> .....	13
<b>I/O Tuning</b> .....	24
<b>Processor Scaling</b> .....	32
<b>Memory Tuning</b> .....	32
<b>Updating InitSID.ora</b> .....	37
<b>Network Tuning</b> .....	38
<b>System Management</b> .....	39
<b>Conclusion</b> .....	39

# OLTP Configuration and Tuning of Oracle8i for Windows NT on Compaq Servers

**Abstract:** This Technical Guide provides guidelines for configuring and tuning the performance of Oracle8i Server and Microsoft Windows NT on Compaq servers in an On-Line Transaction Processing (OLTP) environment.

The system tested by Compaq represents a single node, Transaction Processing Performance Council (TPC) Benchmark C on the *Compaq ProLiant* family of servers.

## Notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements.

Compaq, Deskpro, Compaq Insight Manager, Systempro, Systempro/LT, ProLiant, ROMPaq, QVision, SmartStart, NetFlex, QuickFind, PaqFax, and Prosignia are registered with the United States Patent and Trademark Office.

ActiveAnswers, Netelligent, Systempro/XL, SoftPaq, QuickBlank, QuickLock are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows and Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

Intel, Pentium and Xeon are trademarks and/or registered trademarks of Intel Corporation.

Oracle is the registered trademark of Oracle Corporation.

TPC Benchmark C material from *TPC Benchmark C - Standard Specification, Revision 3.5, August 25, 1999* is under the copyright of the Transaction Processing Performance Council and was copied by permission.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1999 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

OLTP Configuration and Tuning of Oracle8i for Windows NT on Compaq Servers  
Technical Guide prepared by Customer Relationship Management and Business Intelligence Group.

First Edition (October 1999)  
Document Number ECG517/1099

## Introduction

This Technical Guide contains information to help you optimize the performance of your On-Line Transaction Processing (OLTP) system. The document lists tuning goals and provides recommendations for basic tuning and for tuning your I/O performance, processors, and memory. An introduction to the Transaction Processing Performance Council (TPC) Benchmark C and an overview of the key features and architecture of Oracle8i server are provided.

When tuning an OLTP system, it is important to realize that optimizing parameters for one application does not necessarily improve performance for other applications. Tuning is an iterative process; requirements can change as your workload changes.

## TPC Benchmark C (TPC-C)

The TPC-C benchmark is a workload that consists of read-only and update-intensive transactions simulating the activities found in a complex OLTP application environment. TPC-C exercises the simultaneous execution of multiple transaction types spanning a broad spectrum of complexity. These transaction types include:

- Online and deferred transaction execution modes
- Multiple on-line terminal sessions
- Moderate system and application execution time
- Significant disk input/output
- Transaction integrity
- Non-uniform distribution of data access through primary and secondary keys
- Databases consisting of many tables with a wide variety of sizes, attributes, and relationships
- Contention on data access and update

While TPC-C offers a rich environment that emulates many OLTP applications, this benchmark does not reflect the entire range of OLTP requirements. Compaq does not recommend extrapolation to any other environment.

---

**Note:** For more information on TPC-C, contact the Transaction Processing Performance Council on <http://www.tpc.org/cspec.html> or phone 408-295-8894.

---

## Oracle8i Server Overview

The information in this Technical Guide relates to Oracle8i Server Enterprise Edition 8.1.5 for Windows NT 4.0. Because information on the generic tuning of Oracle8i Server is already available, this document focuses on specific tuning suggestions for OLTP on Compaq servers and Windows NT. Wherever possible, this document references additional tuning documentation.

This section provides general information on the architecture of Oracle8i Server.

Oracle8i Server for Windows NT is a 32-bit application that is implemented on Windows NT as a single process, multi-threaded architecture. Each Oracle8i Server instance consists of:

- A single Windows NT process with multiple Oracle8i server threads
- A System Global Area (SGA)
- A Process Global Area (PGA)
- A redo log
- Control and configuration files

This section provides basic information on the Oracle8i Server architecture. A more in-depth discussion is available in the *Oracle8i Server Concepts Manual*.

## Threads

---

**Note:** Under the multithreaded architecture of Windows NT, an Oracle8i Server instance is considered to be a Windows NT process while Oracle8i Server processes<sup>1</sup> are considered to be Windows NT threads. Therefore, for the purpose of this Technical Guide, Compaq refers to an Oracle8i Server instance as a process and an Oracle8i Server process as a thread.

---

An Oracle process consists of **server threads** (handling requests from user connections), **background threads** (consolidating and handling functions that would otherwise be carried out by multiple management tasks running for each user connection), and **shadow threads**<sup>2</sup> (reflecting the number of user connections to the database).

The number and type of threads associated with an Oracle process depend on the number of Oracle8i Server options selected (background threads) and the number of user connections (shadow threads).

### Server Threads

There are two server threads present for each Oracle process, regardless of the options selected and the number of user connections. Thread 0, the Oracle process main thread, acts as a service dispatcher, creating thread 1 to handle the service.

### Background Threads

Each background thread is associated with an Oracle8i Server background process: **PMON**, **DBWn**, **LGWR**, **CKPT**, **SMON**, or **RECO**. In addition, if you enable **ARCHIVELOG** mode, Oracle8i Server creates the **ARCn** background thread. For more information on background threads, see Table 1.

---

<sup>1</sup> Server, background, and shadow processes.

<sup>2</sup> There is one shadow thread for each user connection.

**Table 1. Oracle8i Server Background Threads**

Thread	Mnemonic	Description
Process Monitor	<b>PMON</b>	Process Monitor is responsible for the cleanup of abnormally terminated connections.
Database Writer	<b>DBWn</b>	Database Writer writes modified database blocks to datafiles.
Log Writer	<b>LGWR</b>	Log Writer writes the redo log entries to logfile.
Checkpoint Process	<b>CKPT</b>	Checkpoint Process signals <b>DBWn</b> to perform updates on all database data and control files. If <b>CKPT</b> is not present, <b>LGWR</b> signals <b>DBWn</b> to perform the updates.
System Monitor	<b>SMON</b>	System Monitor performs instance recovery and cleanup.
Recovery Process	<b>RECO</b>	Recovery Process resolves failures associated with the distributed database option.
Archival Process	<b>ARCn</b>	If you enabled <b>ARCHIVELOG</b> mode, Archival Process copies a full redo log file to the archive device.

**Note:** Since Compaq disk subsystems provide asynchronous input/output, Oracle8i Server requires only one Database Writer (**DBWn**) thread and one Log Writer (**LGWR**) thread.

## Shadow Threads

A shadow thread is a separate, dedicated server thread that acts on behalf of a particular user. Oracle8i Server creates a shadow thread for every user that connects to the database and uses this thread to perform database requests made by the user.

**Note:** Oracle8i Server supports a Multi-Threaded Server (MTS) environment. An MTS server thread can service requests from multiple clients.

When performing a database request, the shadow thread first checks to see if the required information is present in the System Global Area (SGA); if the information is not present, the shadow thread reads the appropriate data from datafiles into the SGA. When writing information to the database, the shadow thread writes into the SGA only, **DBWn** later writes these “dirty” data to disk.

Shadow threads can operate in parallel. Parallel execution begins when a shadow thread executes a SQL statement containing operations that can be performed in parallel. This shadow thread becomes the **query coordinator**, which dispatches the execution of the SQL statement to several parallel server threads, coordinates the results, and sends the results to the user. The number of parallel server threads assigned to a single operation is known as the “degree of parallelism.”

The maximum number of shadow threads (user connections) depends on the size and configuration of system memory available to Oracle8i Server.

## Monitoring Process and Thread Activity

This section discusses using Windows NT to monitor process and thread activity.

### Process Status

Windows represents the Oracle8i Server process as a service named OracleServiceSSSS (where SSSS is the system ID) associated with the executable **ORACLE.EXE**. Each Oracle8i Server instance (by which an Oracle8i database is created and accessed) is associated with a specific Windows NT service.

Use Windows NT Control Panel→Services to display the status of services as shown in Figure 1 where “OracleServiceORCL” represents Oracle8i Server instance ORCL.

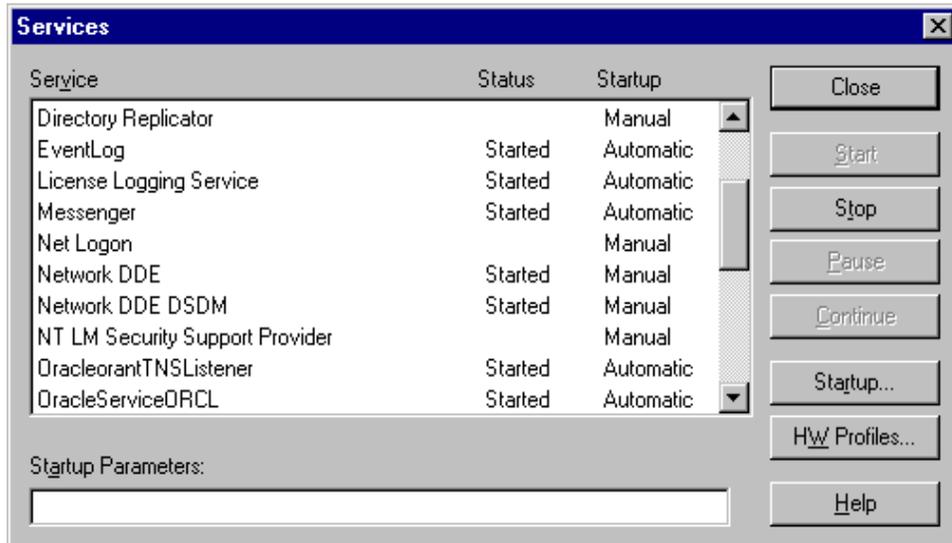


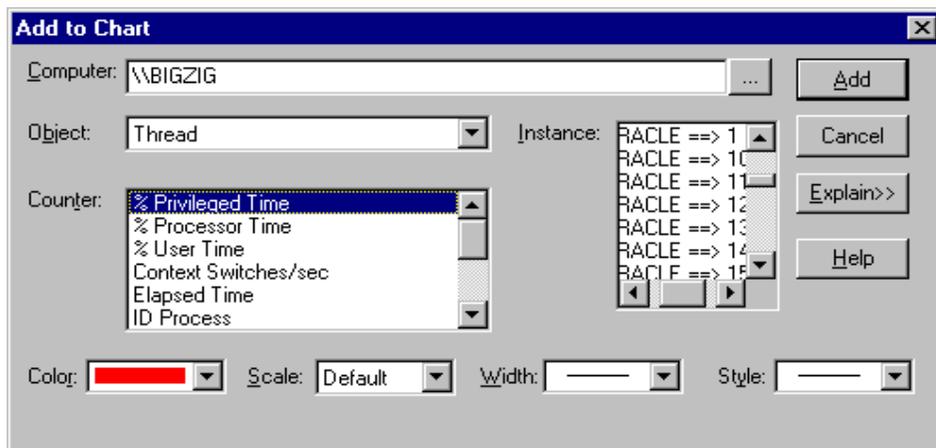
Figure 1. The Services screen showing an Oracle8i Server instance

## Thread Activity

The Windows NT Performance Monitor allows you to monitor Oracle8i Server thread activity.

Select Edit→Add to Chart to obtain the Performance Monitor screen illustrated in Figure 2.

Select “Thread” from the Object list and use this screen to associate Performance Monitor counters with particular thread instances. For example, selecting “ORACLE ==> 3” allows you to add counters for Oracle ID Thread 3.

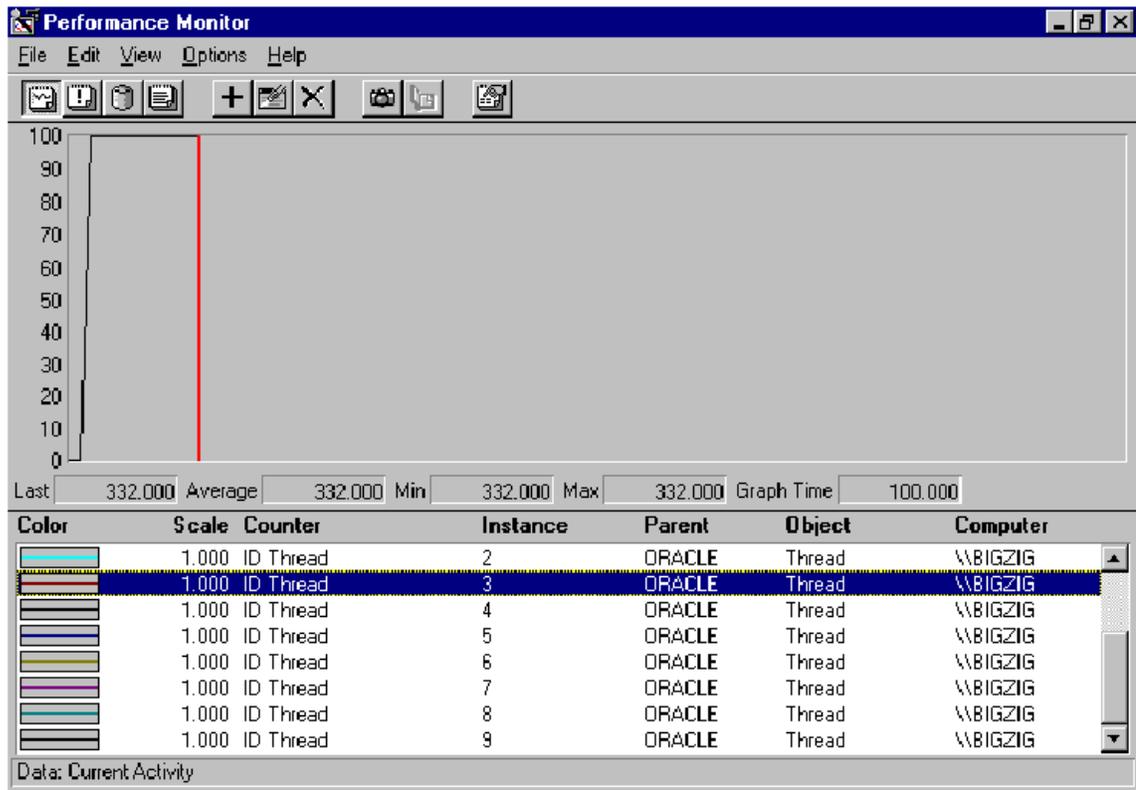


**Figure 2. The Performance Monitor thread object screen**

You can identify a particular Oracle process from the value of the ID Thread Counter. From the screen illustrated in Figure 2, follow these steps:

1. Select “ID Thread” from the Counter list.
2. Select all Oracle instances from Instance list.

The resulting screen provides information on all Oracle threads. See Figure 3.



**Figure 3. Performance Monitor showing the Oracle8i process and threads**

**Note:** Performance Monitor separates Oracle8i Server instances on the chart, log, and report but identifies each instance as “ORACLE,” making it difficult to monitor multiple instances.

In Figure 3, the user has selected ID Thread 3. The value of the selected ID Thread Counter is 332 (as displayed in the “Last,” “Average,” “Min,” and “Max” boxes). To determine which background process corresponds to a particular ID Thread Counter value, you can use this SQL statement:

```
SELECT bg.paddr, bg.name, bg.description, p.spid, p.pid
FROM v_$bgprocess bg, v_$process p
WHERE bg.paddr = p.addr;
```

This statement displays all background processes running on the Oracle8i Server instance. See Table 2 for an example of the information displayed.

**Table 2. Sample Background Process Display**

<b>p.addr</b>	<b>bg.name</b>	<b>bg.description</b>	<b>p.spid</b>	<b>p.pid</b>
1C169E2C	<b>PMON</b>	process cleanup	332	2
1C16A11C	<b>DBW0</b>	db writer process 0	309	3
1C16A40C	<b>LGWR</b>	Redo etc.	288	4
1C16A6FC	<b>CKPT</b>	checkpoint	249	5
1C16A9EC	<b>SMON</b>	System Monitor Process	277	6

The **p.spid** values in the display correspond to values specified in the ID Thread Counter; **bg.name** identifies the background process associated with the particular ID Thread Counter value.

From Figure 3, the ID Thread Counter value associated with ID Thread 3 is 332. From Table 2, a **p.spid** value of 332 corresponds to the **PMON** process.

Knowing which background process is associated with a particular ID Thread allows you to monitor Oracle8i Server threads using any of the thread counters in Windows NT Performance Monitor. See Figure 2.

## System Global Area (SGA)

Table 3 lists the contents of the SGA sharable memory construct.

**Table 3. SGA Contents**

<b>Contents</b>	<b>Description</b>
Database Buffers	The Database Buffers contain the most-recently-used database blocks.
Shared Pool	The Shared Pool contains shared SQL areas and the data dictionary cache.
Redo Log Buffers	The Redo Log buffers record all changes made to the database.
Large Pool	The Large Pool is an optional memory area that can provide large memory allocations for Multi-Threaded Server (MTS) session memory, the Oracle8i Server XA interface, the I/O server processes, and Oracle8i Server backup and restore operations.

## Redo Log File

To facilitate the recovery of a particular instance, the Oracle8i Server redo log contains a history of the database transactions committed by each instance. Oracle8i Server requires at least two redo logs and, if desired, you can use more.

When a redo log fills, a **log switch** occurs. After the log switch, Oracle8i Server records all new redo information in the next redo logfile in line. If the system is running in **ARCHIVELOG** mode<sup>3</sup>, Oracle8i Server copies the previous logfile to an **archive logfile**.

With a recent database backup, the redo logfiles, and the archive logfiles, you can recover an Oracle8i Server instance, if needed.

<sup>3</sup> Recommended by Compaq.

## Control and Configuration Files

Oracle8i Server uses control and configuration files to store information on the state and layout of the database and system tunables. For more information on OLTP configuration, see the “Updating **InitSID.ora**” section of this Technical Guide.

---

**Note:** It is very important to backup your control file when backing up the Oracle8i database.

---

## Key Features of Oracle8i Server

Table 4 presents key features of Oracle8i Server 8.1 that can help to improve the performance of your OLTP application. For more information on these features, refer to the Oracle publication: *Getting to Know Oracle8i*.

**Table 4. Features of Oracle8i Server**

Feature	Description
Partitioned Tables and Indexes	Oracle8i Server can divide large tables and indexes into smaller, more manageable segments that are known as <b>partitions</b> . Oracle8i Server can create, manage, backup, restore, and utilize partitions to improve query and Data Manipulation Language (DML) performance.  Oracle8i Server supports range, hash, and composite partitioning.
<b>ROWID</b> Enhancement	Oracle8i Server <b>ROWIDs</b> are 10-byte IDs that contain the data object number, the data block address, and row number of a particular row. <b>ROWIDs</b> can accommodate partitioned tables and tablespace-relative data block addresses.
Index-Organized Tables	Oracle8i Server stores table data in a B*-tree index structure for better performance and reduced storage.
Data Dictionary	The data dictionary includes Oracle8i Server views that provide additional information on partitions, parallel server, and latches/locks. Some views have additional columns defined to include new features and functionality.
Reverse-Key Indexes	A reverse-key index reverses the bytes of each column indexed, except the <b>ROWID</b> , but maintains the column order. This allows Oracle8i Server to distribute insertions across all leaf keys in the index.
SQL*Loader Partitioned Object Support	SQL*Loader partitioned object support facilitates the loading of partitioned objects into the database.  The direct path has been changed to accommodate the mapping of rows to table partitions and to support local and global indexes.  Parallel direct-path includes concurrent loading of an individual partition as well as support for concurrent loading of a partitioned table.
Multiple Buffer Pools	Oracle8i Server supports the configuration of database buffer cache into multiple buffer pools. You can assign schema objects to specific buffer pools.  The <b>KEEP</b> and <b>RECYCLE</b> buffer pools can retain data blocks in memory or eliminate these data blocks from memory when they are no longer needed. The <b>DEFAULT</b> buffer pool contains data blocks from schema objects that are not assigned to the <b>KEEP</b> or <b>RECYCLE</b> buffer pools.

*Continued*

**Table 4. Features of Oracle8i Server** *continued*

Feature	Description
Locally-Managed Tablespaces	<p>The capability to manage tablespaces locally improves reliability and reduces fragmentation by avoiding the recursive space management that can occur in dictionary-managed tablespaces if consuming or releasing space in an extent results in another operation that consumes or releases space in a rollback segment or data dictionary table.</p> <p>The use of locally-managed tablespaces automatically tracks adjacent free space, eliminating the need to coalesce free extents.</p>

## Tuning Goals

There are several factors involved in achieving the best performance from your system, including the optimization of the hardware, the Oracle8i Server, the operating system, and the application software. This Technical Guide focuses on optimizing these system components:

- Hardware
- Oracle8i Server
- Operating system

---

**Note:** It is important to tune database application software to take full advantage of system resources. However, due to the diversity of database applications, tuning application software is beyond the scope of this document. For more information on tuning applications, refer to the *Oracle8i Application Developers Guide*.

---

A well-tuned Oracle8i Server system has these characteristics:

- I/O wait time is minimal.
- Most of the system processor utilization is allocated to shadow threads and not the system processes threads.
- Most of the system processor utilization is in user mode.
- Users see good response times.

This section discusses the characteristics of a well-tuned system; subsequent sections provide guidelines for achieving these goals.

### Waiting on I/O

Minimal waiting on I/O indicates that the system processors always have work to do while there are outstanding I/Os.

---

**Note:** You can use the Windows NT Performance Monitor to monitor I/O performance. Performance Monitor statistics appear in this format: (XXXX:YYYY) where XXXX is an object, such as Processor, and YYYY is an associated counter, such as %Processor Time.

---

If the processor utilization (Processor:%Processor Time) is 100%, then there is no disk I/O contention.

If the %Idle (100 minus (Processor:%Processor Utilization)) is more than 15% to 20% while the system is under load, you should check the number of outstanding I/Os. You can estimate the number of outstanding I/Os by running Performance Monitor and observing the number of queued transactions (Physical Disk:Avg Disk Queue Length) and the completed transactions (Physical Disk:Disk Transfers/sec).

(Physical Disk:Avg Disk Queue Length) - (Physical Disk:Disk Transfers/sec) should be minimal.

---

**Note:** To observe disk activity using Performance Monitor, run the **diskperf** utility to activate disk performance measurement by the system (**diskperf-y**). Since **diskperf** degrades system performance by up to 1%, you should disable **diskperf** (**diskperf-N**) after you have completed disk monitoring to restore maximum performance.

You must restart your system for the **diskperf** option to become effective.

---

## Processor Utilization

Most of the system processor utilization should be allocated to the shadow threads and not the system processes threads. Use Performance Monitor to check the system processor utilization at the Oracle thread level (Thread:% Processor Time) and make sure that the Oracle process and threads receive most of the system processor time.

## User Mode

Most of the system processor utilization should be available for user applications rather than Windows NT kernel applications. Check the system processor utilization by running Performance Monitor and reviewing the percentage of processor time in privileged time<sup>4</sup> and in user time for the Oracle process (Process:% User Time/% Privileged Time). You should maximize the percentage of user time.

## Response Times

Users should see good response times. A system that appears to be well-tuned but has poor response times could have an inefficient statement in the application or excess latencies in the network or I/O subsystems. Since well-tuned database applications are system processor-bound, a faster processor can improve response times.

---

<sup>4</sup> Think of privileged time as operating system overhead such as time spent in the I/O subsystem or in making system calls.

## Basic Tuning

This section contains basic tuning information relating to the installation and configuration of the *Compaq ProLiant* server hardware, operating system, and Oracle8i Server.

Compaq provides guidelines for:

- Tuning your system at installation time
- Installing Oracle8i Server
- Making changes to your Windows NT configuration:
  - Installing the latest fixes and drivers
  - Configuring physical, virtual, and server memory
  - Removing unwanted network protocols
  - Configuring server and workstation bindings
  - Stopping unwanted services
  - Configuring the file system

Later sections of this Technical Guide discuss the tuning of your OLTP system.

## Installation Time

Table 5 provides guidelines for tuning your system at installation time.

**Table 5. System Tuning at Installation Time**

Action	Description
Updating the flashable ROM	<p>Apply the latest <i>Compaq ROMpaq</i> to your <i>ProLiant</i> after configuring your system.</p> <p>Applying the <i>ROMpaq</i> can take some time if your system has a large number of drives configured.</p>
Configuring the system	<p>Use the <i>Compaq SmartStart</i> CD-ROM to configure your system and create your system partition. After creating the system partition, <i>SmartStart</i> copies the system configuration files to your system partition after rebooting. This allows you to run <i>sysConfig</i> upon bootup (select F10 at the prompt) at any later date.</p> <p>To review and edit your system configuration, reboot and run <i>sysConfig</i>. Select System Configuration→Configure Hardware→Review Details→View or Edit Details to review your boot order. Set the order of your boot drive in "Controller Order" to "First".</p> <p>Note that the CD-ROM (Embedded IDE Controller) is always "Enabled as Second Controller".</p>
Configuring the array controllers	<p>Run the Compaq Array Configuration Utility (on the <i>SmartStart</i> CD-ROM) against all controllers to be used in your OLTP system. See Figure 4. "Controller Selection" displays the controller and slot number currently viewed on the screen.</p> <p>Apply these settings:</p> <p><b>Operating System</b> Select Controller Settings→Operating System→Windows NT 3.5x, 4.0, or later.</p> <p><b>Accelerator Ratio</b> The accelerator ratio allocates memory reserved for the accelerator's read ahead/write-posting functionality. The amount of memory for the accelerator and the ratio for the accelerator cache depend on the controller board. For more information, see "Array Controllers" later in this Technical Guide.</p> <p>Use Controller Settings→Accelerator Ratio to review and edit the accelerator ratio.</p> <p><b>Array Accelerator Cache</b> Enable the array accelerator cache to increase controller performance by writing data to RAM on the controller rather than directly to disk. The RAM on the controller is mirrored and battery-protected for increased reliability. Enable the Array Accelerator Cache for both log drives and non-log drives under Logical Drive→Array Accelerator.</p>

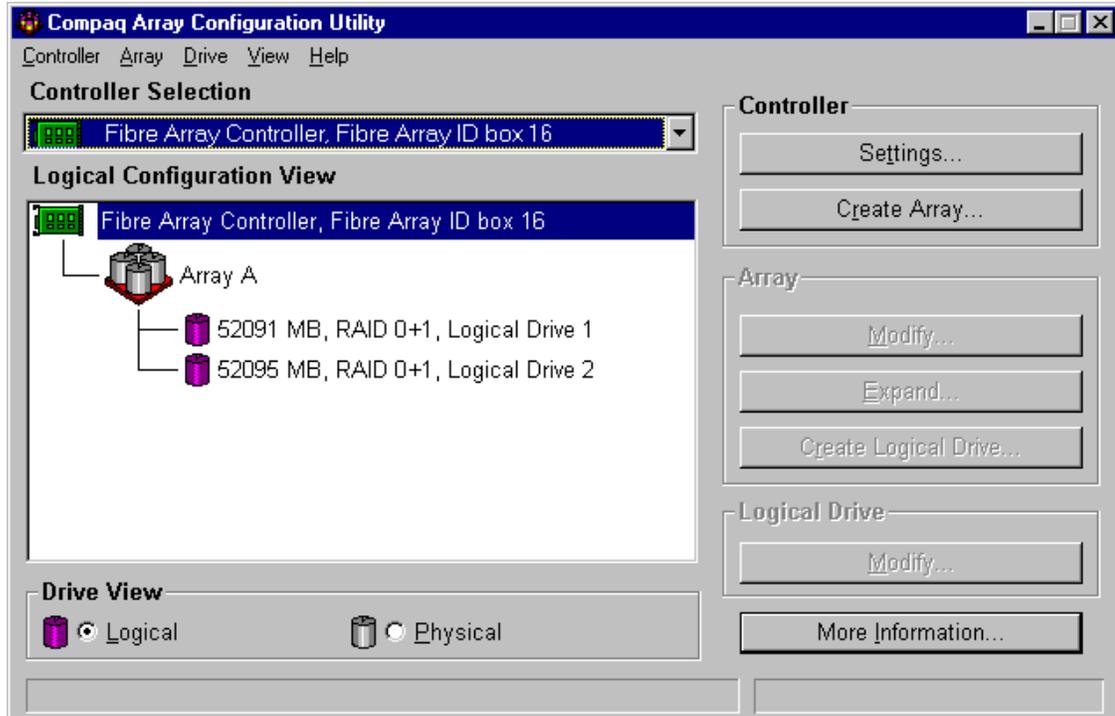


Figure 4. The Compaq Array Configuration Utility screen

## Installing Oracle8i Server

Follow these guidelines when installing Oracle8i Server:

- Install the partition option to enable Oracle8i Server partition functionality.
- Use one of these methods to create an Oracle8i Server instance:
  - Use the Oracle Database Configuration Assistant to create the instance automatically.
  - Use the Oracle Database Instance Manager (ORADIM) utility to create the instance manually and assign the administrator ID and password. If you are recreating an instance, you should first delete the old instance using:

```
oradim -delete -sid SSSS
```

before creating the new instance using:

```
oradim -new -sid SSSS -intpwd PWD -maxusers U
        -startmode manual
```

where **SSSS** is your instance ID, **PWD** is your internal password, and **U** is the maximum number of internal users.

- Use the Windows NT Registry Editor to make any necessary changes to Oracle8i Server registry parameters such as **ORACLE\_SID**, **LOCAL**, and **ORACLE\_HOME**, as well as other Oracle registry parameters. You should make the changes to:  
`\SOFTWARE\ORACLE\ORA(OracleID)` where **OracleID** is the ID number found in the `oracle.key` file in the `oracle_home\bin` directory. See Figure 5.

**Note:** Set the **ORA\_MAX\_ALLOC** registry parameter to 32,768 for Oracle8i Server 8.1.5.0 to specify that 32,768 bytes of guaranteed allocatable memory is available to Oracle8i Server.

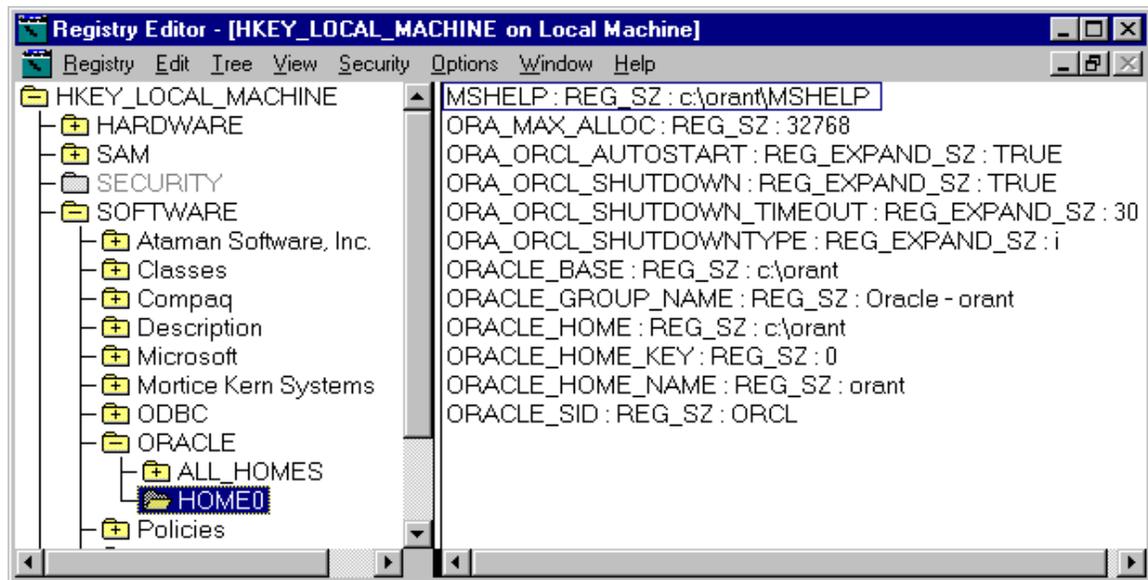


Figure 5. Registry Editor showing Oracle8i Server configuration parameters

## Windows NT Tuning

Prior to or immediately following the installation of Oracle8i Server, Compaq recommends making these changes to Windows NT:

- Applying the latest operating system fixes from Microsoft
- Installing the latest Compaq drivers
- Installing the appropriate Compaq SMART Array controller drivers

### Applying the Latest Fixes

Install the latest version of the Microsoft Windows NT 4.0 Workstation and Server Service Pack to apply the latest operating system fixes from Microsoft.

### Installing the Latest Compaq Drivers

Install the latest version of the Compaq Software Support Diskette (SSD) to apply the latest drivers for Windows NT on Compaq hardware. You can obtain a complete listing of available drivers from: <http://www.compaq.com/support/files/server/softpaqs/winnt/NTSSD400.htm>

## Installing Compaq SMART Array Controller Drivers

You can install miniport or monolithic drivers for the SMART Array controllers:

- **Miniport driver:** Microsoft distributes a default miniport driver, **CPQARRAY2.SYS**<sup>5</sup>, with Windows NT; Compaq distributes drivers and updates on the SSD.

---

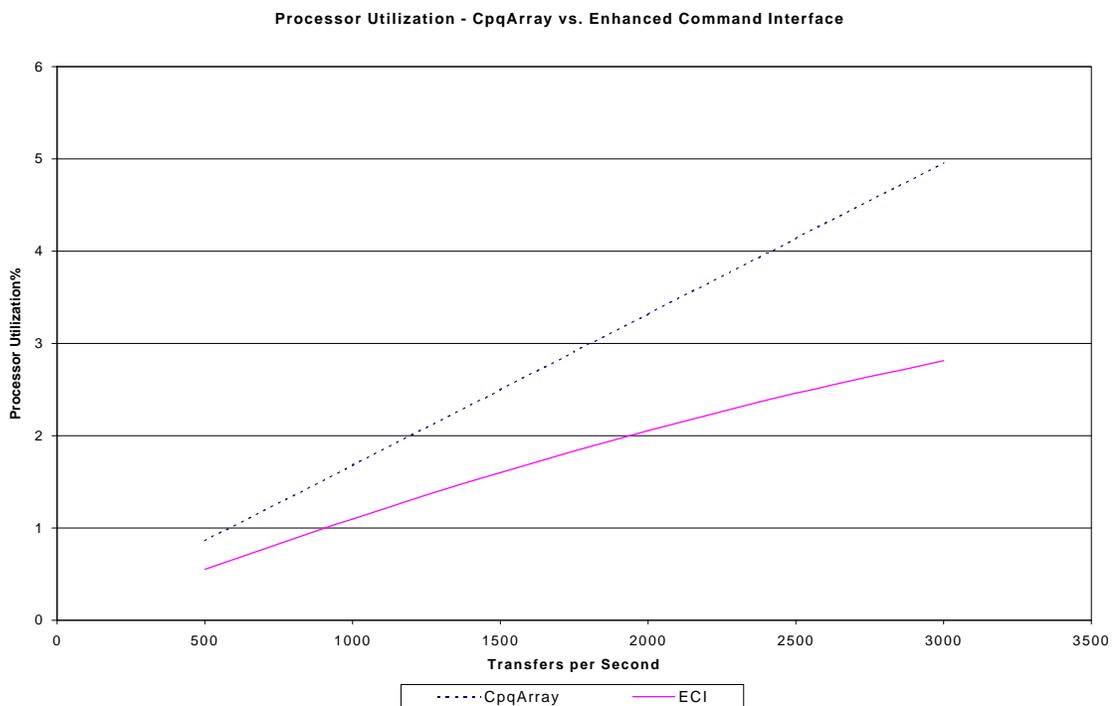
**Notes:** The driver for SMART Array 42xx controller is on SSD only. If you are installing Windows NT on a boot drive controlled by a SMART Array 4200 or 4200ES controller, you must start the Windows NT install from the installation diskettes.

---

- **Monolithic driver:** The Enhanced Command Interface (ECI) driver is a monolithic device driver that, in some cases, can improve performance. The ECI driver, **CPQSMRT2.SYS**, is available from:

<http://www.compaq.com/support/files/server/softpaqs/Rompaq/smart2eci.html>

The ECI driver requires less processor overhead for performing I/O and can decrease overall processor use, allowing the processor(s) to accommodate a heavier workload. If your system is processor-bound, you should consider using the ECI driver. See Figure 6 for a comparison between processor usage with miniport and monolithic drivers.



**Figure 6. Miniport (CPQARRAY) and monolithic (ECI) driver performance**

---

<sup>5</sup> **CPQARRAY.SYS** for the SMART Array 3200 controller, or older.

You should be aware of these restrictions on the use of ECI drivers:

- The ECI driver supports all SMART Array controllers except for Fibre Channel storage solutions.
- These applications do not currently supported the ECI driver:
  - *Compaq Insight Manager* agents for Windows NT and *Compaq Insight Manager*
  - Compaq Standby Recovery Server Option
  - Compaq On-Line Recovery Server Option
  - Compaq Array Configuration Utility (when used on-line under Windows NT)

---

**Note:** If you require any of these applications to support your SMART Array controllers, use the miniport driver.

---

- Although installed through the Control Panel → SCSI Adapters applet, the ECI driver is not a SCSI driver so, rather than being listed under the Devices tab of the applet as when using the Compaq Drive Array driver, the driver is listed under the Drivers tab.

## Configuring Physical Memory

Windows NT Enterprise Edition increases the per-process address limit from 2GB to 3 GB. This benefits applications that run on *Compaq ProLiant* servers with more than 2GB of physical RAM and can take advantage of a larger address space.

If you have more than 2GB of RAM, install the Windows NT 4.0 Enterprise Edition and modify the **boot.ini** file by adding the **/3GB** parameter to the startup line.

---

**Note:** Oracle8i Server 8.1.5 is limited to 3GB of memory. Do not exceed this limit.

---

For example:

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(2)\WINNT
[operating systems]
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Server
Version 4.00" /3GB
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Server
Version 4.00 [VGA mode]" /basevideo /sos/basevideo /sos
```

You must re-boot after modifying the **boot.ini** file.

To take advantage of the increased Windows NT address limit, run **imagecfg** with the **-l** option on the Oracle8i Server executable. This enables Oracle8i Server to use an address space larger than 2GB:

```
imagecfg -l oracle.exe
```

---

**Note:** You can find **imagecfg** in the Support folder of the Windows NT 4.0 Enterprise Edition CD-ROM.

---

## Configuring Virtual Memory

The size of your virtual memory is the size of the physical memory (RAM) plus your swap file size. Ideally, the size of the swap file should be approximately one to one and a half times that of the RAM and should not exceed twice the size of the RAM. In this case, the paging rate may increase during peak times, resulting in decreased performance.

You can set virtual memory under Control Panel→System→Performance→Virtual Memory.

---

**Note:** If you are using amounts of RAM larger than the disk space that can be spared, it may not be practical to set your swap file to one to one and a half times that of your physical memory. When setting your virtual memory and swap file size, be aware that you should reduce the amount of paging and swapping.

---

## Configuring Server Memory

To optimize server memory for Oracle8i Server, use Control Panel→Network→Server to re-configure server memory. Change the relationship<sup>6</sup> between the memory allocated to network connections and the memory allocated to applications running on the server by selecting "Maximize Throughput for Network Applications." See Figure 7.

**Note:** Oracle8i Server performs its own memory management for caching file and network I/O.

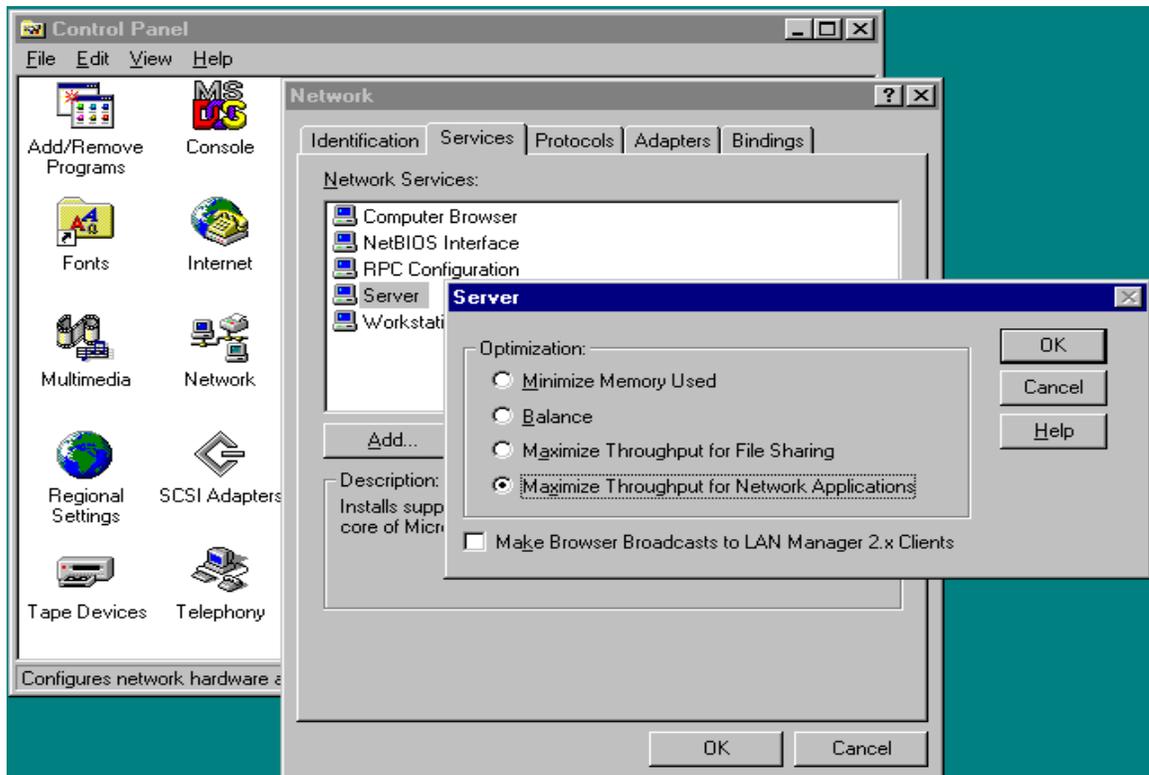


Figure 7. Configuring server memory

<sup>6</sup> The default is "Maximize Throughput for File Sharing."

## Removing Network Protocols

Use Control Panel→Network→Protocols to obtain the Network Protocols list. Remove any protocol not required for server activities. See Figure 8.

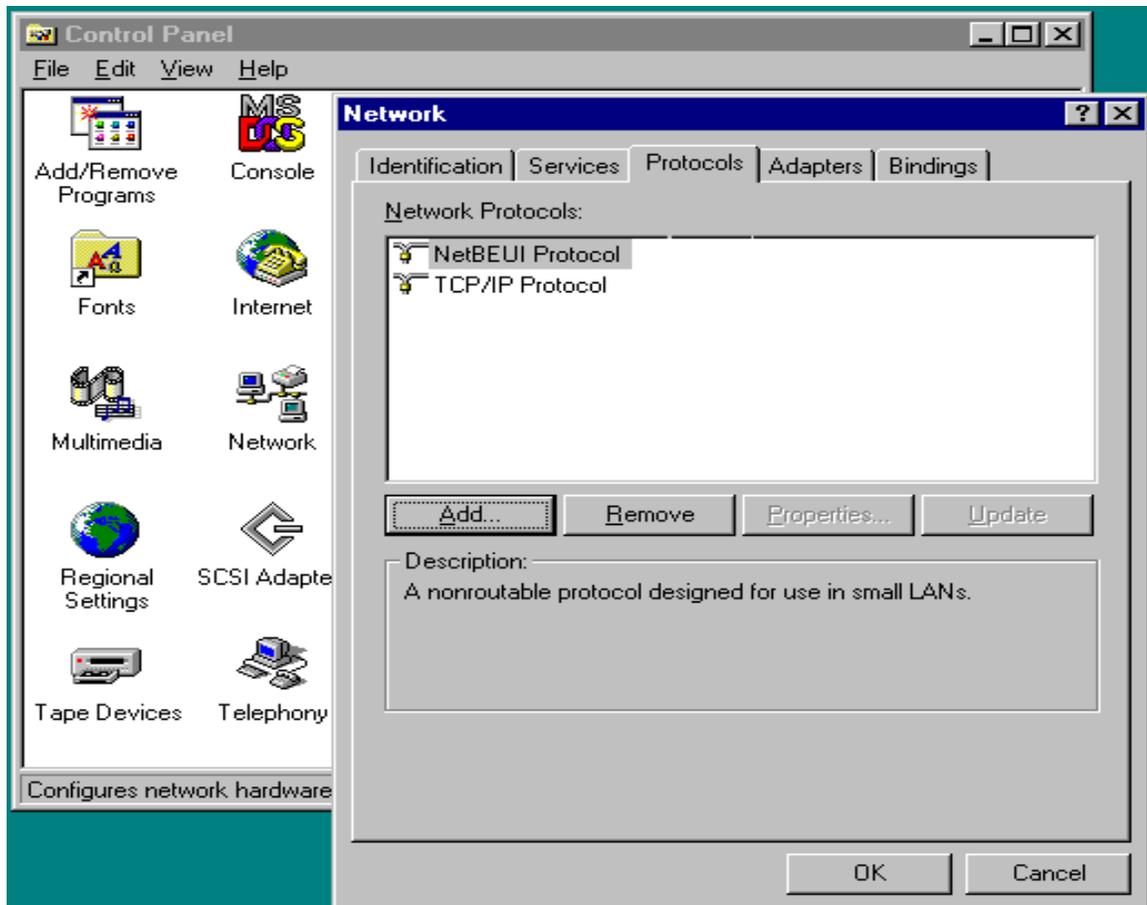


Figure 8. Removing unused network protocols

## Configuring the Server and Workstation Network Bindings

Use Control Panel→Network→Bindings to configure the Server and Workstation network bindings path. Move the most important protocol to the top of the list to reduce the response time for that protocol; position the remaining protocols in descending order of importance. See Figure 9.

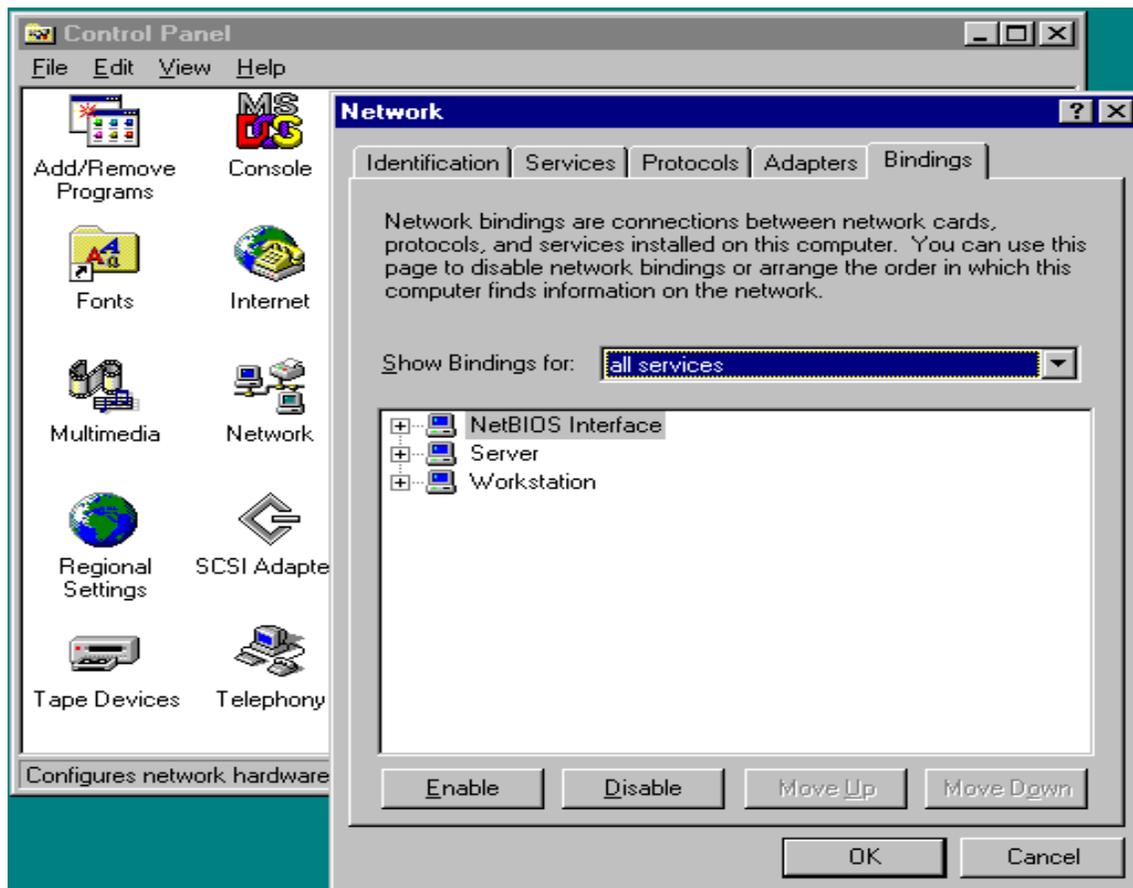


Figure 9. Configuring the protocol binding order

## Stopping Services

If the Oracle8i Server instance is dedicated to being a database server only, you can stop unnecessary services to reduce overhead on the processor(s). The minimum services required for Oracle8i Server are:

- OracleServiceSSSS (where SSSS is the system ID)
- Net8 requires one or more services (depending upon the version of Net8 and the protocols supported) to support user connections such as OracleorantTNSListener.
- Compaq recommends starting the Server, EventLog, and Workstation services, as well as any management services you have installed. You can use the Compaq Software Support Disk to install the management services, which start automatically after installation.

You can turn services off using Control Panel→Services. If desired, you can change the Startup options for your services to Manual. See Figure 10.

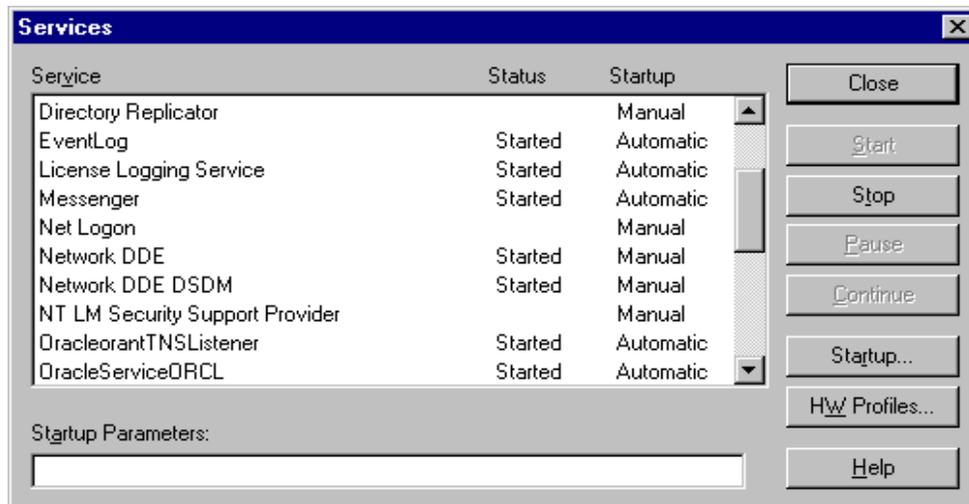


Figure 10. Starting and stopping services

## Configuring the Windows NT File System

Compaq recommends:

- Using FAT for the initial boot partition to allow booting from a DOS diskette for emergency repairs.
- Placing Windows NT and Oracle8i Server executables on the FAT partition as well.
- Placing Oracle8i Server system, log, and data files on NTFS or raw file partitions.

For an I/O-bound system, Compaq testing indicated a 4% - 7% performance gain when using raw devices over FAT or NTFS file systems in Windows NT. If your system is I/O bound, you can switch the log and data files to raw partitions to improve I/O performance but you are limited by the restrictions on using raw devices such as backup and copy. You can minimize the impact of these restrictions by using Oracle8i Server Recovery Manager's copy command and products such as MKS Toolkit's `dd` utility.

Each raw partition equates to a single file and can be represented by a single drive letter (such as \\.\E:) or a logical physical drive number (such as \\.\PhysicalDrive0) or, if you have many partitions per disk, by symbolic links (such as \\.\ORA\_PARTITION\_1 where ORA\_PARTITION\_1 is defined using Oracle's `setlinks` utility).

## I/O Tuning

In most well-tuned Oracle systems, the I/O subsystem is not the performance bottleneck. While adding disk spindles often reduces I/O contention due to their random I/O access, optimizing the OLTP database layout makes a significant impact on performance.

This section focuses on tuning I/O for your OLTP system and provides information on these topics:

- Asynchronous I/O
- Sequential and random I/O
- Disk layout and partitioning
- Physical disk I/O limits
- Array controllers
- Availability versus performance

### Asynchronous I/O

Oracle8i Server uses the asynchronous I/O capabilities of Windows NT by allowing the `DBWn` and `LGWR` threads to issue multiple I/Os simultaneously. This reduces the I/O time and allows the `DBWn` and `LGWR` threads to wake up any waiting Oracle threads sooner, thereby improving throughput and reducing response times.

Using the asynchronous I/O capabilities of Windows NT and Oracle8i Server means that only one `DBWn` thread is necessary in the system.

### Sequential and Random I/O

To achieve maximum performance when Oracle8i Server is accessing datafiles sequentially, the disk(s) must be dedicated to this purpose. Oracle8i Server uses sequential disk access for these activities:

- Accessing transaction logs
- Making full table scans
- Making full index scans
- Making direct loads with SQL\*Loader (reading from input files and writing to datafiles)
- Redo log archiving

Of primary importance are the Oracle8i Server transaction log files, which are generally accessed in a sequential, write-only fashion; when archiving logs or recovering the database, Oracle8i Server accesses transaction logfiles in a sequential, read-only fashion.

In a typical multi-user database system, file access is random. You should spread the datafiles over as many physical disks as necessary to achieve random I/O rates that do not exceed the recommendations given in Table 7. This is best achieved using the disk striping available with a Compaq SMART Array controller, which ensures that the load is balanced equally across all the disks in a volume and allows a high degree of parallelism to occur on accesses. For more information on optimizing array configurations, refer to *Configuring Compaq RAID Technology for Database Servers* which is available on the *Compaq ActiveAnswers* website at <http://www.compaq.com/activeanswers>.

## Disk Layout and Partitioning

Disk layout and partitioning are two of the most important areas of I/O tuning and OLTP performance optimization and Compaq offers these recommendations based on testing carried out on *ProLiant* servers.

### Disk Layout

You should consider these issues when attempting to optimize OLTP performance:

---

**Note:** It is helpful to know the partitioning requirements of your system prior to determining the disk layout.

---

- **Redo log access:** Redo log access is 100% sequential I/O and should be isolated, if possible. Ensuring fast access is essential so Compaq recommends configuring the drives containing the redo logs for hardware fault tolerance using RAID-1 to maximize performance and reliability. For more information, refer to *Configuring Compaq RAID Technology for Database Servers* which is available on the *Compaq ActiveAnswers* website at <http://www.compaq.com/activeanswers>.
- **Redo log archive files:** Redo log archive files are also 100% sequential I/O access. You should isolate these files for maximum performance and configure the drives containing the files for hardware fault tolerance using RAID-1.
- **Data files:** Data file access is usually random and should be balanced across all drives allocated to data and indexes. Increasing the number of physical drives allows you to achieve faster I/O rates; using a striped array ensures that I/Os are well distributed.

### Disk Partitioning

Once you have determined your disk layout, you can configure the disks and create partitions. Use the Compaq Array Configuration Utility to create logical disks and configure the array controller cache and RAID levels; use the Windows NT Disk Administrator to split the logical drives into partitions. See Figure 11.

---

**Note:** To create a partition, select the disk and invoke Partition→Create Logical.

---

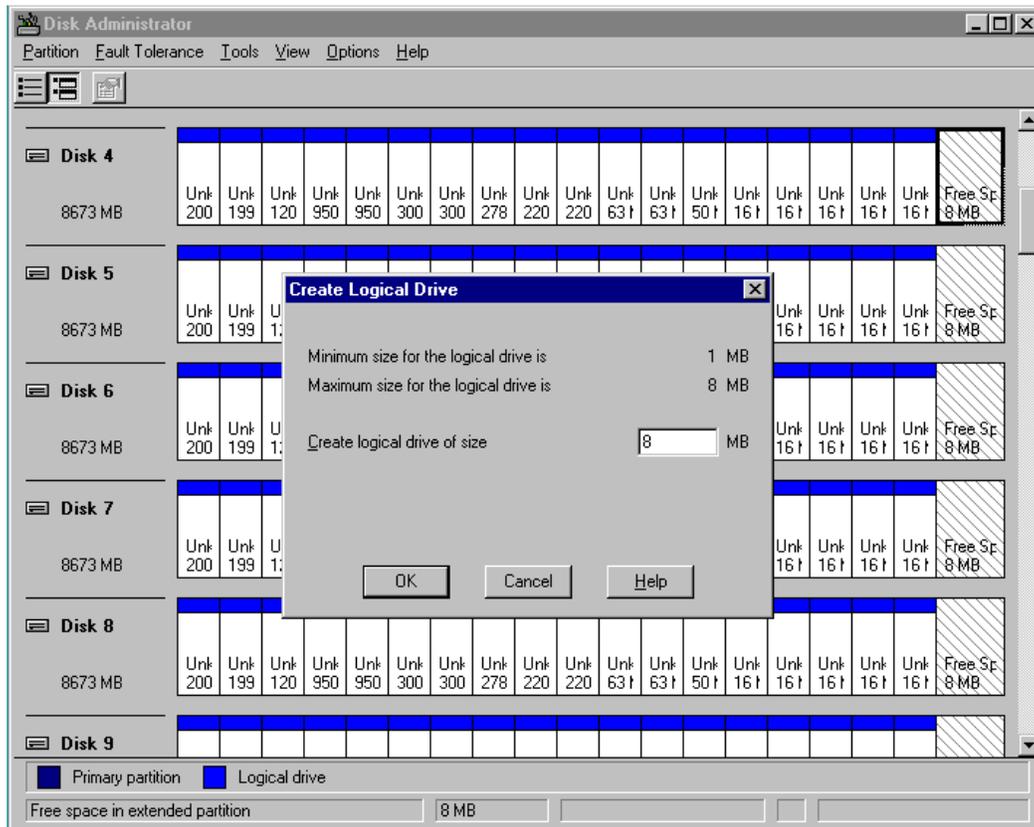


Figure 11. The Windows NT Disk Administrator

Partitioning tables and indexes in Oracle8i Server provides a more granular table and index layout and allows the better management of table and disk storage. When setting up partitions, you should consider these issues:

- **Table partitioning:** You can partition any table that is not part of a cluster and does not contain Long or Long Raw data types. Oracle8i Server supports range, hash, and composite partitioning:
  - **Range partitioning** uses a range of column values (partitioning key) to map rows or index entries to partitions. The partitioning key consists of an ordered list of up to sixteen columns and is based on the partitioning column values. You should determine which data values you wish to partition; the most common range partitioning is by date.
  - **Hash partitioning** uses a hash function on the partitioning columns to stripe data into partitions. It is beneficial to use hash partitioning if you do not know the amount of data in a range, if the size of range partitions differ substantially, or if partition pruning and partition-wise joins on a partitioning key are important.
  - **Composite partitioning** uses the range method to partition data and, within each partition, uses the hash method to sub-partition the data. Composite partitioning provides both management benefits and data placement and parallelism advantages.

- **Index partitioning:** You can partition an index that is not a cluster index and is not defined on a cluster table. Oracle8i Server supports two primary types of index partitioning, local and global:
  - In a **local** partitioned index, all keys in a particular index partition refer only to rows stored in a single underlying table partition<sup>7</sup>.

---

**Note:** A bitmap index on a partitioned table must be a local index.

---

- In a **global** partitioned index, the keys in a particular index partition may refer to rows stored in more than one underlying table partition<sup>8</sup>.

In addition, Oracle8i Server supports these index types:

- Prefixed and non-prefixed local indexes
- Prefixed global indexes

Prefixed indexes are partitioned on a left prefix of the index columns, while a non-prefixed index is partitioned on something other than a left prefix. For example: a **SALES** table and its index **SALES\_INDEX** are both partitioned on the **MONTH** column, making **SALES\_INDEX** a local index. If the **SALES\_INDEX** is defined on the columns **MONTH**, and **SALES\_AMT**, then **SALES\_INDEX** is local prefixed because **MONTH** is the left prefix of the index columns. If **SALES\_INDEX** is defined on the **QTR** column, then **SALES\_INDEX** is a local non-prefixed index. See Table 6.

**Table 6. Comparison of Prefixed and Non-Prefixed Indexes**

Local Prefixed Index	Sales Table	Local Non-Prefixed Index
<b>SALES_INDEX</b>	<b>SALES</b>	<b>SALES_INDEX</b>
Partitioned by <b>MONTH</b>	Partitioned by <b>MONTH</b>	Partitioned by <b>MONTH</b>
<b>(MONTH, SALES_AMT)</b>	<b>(MONTH, QTR, SALES_AMT)</b>	<b>(QTR, SALES_AMT)</b>

A global index follows the same left prefix rule as a prefixed local index.

Local indexes facilitate index manageability and provide more availability when there are partition or subpartition maintenance operations on the table. When partitions in an underlying table are added, dropped, merged, or split, or when hash partitions or subpartitions are added or coalesced, Oracle8i Server maintains the index partitioning for local indexes on that table to ensure that the local index remains equipartitioned<sup>9</sup> with the table.

In an OLTP environment, global and local prefixed indexes provide better performance than local non-prefixed indexes because they minimize the number of index partition probes. Local non-prefixed indexes are useful for historical databases<sup>10</sup> where the index is defined on one column (such as **SALES\_NO**) for faster access but partitioned on another column (such as **MONTH**) to ease partition maintenance.

---

<sup>7</sup> All keys in a particular index partition map to all rows in a single underlying table partition.

<sup>8</sup> All indexes in one partition may map to rows in multiple table partitions.

<sup>9</sup> The index and table are equally partitioned.

<sup>10</sup> Databases that track information over a time interval.

- **Storage parameters:** All partitions of a table or index have the same logical attributes but their physical attributes can differ. You can assign each partition to its own tablespace and specify unique physical attributes (such as **PCTFREE**, **PCTUSED**, **INITTRANS**, **MAXTRANS**) and storage parameters (**INITIAL**, **NEXT**, **MINEXTENTS**, **MAXEXTENTS**, **FREELIST**, **FREELIST GROUPS**).

For more information on partitioning, please refer to the *Oracle8i Server Concepts Manual* or the *Oracle8i Server Application Developer's Guide*.

## I/O Limits

Typically, an Oracle8i OLTP system can be characterized as many users acting on a small subset of data throughout the database. The resulting I/O profile includes very heavy random reads and writes across the data and index files while transaction logs receive a heavy stream of sequential write operations of 2KB or less. In this case, it is important to monitor the I/O rate on the disk subsystem and ensure that individual disks are not overloaded with random I/Os.

The (Physical Disk:Disk Transfers/sec), (Physical Disk:Disk Reads/sec), and (Physical Disk:Disk Writes/sec) counters in the Windows NT Performance Monitor display the number of I/O operations per second handled by the disk subsystem. If you have a Compaq SMART Array controller with several drives allocated to a logical volume, the counters monitor the total number of disk transfers to the entire volume. To calculate the number of transfers per second for each individual disk drive, you must first divide the number of I/O operations per second by the number of disk drives and then factor in the appropriate RAID overhead.

When monitoring and optimizing the drive subsystem, Compaq recommends that you do not exceed the values for I/O requests (disk transfers) per second per disk drive given in Table 7. For more information, refer to *Configuring Compaq RAID Technology for Database Systems* available on the *ActiveAnswers* website at <http://www.compaq.com/activeanswers>.

**Table 7. Average Disk Seconds per I/O Operation**

Type of 2KB I/O Request	Compaq Hard Drive					
	4.3-GB (7200 RPM)	4.3-GB (10K RPM)	9.1-GB (7200 RPM)	9.1-GB (10K RPM)	18.2-GB (7200 RPM)	18.2-GB (10K RPM)
Sequential Write	180	290	210	290	210	290
Random Read/Write	55-60	75-80	70-75	75-80	70-75	75-80

Use this procedure to estimate the average number of I/O requests per second for each disk drive:

1. Using Windows NT Performance Monitor, determine the number of (Physical Disk:Disk Transfers/sec), (Physical Disk:Disk Reads/sec), and (Physical Disk:Disk Writes/sec). Average these figures out over an extended time period rather than selecting a few minutes of peak activity (such as a checkpoint).
2. Use the following equations to calculate I/Os per second per disk drive.
  - **RAID 0:** I/Os per second per drive = (Physical Disk:Disk Transfers/sec) / Number of drives in the volume
  - **RAID-1:** I/Os per second per drive = ((Physical Disk:Disk Reads/sec) + 2\*(Physical Disk:Disk Writes/sec)) / Number of drives in the volume

- **RAID-5:** I/Os per second per drive = ((Physical Disk:Disk Reads/sec) + 4\*(Physical Disk:Disk Writes/sec)) / Number of drives in the volume
3. Repeat these steps for each logical volume. For more information on planning and configuring Compaq RAID technology, please refer to *Configuring Compaq RAID Technology for Database Servers* on the *ActiveAnswers* website at <http://www.compaq.com/activeanswers>.

In the example displayed in Figure 12, there is no fault tolerance (RAID-0) on logical NTFS volume G. The average I/O rate is 478.262 I/Os per second. If the drive set consists of seven 4.3-GB (7200 RPM) drives, then the I/O rate is  $478.262 / 7 = 68.32$  I/Os per second per drive. The test purposely flooded the logical drive with I/Os giving a scenario above the recommended 60 I/Os per second per disk.

If the average I/O rate values significantly exceed those suggested in Table 7, increase the speed of your disk subsystem by adding more drives or using faster drives. Replacing the 4.3-GB (7200 RPM) drives with any of the remaining drive types listed in Table 7 is appropriate since the minimum disk I/O per second rate for these drives is 70 or greater.

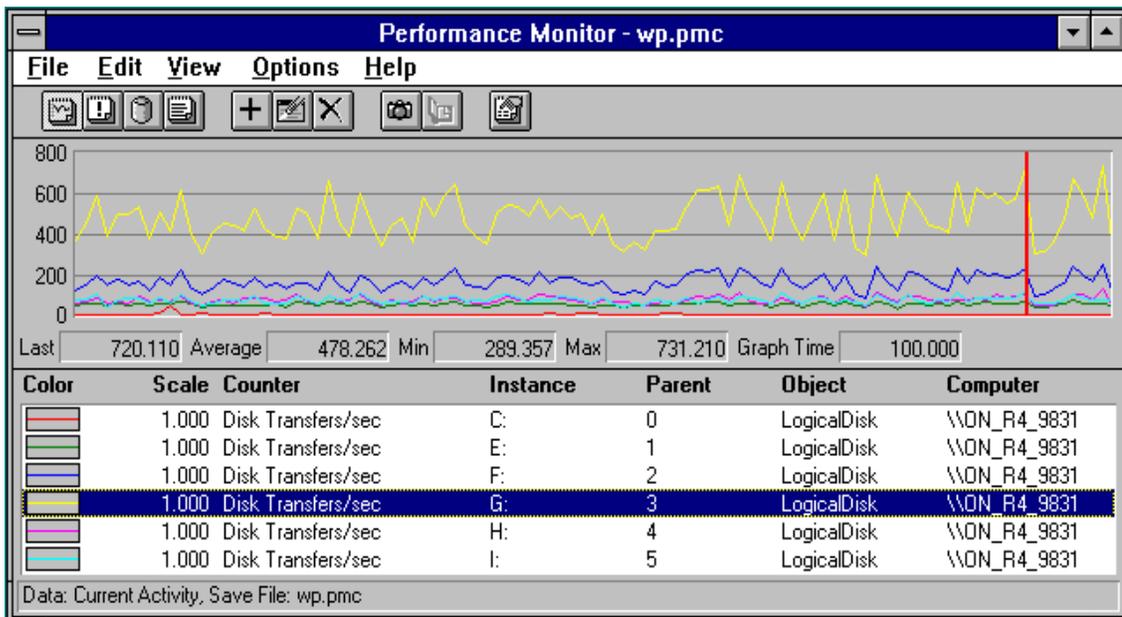


Figure 12. The Windows NT Performance Monitor showing disk I/O

## Array Controllers

Compaq recommends using SMART Array controller technology rather than generic SCSI controllers. SMART Array controller technology is present in Compaq SCSI and Fibre Channel storage solutions and brings these benefits:

- Both Compaq storage solutions provide higher throughput, large drive arrays, and multiple disk striping.
- The Compaq Fiber Channel storage solution provides shared and distributed storage benefits.
- Both Compaq storage solutions provide hardware RAID protection, including multiple RAID levels on a single drive array and hardware RAID protection for your boot drive(s).

## Compaq SCSI Storage Solutions

Consider these issues when deploying Compaq SMART Array 3200, 4200, and 4250ES<sup>11</sup> solutions:

- You can configure up to 56 disks as one or more logical drives. When multiple disks are collected into a single array, the data on logical drives created on that array are striped across all disks for even distribution. This balances your I/O load more evenly and works well for random I/O loads.
- You can configure multiple logical drives on each disk or array of disks. The maximum number of logical drives for a single controller is 32.
- You can configure the 56MB accelerator cache as either read-ahead cache or write-posting cache or a combination of the two<sup>12</sup> depending on the number of database writes. The cache memory is battery-protected.
- Rather than using software RAID on Windows NT, you can define RAID levels at the controller level to relieve system processor overhead.

You can find more information on Compaq SMART Array controllers at [http://www.compaq.com/products/storageworks/array-and-scsi-controllers/controller\\_options.html](http://www.compaq.com/products/storageworks/array-and-scsi-controllers/controller_options.html).

## Compaq Fibre Channel Storage Solutions

The Fibre Channel RAID Array 4000 storage solution combines the benefits of a Fibre Channel with the Compaq SMART Array controller. Consider these issues when deploying a Fibre Channel solution:

- You can configure up to 12 disks as one or more logical drives. When multiple disks are collected into a single array, the data on logical drives created on that array are striped across all disks for even distribution. This balances your I/O load more evenly and works well for random I/O loads.
- You can configure multiple logical drives on each disk or array of disks. The maximum number of logical drives for a single controller is 8.
- You can configure the accelerator cache as either read-ahead cache or write-posting cache or a combination of the two<sup>13</sup> depending on the number of database writes. The accelerator memory is 64MB of which 48MB is battery-powered and is available for read/write; 64MB is available for read-ahead cache.
- Rather than using software RAID on Windows NT, you can define RAID levels at the controller level to relieve system processor overhead.

You can find more information on the Compaq Fibre Channel RAID Array 4000 storage solution at <http://www.compaq.com/products/storageworks/raidstoragesystems.html>.

---

<sup>11</sup> This is an internal controller. The number of drive bays in the system limits the number of drives supported.

<sup>12</sup> Your choices for allocating the memory are 100/0 (100% read-ahead cache with 0% write-posting cache), 75/25, 50/50, 25/75, or 0/100.

<sup>13</sup> Your choices for allocating the memory are 100/0 (100% read-ahead cache with 0% write-posting cache), 90/10, 80/20, 70/30, 60/40, 50/50, 40/60, 30/70, 25/75, and 0/75.

## Availability versus Performance

Once you have calculated the disk space required for your OLTP system, it is important to determine the appropriate level of protection for your drives. Database availability, performance, and cost are considerations when determining the Compaq SMART Array controller RAID level to use.

---

**Note:** The RAID level you select directly impacts system performance and the physical disk space required for your OLTP system.

---

Compaq solutions support these RAID levels:

- **RAID-0:** RAID-0 has a 1:1 write ratio (one physical write for each logical write) and a 1:1 read ratio (one physical read for each logical read). All RAID-0 data are stored in one physical location. RAID-0 offers no fault tolerance for your drives but, since the number of drives required is equal to the number of logical drives, this is the lowest cost option.
- **RAID-1:** RAID-1 has a 2:1 write ratio (two physical writes for each logical write) and a 1:1 read ratio. The **data mirroring** feature of RAID-1 provides fault tolerance by storing data in two physical locations and provides the additional benefit of **split reads**.

Split reads provide a performance boost, especially when there are not many writes, since the database can read from either copy of the data.

Since twice the storage is required, RAID-1 is the highest cost option but is the most fault tolerant.

- **RAID-5:** RAID-5 has a 4:1 write ratio (two physical reads and two physical writes for each logical write) and a 1:1 read ratio. The distributed **data guarding** feature of RAID-5 provides fault tolerance but, since all RAID-5 data are stored in one physical location, there is no benefit from split reads. Parity is striped across all drives in the RAID-5 array.

The number<sup>14</sup> of disk drives required for a RAID-5 configuration is the number of drives required for datafiles plus one additional drive.

There is a substantial performance penalty of 20% - 35% during writes, depending upon your application.

With RAID-5, Compaq recommends limiting the number of physical drives to 14 for each logical drive.

Compaq SMART Array controllers have the flexibility to use any of these RAID levels either individually or in combination.

Depending upon your database availability, performance, and cost requirements, it may be beneficial to use RAID-1 on critical drives and RAID-0 or RAID-5 on other drives. Compaq recommends placing log drives on a RAID-1 drive to optimize performance and ensure that no log files are lost in the case of a single disk failure.

For more information on Compaq SMART Array controller RAID selection and implementation, refer to *Configuring Compaq RAID Technology for Database Servers*, which is available on *ActiveAnswers* at <http://www.compaq.com/activeanswers>.

---

<sup>14</sup> There must be a minimum of three drives for a RAID-5 array.

## Processor Scaling

While no operating system or hardware provides 100 percent scalability, Windows NT and Oracle8i Server provide above-average levels of scalability so that, if your system is processor-bound, you can improve system performance by adding processors.

Because well-tuned Oracle8i Server instances are typically system processor-bound, most applications see significant improvements after the addition of more powerful processors, more processors, or more processor cache. However, if the system is memory- or I/O-bound, adding system processors will not provide a significant improvement.

Figure 13 provides a graphical representation of system processor scalability on the *Compaq ProLiant 8500* using tests that are both disk- and memory-intensive.

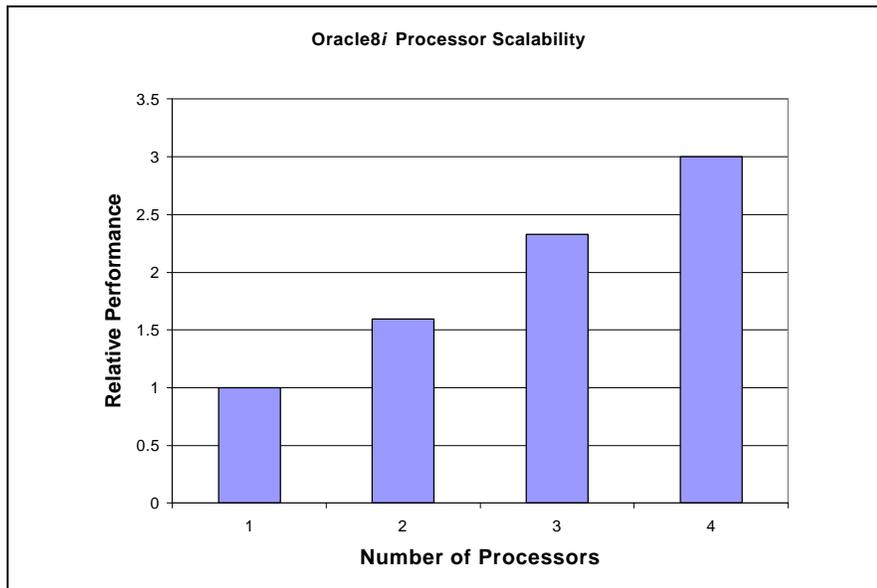


Figure 13. Graphical representation of system processor scalability

## Memory Tuning

Since the 3GB memory limit imposed by Oracle8i Server 8.1.5 includes memory allocated to user threads, it is important to tune your Oracle8i Server so there is enough memory for the user threads (PGA) and the SGA.

---

**Note:** Allocating too much memory to the SGA decreases the memory available for the PGA.

---

Compaq offers memory tuning recommendations in this section for:

- Database block buffers
- Multiple buffer pools
- Database block size
- SGA size
- User capacity

- PGA size
- Number of threads
- Windows NT memory

## Database Block Buffers

To reduce excess I/O, tune the Oracle8i Server cache buffers to maximize the use of the available resources using the Oracle8i Server initialization file (*InitSID.ora*) parameter

**DB\_BLOCK\_BUFFERS.**

Since Oracle8i Server is a single process, multi-threaded application, all threads share the memory from the single process. To determine how much memory is available for buffers, check the Performance Monitor (Memory: Available Bytes) to identify the amount of memory available in the system. See Figure 14.

You should allocate a minimum of 128MB to Windows NT and, if your system is not a dedicated database server, you should determine how much memory is required for other applications running on the system. Adding more physical memory to your system<sup>15</sup> allows you to increase the size of **DB\_BLOCK\_BUFFERS.**

---

**Note:** If you use all available memory for buffers, **swapping** may occur. The performance degradation caused by swapping is greater than the improvement obtained by tuning the Oracle8i Server buffers.

---

---

<sup>15</sup> Up to 4GB.

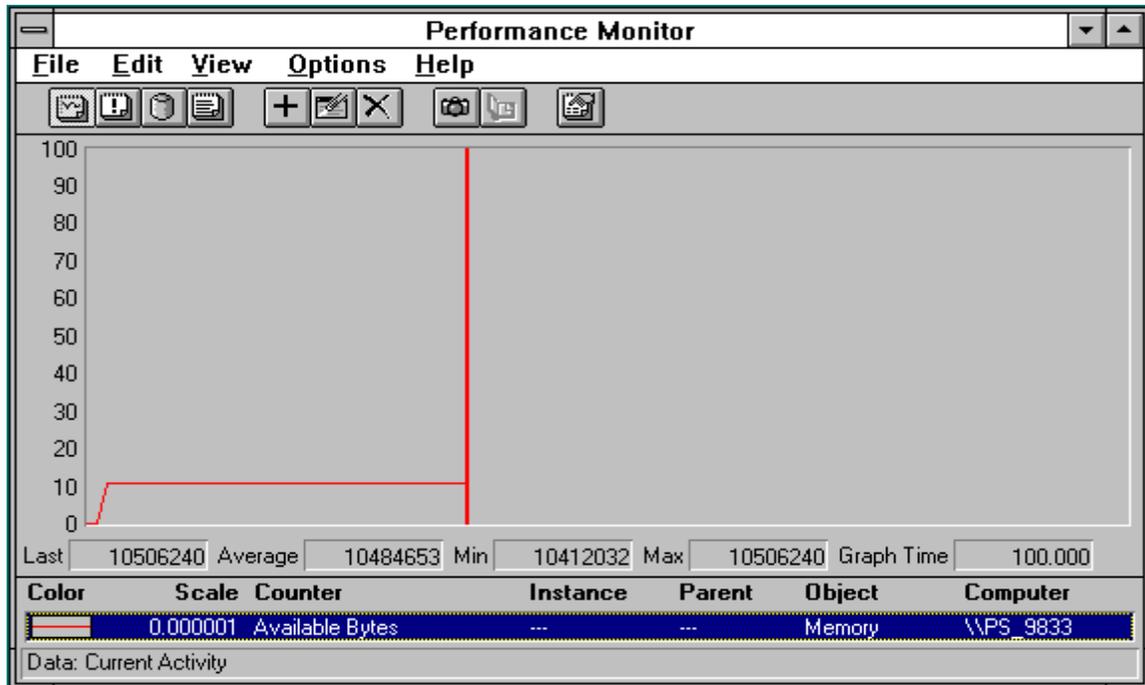


Figure 14. Performance Monitor showing available memory

## Multiple Buffer Pools

You can allocate schema objects to particular cache buffer pools. Oracle8i Server places all blocks read from the object in the appropriate buffer pool.

Oracle8i Server supports these buffer pools:

- **KEEP:** The **KEEP** buffer pool maintains objects in the buffer cache. Small, frequently accessed tables that require a quick response time are good candidates for the **KEEP** pool.
- **RECYCLE:** The **RECYCLE** buffer pool recycles objects to eliminate the consumption of unnecessary cache space. A very large, infrequently accessed table with random I/O is a good candidate for the **RECYCLE** pool.
- **DEFAULT:** Oracle8i Server maintains the **DEFAULT** buffer pool for objects that have not been assigned to either the **KEEP** or **RECYCLE** pools.

For more information on enabling multiple buffer pools, refer to the *Oracle8i Tuning Guide*.

## Database Block Size

You should select the value for the database block size parameter, **DB\_BLOCK\_SIZE**, based on your most critical type of transaction load.

---

**Note:** The default value for **DB\_BLOCK\_SIZE** is 8192.

---

Compaq makes these recommendations for database block size:

- If performance is most critical during applications that primarily access the database in a random fashion (small reads or updates scattered throughout the database), use a block size of 2048.
- If most applications are accessing the database sequentially when performance is most critical, use a block size of 16384.
- If you have a split of random and sequential access during the most critical time, use a block size of either 4096 or 8192 depending upon the ratio of random to sequential access.

## System Global Area (SGA) Size

---

**Note:** Since the SGA resides in server memory, you must allocate this memory at the startup of the instance.

---

The SGA is made up of the following components:

- **Fixed size:** The Oracle products that you have installed determine the fixed size.
- **Variable size:** Oracle8i Server initialization file (**InitSID.ora**) parameters such as **SHARED\_POOL\_SIZE**, **DB\_BLOCK\_BUFFERS**, and **DB\_BLOCK\_LRU\_LATCHES** determine the variable size. In addition, the variable size includes the buffer headers for the database block buffers.
- **Database block buffers:** You can determine the size of the database block buffer component by multiplying **DB\_BLOCK\_BUFFERS** by **DB\_BLOCK\_SIZE**.
- **Redo Buffers:** You can determine the size of the redo buffer component by adding approximately 80KB of overhead to the value of the **LOG\_BUFFER** parameter.

## User Capacity

The number of users you wish to connect has an impact on tuning. A number of parameters are necessary to allow a specified number of connections into the database engine and there may be several areas of operating system tuning required based on your configuration.

Oracle8i Server initialization parameters that impact user capacity include **PROCESSES**, **SESSIONS**, **CURSORS**, and **TRANSACTIONS**.

If the Oracle available memory (SGA and PGA) has been exceeded, users may not be able to connect to the database. If you exceed your user limit or if you wish to add more concurrent users, you can either decrease the SGA or PGA, increase the Oracle available memory, or reduce the memory required by your Oracle application(s).

## Program Global Area (PGA) Size

The PGA is a region of memory that contains data and control information for a single connection. Since this memory must be available at connect time for a particular user, the amount of free server memory limits the number of concurrent connections.

The values of these parameters impact the size of the PGA:

- **OPEN\_LINKS**
- **DB\_FILES**
- **LOG\_FILES**
- **HASH\_AREA\_SIZE**: If you set **HASH\_AREA\_SIZE** to 10MB but only 1MB is needed for your hash area, Oracle8i Server adds only 1MB to the PGA.
- **SORT\_AREA\_SIZE**: If you set **SORT\_AREA\_SIZE** to 10MB but only 1MB is needed to sort your data, Oracle8i Server adds only 1MB to the PGA.

## Number of Threads

You should update the value of the Oracle8i Server initialization file (**InitSID.ora**) parameter **PROCESSES**, which specifies the number of threads that Oracle8i Server can create. This value must include the maximum number of concurrent connections, the two Oracle service threads, and the background threads.

---

**Note:** The **v\_\$PROCESS** and **v\_\$BGPROCESS** views contain information about the Oracle8i Server threads being used.

---

## Windows NT Memory

In addition to the memory required for Oracle8i and its processes, you must tune the memory allocated to Windows NT.

---

**Note:** You can monitor memory usage via the Windows NT Task Manager→Performance window.

---

The primary goals in tuning Windows NT memory are:

- **Reducing the amount of paging and swapping:** Windows NT uses paging or swapping to move information from one storage location (such as real memory, virtual memory, or disk) to another. Excessive paging or swapping can reduce performance.

To monitor paging, use the Windows NT Performance Monitor→(Memory:Pages Input/sec) where (Memory:Pages Input/sec) records the number of pages read in from disk to resolve memory references to pages that were not present in memory when referenced. If paging-in<sup>16</sup> occurs, either increase the total memory on your system or decrease the amount of memory you have allocated to the SGA.

---

<sup>16</sup> (Memory:Pages Input/sec) is greater than zero.

- **Fitting the SGA and PGA into main memory:** Since the purpose of the SGA is to store data for fast access, the SGA should always be part of main memory. You can cause Oracle8i Server to read the entire SGA into memory when you start your instance by setting the value of the `InitSID.ora` parameter `PRE_PAGE_SGA` to `YES`. This may increase the amount of instance startup time but typically decreases the time taken for Oracle8i Server to reach its full performance capacity after startup.

---

**Note:** Setting `PRE_PAGE_SGA` to `YES` does not prevent Windows NT from paging or swapping the SGA after it is initially read into memory.

---

As stated earlier in the “Memory Tuning” section of this Technical Guide, it is important for your entire SGA and PGA to fit into Oracle available memory to prevent paging or swapping.

## Updating InitSID.ora

Earlier sections of this Technical Guide have discussed the tuning of certain Oracle8i Server initialization file (`InitSID.ora`) parameters; this section discusses additional `InitSID.ora` parameters and provides further information on the previously discussed parameters.

Table 8 provides information on tuning `InitSID.ora` parameters.

**Table 8. InitSID.ora Parameters**

Parameter	Description
<code>BUFFER_POOL_KEEP</code>	The <code>BUFFER_POOL_KEEP</code> parameter defines the <code>KEEP</code> buffer pool. The number of buffers in the <code>KEEP</code> buffer pool must be at least 50 times the number of <code>DB_BLOCK_LRU_LATCHES</code> allocated to this pool.
<code>BUFFER_POOL_RECYCLE</code>	The <code>BUFFER_POOL_RECYCLE</code> parameter defines the <code>RECYCLE</code> buffer pool. The number of buffers in the <code>RECYCLE</code> buffer pool must be at least 50 times the number of <code>DB_BLOCK_LRU_LATCHES</code> allocated to this pool.
<code>CURSOR_SPACE_FOR_TIME</code>	<p>The <code>CURSOR_SPACE_FOR_TIME</code> parameter specifies whether or not Oracle8i Server can de-allocate a shared SQL area from the library cache to make room for a new SQL statement.</p> <p>If there are no library cache misses in your OLTP system, setting this value to <code>TRUE</code> may slightly improve the performance of execution calls while preventing the de-allocation of private SQL areas until associated application cursors are closed.</p>
<code>DB_BLOCK_BUFFERS</code>	<p>The size of the buffer cache is equal to the value of the <code>DB_BLOCK_BUFFERS</code> parameter multiplied by the value of the <code>DB_BLOCK_SIZE</code> number.</p> <p>You can calculate the <b>cache hit ratio</b> from:  <math display="block">\text{Cache hit ratio} = (1 - (\text{physical reads} / (\text{db block gets} + \text{consistent gets})))</math>           You can obtain the values for “physical reads,” “db block gets,” and “consistent gets” from the <code>V\$SYSSTAT</code> view.</p> <p>If the cache hit ratio is less than 0.9, you should increase the <code>DB_BLOCK_BUFFERS</code> value, if possible. If an increase to the value of <code>DB_BLOCK_BUFFERS</code> improves the performance of your OLTP system, try another increase.</p>

*Continued*

**Table 8. InitSID.ora Parameters** *continued*

Parameter	Description
<b>DB_BLOCK_SIZE</b>	<p>You must specify the <b>DB_BLOCK_SIZE</b> parameter before creating the database since a change in <b>DB_BLOCK_SIZE</b> requires the database to be rebuilt.</p> <p>Depending on the application, the value of <b>DB_BLOCK_SIZE</b> should be between 2KB and 16KB.</p>
<b>DB_FILE_MULTIBLOCK_READ_COUNT</b>	<p>The value of the <b>DB_FILE_MULTIBLOCK_READ_COUNT</b> parameter determines how many database blocks are read with a single operating system read.</p> <p>Since the maximum I/O size of the Compaq SMART Array controller is 64KB, you should set the <b>DB_FILE_MULTIBLOCK_READ_COUNT</b> to 64KB divided by the value of the <b>DB_BLOCK_SIZE</b> parameter (or a multiple thereof). For a <b>DB_BLOCK_SIZE</b> of 4KB, Compaq recommends a <b>DB_FILE_MULTIBLOCK_READ_COUNT</b> of 16.</p>
<b>DB_BLOCK_LRU_LATCHES</b>	<p>The maximum number of <b>DB_BLOCK_LRU_LATCHES</b> is the lower of (# CPUs * 2 * 3) or (<b>DB_BLOCK_BUFFERS</b> / 50).</p> <p>If you are using multiple buffer pools, you should set <b>DB_BLOCK_LRU_LATCHES</b> to a multiple of the number of processors in the system.</p>
<b>LOG_BUFFER</b>	<p>The value of the <b>LOG_BUFFER</b> parameter reserves a fixed space for the redo log buffer.</p> <p>If the ratio of redo log space requests to redo entries is greater than 1:5000, you should increase the size of <b>LOG_BUFFER</b> until the space request ratio stops falling.</p>
<b>SHARED_POOL_SIZE</b>	<p>The value of the <b>SHARED_POOL_SIZE</b> parameter should be large enough to hold all data dictionary, shared SQL, and compiled objects.</p>

For more information on *InitSID.ora* parameters, refer to the *Oracle8i Tuning* manual.

## Network Tuning

Oracle8i Server uses Net8 to support connections using LU6.2, Named Pipes, SPX/IPX, and TCP/IP protocols. For more information on the setup and configuration of Net8, see the *Oracle Net8 Administrator's Guide*.

## System Management

Since managing the operating system, the database, and the hardware is a great concern, Compaq recommends using *Compaq Insight Manager*.

---

**Note:** If the system is used in a mission-critical role, it is especially important that users notify the administrator of any fault in the system so that corrective action can be taken quickly.

---

### *Compaq Insight Manager*

*Insight Manager* is a Windows-based utility<sup>17</sup> that uses SNMP in conjunction with operating system and driver agents on the server to monitor system hardware and operating system components and report alarms (hardware failures, imminent failures, or system degradation caused by a hardware problem). You can configure *Insight Manager* to page the system administrator if an alarm occurs.

Using the *Insight Manager* pre-failure warranty allows you to replace a failing hardware component under warranty before it can impact your system.

---

**Note:** If you installed the Compaq ECI driver on your system, you cannot use *Insight Manager*.

---

## Conclusion

The information in this paper is not a complete tuning guide but supplements other tuning information provided by Oracle and Microsoft.

To achieve an optimal configuration, there are several areas that you should carefully plan and tune: the application, the hardware, the operating system, and the network.

The tuning process is iterative and should be performed several times to achieve the best performance possible. Compaq expects the information provided in this Technical Guide to facilitate the tuning process but, although all the guidelines provided here have been tested extensively, you should not assume that tuning a particular parameter always gives the desired result.

Compaq welcomes feedback on your configurations and experiences to help improve future information products. Please send any comments or suggestions to Compaq on the attached form, using additional sheets if necessary.

---

<sup>17</sup> The utility is available on the *Compaq SmartStart* CD-ROM.

## User Registration/Evaluation Form

Please fill out and return this form to help Compaq keep you up to date with future revisions of this document and related new information products.

Your efforts help to improve the quality of Compaq information products.

**Name:** \_\_\_\_\_  
**Title:** \_\_\_\_\_  
**Company:** \_\_\_\_\_  
**Address:** \_\_\_\_\_  
 \_\_\_\_\_  
 \_\_\_\_\_  
**Phone Number:** \_\_\_\_\_

Please evaluate the quality of this document:

	Excellent					Poor
<b>Technical Accuracy</b>	1	2	3	4	5	5
<b>Organization</b>	1	2	3	4	5	5
<b>Clarity</b>	1	2	3	4	5	5
<b>Completeness</b>	1	2	3	4	5	5
<b>Timeliness</b>	1	2	3	4	5	5

Please indicate the type of environment you have at your site:

Operating Systems	RDBMS	Processing Type
Microsoft Windows NT	Microsoft SQL Server	On-Line Transaction Processing
IBM OS/2	Sybase ASE	Decision Support
Novell Intranetware	Oracle8i	Batch Processing
SCO UnixWare	IBM DB/2	Other:
Compaq Tru64 UNIX	Other:	
Other:		

Please indicate the type of information you would like us to provide in the future:

Topic	Operating Systems	RDBMS
Configuration and Tuning	Microsoft Windows NT	Microsoft SQL Server
Capacity Planning	IBM OS/2	Sybase ASE
Integration Information	Novell Intranetware	Oracle8i
Competitive Analysis	SCO UnixWare	Other:
Systems Management	Compaq Tru64 UNIX	
Other:	Other:	

**Additional Comments:**

**Return to:**  
 Database Performance Engineering  
 Compaq Computer Corporation  
 MailCode 150402  
 20555 SH 249  
 Houston, Texas 77070