

Sizing Oracle on Microsoft Windows and Dell PowerEdge Servers,

a White Paper sponsored by Microsoft, Dell and Oracle

*By Larry Pedigo, Performance Tuning Corporation
December 2004*

Introduction

Sizing server requirements for deploying a new corporate application is one of the least enviable tasks in the IT field. Expectations are always high, and the responsibility for providing sizing estimates often falls to one person. Whether you have a background in Database Administration, Systems Administration, Operations, or Capacity Planning, nothing adequately prepares you for this responsibility. It often seems that you are forced to make what can be described in the most complimentary light as “educated guesses”.

It doesn't help that the available literature often falls short in providing the tools that you need to prepare sizing estimates. One of the key problems is the common confusion between the terms “Sizing” and “Capacity Planning”. Sizing is best described as the process of estimating the hardware requirements of a planned application, based on technical descriptions of the customer's needs. Capacity Planning is the process of measuring resource requirements from existing applications, and projecting the amount and kind of hardware required to support larger workloads. Most of the available literature is actually focused on capacity planning, rather than sizing. For example, it is easy to find whitepapers and articles on how to “size” system resources for applications. However, most of these whitepapers and articles assume that you have a thorough knowledge of your application's behavior in the real world, along with copious metrics. The problem is that at the beginning of a brand new project, before you have purchased the first server, you may have almost no idea what resources will be required by the new application. Capacity planning guidelines will be of little use to you.

Fortunately, for those planning Oracle® database deployments on Dell™ PowerEdge™ servers, a great deal of help is already available. Oracle Corporation, Microsoft Corporation, Dell, and Intel have teamed up to provide useful guidelines for real world performance characteristics of their respective system components. These guidelines can be used to prepare accurate sizing estimates for corporate deployments. Most of the available tools can be found on the Dell website at www.dell.com/Oracle. This link is centered on Oracle deployments on Dell PowerEdge servers. Here you may find detailed information about a range of PowerEdge servers, equipped with Intel processors and

featuring Microsoft® Windows® Server Operating Systems. This information includes product guides, technical specifications, whitepapers, and real-world case studies.

This whitepaper should provide further assistance for those sizing Oracle databases on Microsoft Windows and Dell PowerEdge servers. Both Oracle9i and Oracle10g databases are discussed, as well as Microsoft Windows Server 2000 and 2003 versions. However, most of the guidelines presented here are not necessarily version specific, and should be useful across a range of software. The focus is on providing “rules of thumb” for sizing Dell PowerEdge servers for first-time Oracle database deployments. On the other hand, if you are truly involved in “capacity planning”, and already have a good “baseline” of your application’s performance characteristics, then the information presented here should also prove very useful. However, no assumptions are made of extensive prior knowledge of your application’s performance characteristics. If you are starting from scratch, then you are reading the right paper.

Overview of Oracle on Microsoft Windows

Before examining sizing techniques for Oracle servers, it is useful to cover some background information on how Oracle is implemented on the Microsoft Windows platform. The following sections describe how Oracle database technology functions within the Microsoft Windows operating system, both from a Corporate IT perspective and a technical perspective.

Let’s begin by introducing some terminology that will be used throughout this paper:

Oracle Enterprise Manager – Management services for the entire Oracle infrastructure. Included with all editions of Oracle Database.

Oracle Real Application Clusters (RAC) – Allows administrators to run Oracle on two or more systems in a cluster while concurrently accessing a single shared database. This creates a single database system that spans multiple hardware systems, yet appears to the application as a single unified database system, extending high availability and broad scalability to applications. Optional for Enterprise Edition, included with Standard Edition.

Oracle Database Editions:

Enterprise Edition (9i, 10g; server-side) – For unlimited servers and processors. Optimal scalability and high availability with Oracle Real Application Clusters option.

Standard Edition (9i, 10g; server-side) – For servers with up to four processors. Includes Oracle Real Application Clusters at no extra cost.

Automated Storage Management (ASM) – Manages disk striping and mirroring within Oracle Database 10g. Included with all editions of Oracle Database.

Data Guard – Allows administrators to set up and maintain a standby copy of their production database to more quickly recover and maintain continuous database service, even in the event of a disaster. Available with Oracle Database Enterprise or Personal Editions.

Flashback Query – Allows administrators or users to query any data at some point-in-time in the past and can be used to view and reconstruct lost data that may have been deleted or changed by accident. Included with all editions of Oracle 9i Database and Oracle Database 10g.

Oracle Data Provider for .NET (ODP.NET) – .NET data provider designed exclusively by and for Oracle for optimal data access from .NET to Oracle 9i Database and Oracle Database 10g. Can be used from any .NET language, including C# and Visual Basic .NET.

Why Oracle on Windows?

Oracle is dedicated to providing open solutions that integrate with their customer's preferred development and production platforms. A large number of Oracle customers choose Microsoft Windows as their primary application development platform. Many of these customers utilize Windows Servers in their production environment. It is no great surprise then that Windows Server is a tier 1, fully supported platform for Oracle. In fact, Windows Server is a base development platform for Oracle.

Oracle is fully committed to providing the most complete, fully integrated database available for the Microsoft Windows Server platform. Integration features include:

- The Oracle Data Provider for .NET (ODP.NET), a native driver for the .NET development environment.
- Oracle Services for Microsoft Transaction Server (MTS)
- Oracle Objects for OLE (OO4O), a COM-based API
- The Oracle Provider for OLE DB
- Oracle Open Database Connectivity (ODBC) an optimized driver for ODBC access to the database from Windows clients
- The COM automation feature, which allows Oracle databases to communicate with COM Automation features (such as Microsoft Office) using SQL, PL/SQL and Java stored procedures.
- Very Large Memory support: up to 16 GB for Windows NT, up to 64 GB for Windows 2000 and higher
- Integration with Microsoft Active Directory, native Windows authentication, and Oracle Internet Directory. This enables Single Sign-On.
- Oracle Fail Safe high availability for Windows Clusters

Oracle Corporation has a long history with Microsoft Windows. Oracle 7 was the first database to ship on Windows in 1994. Since that time, integration with Windows Server

administration and development tools has been one of the primary focuses of Oracle Engineering. More recently, Oracle was able to provide support for both 32-bit and 64-bit versions of Windows 2003 Server on the same day as the OS release. Oracle and Microsoft also work closely together at several levels; the two companies engage with commodity hardware partners and value-added resellers such as Dell to offer a low-cost platform for managing data; Oracle is a member of the Microsoft Visual Studio Industry Partner (VSIP) program to support integration with Visual Studio .NET; and they actively work together to provide Oracle on Windows and .NET resources for administrators and developers through their respective websites, industry events, publications and press, and other outreach activities. Oracle Database customers deploying on Microsoft Windows Server platforms can benefit from improved scalability, reliability, security, productivity, and overall cost savings as a result of this cooperation between Oracle and Microsoft.

Supported Microsoft Versions

The following chart illustrates support for Microsoft Windows versions for available editions of Oracle9i and Oracle10g. Windows 2000 support includes Professional, Server, Advanced Server, and Datacenter editions. Windows Server 2003 support includes Standard, Enterprise, Datacenter and Web editions for 32-bit, and Enterprise and Datacenter editions for 64-bit.

Oracle Database 10g (10.1)		Supported Windows Operating Systems
Enterprise Edition (Server – unlimited processors)		Windows NT/2000/XP/Server 2003
Standard Edition (Server – up to four processors)		Windows NT/2000/XP/Server 2003
Standard Edition One (Server – up to two processors)		Windows NT/2000/XP/Server 2003
Personal Edition		Windows NT/2000/XP/Server 2003
Client		Windows NT/2000/XP/Server 2003
Oracle9i Release 2 (9.2)		Supported Windows Operating Systems
Enterprise Edition (Server)		Windows NT/2000/XP/Server 2003
Standard Edition (Server)		Windows NT/2000/XP/Server 2003
Personal Edition		Windows NT/2000/XP/Server 2003
Client		Windows NT/2000/98/XP/Server 2003
Oracle9i Release 1 (9.0.1)		Supported Windows Operating Systems
Enterprise Edition (Server)		Windows NT/2000/XP
Standard Edition (Server)		Windows NT/2000/XP
Personal Edition		Windows NT/2000/98/XP
Client		Windows NT/2000/98/XP

Brief Summary of Oracle Architecture on Windows

At a high level, Oracle architecture is essentially the same on all OS platforms. Similar memory components and I/O mechanisms are used on each platform. However, different

OS architectures impose different implementations of memory, processes, and I/O. In addition, Oracle takes advantage of OS optimizations, where available.

Due to these factors, the details of the Oracle architecture on Microsoft Windows are somewhat different than other operating systems, such as UNIX. One key difference is that Oracle is implemented as a multi-thread architecture under Windows, as opposed to a multi-process architecture under UNIX. In the Windows OS, Oracle runs as a single visible process, which is managed as a Windows Service.

In terms of I/O, both the NTFS file system and raw logical or physical partitions are supported. Raw partitions offer faster performance than NTFS, because it bypasses the file system driver. Internally, I/O is 64-bit, whether the Windows version is 32-bit or 64-bit. Asynchronous I/O to files has been supported since Oracle7. The maximum size for a single file is 64 GB. The theoretical maximum database size is 4 petabytes. The largest implementation on Windows currently in production is 10 TB. With the advent of 64-bit versions of Windows, scalability is virtually unlimited.

Internally, Oracle does not impose any limits on memory, number of connections, or other resources. However, the specific OS version may impose resource limits. This is particularly noticeable for memory, where the differences between 32-bit architecture and 64-bit architectures come into play. Under 32-bit versions of Windows (up through Windows 2003 Server), the maximum memory space available is 4GB. This includes combined memory for the OS and the database instance. Actually, the maximum memory space allocated for the SGA, database threads and code is only 3 GB. Very Large Memory support allows access to memory over 4 GB (covered later in this paper).

The 64-bit version of Windows 2003 server removes the 4 GB memory limitation. In addition, Windows 2003 Server has been optimized for Intel® Itanium® 2 processors. On 64-bit Windows 2003, Oracle also benefits from the improved parallelism and efficiency built into the Itanium 2 design. Migration from a 32-bit Oracle database to a 64-bit Oracle database is very easy and involves no export or import, and no database re-creation is required.

In summary, Windows Server provides a powerful platform for Oracle databases. Oracle is designed for optimum performance with both 32-bit and 64-bit versions of Windows Server.

Choosing the Best Oracle Edition

Oracle offers several database editions for different data management needs. The most inclusive edition is the Oracle Database Enterprise Edition, which offers unlimited scalability on servers with any number of processors. Enterprise Edition includes an expanded version of Oracle Enterprise Manager (OEM) with significant management and performance monitoring tools. (Many companies elect to use the Enterprise Edition just to get all of the management tools with OEM.) All high availability options are supported, including Oracle Real Application Clusters (RAC). Enterprise Edition is available for both Oracle9i and 10g, and is also available as a 64-bit version for Windows Server 2003.

Oracle Database Standard Edition offers the same core features as Enterprise Edition. Standard Edition may be installed on servers with four processors or less. The high availability features and security features are somewhat more restricted. Many of the data warehouse features are not available with Standard Edition; however, Oracle Real Application Clusters are included with Standard Edition. The Oracle Enterprise Manager in Standard Edition is more restricted than with Enterprise Edition. Standard Edition is available for Oracle9i and 10g. One of the attractions of the Standard Edition is the lower cost compared to the Enterprise Edition.

Oracle Standard Edition One (SE One) is available for servers with two processors or less. It has essentially the same features as Standard Edition, without clustering support. This edition is targeted for small businesses and departmental deployments that don't need the complexity and cost of Enterprise Edition. SE One is competitively priced compared with any commercial database on the market. At the time of this paper, Dell offers SE One on Windows factory installed on PowerEdge 2600 and 2650 servers, with optional PowerVault™ 220 external SCSI storage.

There are two other Oracle Database editions available: Personal Edition and Lite Edition. Personal Edition offers all of the features of Enterprise Edition, but only for a single user. Lite Edition is designed for mobile database applications.

High Availability Options

Oracle Database Enterprise Edition offers industry-leading high availability features, including Fail Safe, Flashback Query, Flashback table (database and transaction query), and Data Guard to protect businesses from server failure, site failure, storage failure, and human error. While Standard Edition and SE One also provide high availability, only the Fail Safe and Flashback Query features are available for those editions.

RAC offers both high availability as well as scalability, and is an add-on option for Enterprise Edition, but is automatically included with Standard Edition.

Scalability Options

In addition to RAC, both Standard Edition and Enterprise Edition offer Integrated Clusterware, Automatic Workload Management, and PL/SQL native compilation. Integrated Clusterware and Automatic Workload Management are add-on options for the Enterprise Edition, and both require the RAC option.

Oracle on Dell PowerEdge Servers

Dell offers a wide variety of servers that are capable of hosting Oracle databases. It is possible to find the right size server for everything from a single user database to a transactional system serving thousands of users. The brief list below is current as of this paper. Check the Dell website for the latest updates. The most popular servers for corporate environments are the rack-dense servers, and they are listed below. Equivalent Performance Tower servers are also available (see the Dell website).

- **PowerEdge 750** – Single Processor 1U Server
 - **Processor** – Single Intel Pentium® 4 processor up to 3.4 GHz or Single Intel Celeron® processor at 2.4 GHz
 - **Memory** – Up to 4GB DDR400 ECC SDRAM
 - **Storage** – Up to 500 GB maximum internal storage
 - **Oracle Databases** – Small test or production databases with a small number of users
- **PowerEdge 1750** – Dual Xeon™ Processor 1U Server
 - **Processor** – Up to 2 Intel Xeon processors up to 3.2 GHz
 - **Memory** – Up to 8GB ECC DDR SDRAM
 - **Storage** – Up to 500 GB maximum internal storage
 - **Oracle Databases** – Small to medium size test or production databases with a moderate of users, single node or RAC
- **PowerEdge 1850** – Dual Xeon™ Processor 1U Server
 - **Processor** – Up to 2 Intel Xeon processors up to 3.6 GHz
 - **Memory** – Up to 12GB ECC DDR-2 SDRAM
 - **Storage** – Up to 500 GB maximum internal storage
 - **Oracle Databases** – Small to medium size test or production databases with a moderate of users, single node or RAC
- **PowerEdge 2650** – Dual Xeon Processor 2U Server
 - **Processor** – Up to 2 Intel Xeon processors up to 3.2 GHz
 - **Memory** – Up to 12 GB DDR SDRAM
 - **Storage** – Up to 730 GB maximum internal storage
 - **Oracle Databases** – Medium to large size test or production databases with a moderate to large number of users, single node or RAC

- **PowerEdge 2850** – Dual Xeon Processor 2U Server
 - **Processor** – Up to 2 Intel Xeon processors up to 3.6 GHz
 - **Memory** – Up to 12 GB DDR SDRAM
 - **Storage** – Up to 1.8TB maximum internal storage
 - **Oracle Databases** – Medium to large size test or production databases with a moderate to large number of users, single node or RAC
- **PowerEdge 6650** – Quad Xeon Processor MP 4U Server
 - **Processor** – Up to 4 Intel Xeon processors MP up to 3.0 GHz
 - **Memory** – Up to 32 GB DDR266 ECC SDRAM
 - **Storage** – Up to 730 GB maximum internal storage
 - **Oracle Databases** – Large to very large size production databases with a large number of users, single node or RAC
- **PowerEdge 7250** – Quad Itanium 2 4U Server
 - **Processor** – Up to 4 Intel Itanium 2 processors at up to 1.5 GHz
 - **Memory** – Up to 32 GB DDR266 ECC SDRAM
 - **Storage** – Up to 219 GB maximum internal storage
 - **Oracle Databases** – large to very large size production databases with a large number of users, single node or RAC. Supports large Oracle SGAs.

Sizing Methodology

The key to successful sizing is following a well defined methodology. Even though many of the steps that you will need to follow are just “common sense”, it pays to have an organized approach, with a formally defined process. The section below outlines a simple, but effective sizing methodology that should be applicable to any sizing project.

Elements of Sizing

There is a tendency to treat sizing servers as some form of “Rocket Science”. The assumption is that the more statistics, benchmarks, and technical specifications that you can arm yourself with, the better. Contrary to popular opinion, donning a lab coat is not the first step to sizing. Technical analysis is just one element of sizing. It is more important to properly understand the business requirement your systems will fulfill than it is to understand the CPU and RAM specifications of every server on the market.

A properly defined sizing methodology should include the following elements:

- Customer input – although this should be obvious, too many sizing exercises suffer from a poor understanding of customer requirements
- Experience and best practices – a formal attempt should be made to gather relevant documentation from previous projects. In many cases, experience with

sizing methodology on one platform can transfer to another platform. During the project, an effort should be made to document experiences and add to the best practices list.

- Assumption and defaults – every sizing exercise includes assumptions, whether explicit or implicit. It is very helpful to prepare a formal list of assumptions and default values used for the project. As the project continues, you should periodically check this list to see if your assumptions still appear to be valid or need to be modified, given your evolving understanding of the project.
- Scalability benchmarks – yes, benchmarks do in fact play an important role. Count yourself fortunate if you are able to obtain benchmark measurements that are completely relevant to your situation. More commonly, the available metrics do not exactly match what you are looking for. It is okay to employ a certain amount of extrapolation from benchmark results to obtain the justification for your sizing. Of course, you should carefully document your assumptions, as noted above.

Sizing Steps

Regardless of the hardware platform or application software considered during a sizing exercise, there are six steps that all sizing exercises have in common:

Define Customer Requirements

This is undoubtedly one of the most important steps. Don't expect accurate customer requirements to be delivered to you; you must be proactive in obtaining the input you require. One of the most common ways of soliciting customer input is to distribute a questionnaire. Another common method is to hold targeted customer interviews. You shouldn't expect all of your customers to understand technical jargon. You must be prepared to translate from business terms into technical terms.

Collect Customer Information

Again, the key is to proactively gather the information you need. Try to identify any existing systems that have characteristics similar to the system that you are seeking to build. Identify the key legacy applications that your new system will replace or emulate. On those systems, measure the user activity as accurately as you can. Also measure the overall workload. Use these measurements as the "baseline" for your sizing exercise.

Analyze System Data

You should always gather as much technical information as possible from legacy or comparison systems. This is true even if the hardware platform will be different than your target system or even if you expect your new system to have significantly improved performance characteristics. In particular, gather statistics for CPU utilization, memory utilization, and I/Os per second (IOPS). If the comparison system is Microsoft Windows based, you may use the Performance Monitor to obtain the required data.

System Determination

The next step is to extrapolate the technical data you have gathered to design a system that will meet the customer requirements. You will need to predict total CPU usage, total memory usage, an estimate the total I/O load. Based on the I/O load and the total size estimate of customer data, you may estimate storage resources required for capacity and performance.

You will need to pull together all of your available sizing guidelines at this step. This includes any relevant benchmarks and “rules of thumb”. Where possible, compare your requirements to similar existing implementations.

System Test (optional)

When possible, you should include time and budget for a test phase. The earlier this is done in the sizing process the better. A “proof of concept” approach can be used with a system that closely approximates your target system. Even more useful is a “benchmark” approach, where multiple configurations are tested to see which most closely matches the requirements.

Implement

Implementation schedules are always tight; deadlines are always sooner than you would like. When possible, compromise by delivering incremental steps earlier, and then deliberately stagger the full implementation. This allows you to tune performance as you implement, with continuous system monitoring.

Key Assumptions for Oracle Deployments on Dell PowerEdge Servers

While the guidelines for sizing listed above are generally applicable for all sizing projects, the rest of this white paper will focus on the specifics of sizing an Oracle deployment on Dell PowerEdge servers and Dell storage. As mentioned above, it is always desirable to list the assumptions underlying your sizing recommendations. In this paper, there are some key assumptions about how the servers and the oracle database will ultimately be configured:

- The OS is optimally tuned on comparison systems, test systems, and the final implementation. There is no resource contention on the system bus, and I/O is efficiently distributed across HBAs, controllers, and disks.
- The network is optimally tuned for Oracle Net and application network traffic.
- The Oracle instance is optimally tuned:
 - In general, it is recommended that there be only one Oracle instance per PowerEdge server. If possible, plan for multiple, smaller servers, with each server containing a single instance rather than a single massive server

containing all instances. Budgetary and practical considerations may provide exceptions to this rule. However, be aware that sizing and tuning the PowerEdge servers will become a more complicated exercise with multiple instances per server. If you plan on including multiple instances per server, you should document your reasons for making this decision.

- The Oracle System Global Area is properly sized.
- Oracle database storage parameters are configured for efficient block usage without row chaining or row migration
- Applications are tuned for efficient database access. ODBC and JDBC (Java Database Connectivity) drivers should be up-to-date for Oracle9i or 10g, as appropriate. Consider the impact of “thick” client drivers versus “thin” web drivers, as well.
- Applications are not hosted on the database server. The impact of multiple applications residing on the database server can be particularly severe, and should be avoided, if possible.
- OLTP and batch activity occur at different times. If you must perform batch activity such as database loading while users are performing transactions, the impact should be considered in the planning stage. The final implementation should be carefully monitored during overlap periods.

Gathering Customer information

One of the initial steps in the sizing methodology defined above is to gather customer information. Please remember that this is NOT the same thing as defining customer requirements. Presumably, you have already done that and you have already obtained a good understanding of the customer’s business requirements. For gathering customer information, you will certainly need the assistance of the user community. But you will probably need to gather much of this information yourself or with the assistance of knowledgeable IT staff.

Key Questions

Several key questions need to be answered during this step. These include:

- What are the customer’s throughput requirements (amount of work per unit time)?
- What are the customer’s scalability requirements (including expected growth)?
- What is the average number of concurrent connected users (active and inactive sessions). The maximum?
- What is the average number of active concurrent user sessions? The maximum?
- What is the OLTP:batch ratio? How does this vary during the day?
- What are the types and frequency of batch operations?
 - What is the percentage of ad hoc queries?
 - What is the percentage of “canned” queries?

- What is the nature of application interfaces, and how does it impact Oracle connectivity?

Transaction Activity

You will also need to analyze transaction activity. Even though a legacy application's transaction model may not resemble Oracle's transaction model, there is always the concept of a business transaction. At a high level, define what constitutes a business transaction within the legacy application. It is also helpful to group these transactions into types. Break these transactions into atomic steps. You can then predict how these business transactions will translate into Oracle database transactions.

Sizing CPU Requirements

Dell PowerEdge servers feature Intel processors. Internally, Intel processors utilize the Common Instruction Set Chip architecture (CISC). This architecture is a popular alternative to Reduced Instruction Set Chip (RISC) architectures because of its competitive performance at low cost. Since database operations tend to be CPU-intensive, the Intel architecture is a compelling choice.

If you are migrating your applications to Dell PowerEdge servers with Intel processors from a server with slower processors, you can benefit in two ways: increased application speed and increased scalability. The typical impact of faster Intel processors is speed-up of individual application sessions. To achieve scalability, simply add processors. For example, by utilizing a Dell PowerEdge 4-way server, you can achieve scale-up for more concurrent users and more concurrent applications.

Scaling with Multiple CPUs

Windows 2000 and Windows Server 2003 offers support for multiple Intel x86 and Intel Itanium processors with the following Editions:

- Server – up to 4 CPUs. Dell offers a variety of multi-processor servers, ranging from the PowerEdge 1750 with a maximum of two Intel Pentium 4 processors to the Dell 7250, with up to four Itanium 2 (64-bit) processors
- Advanced Server – up to 8 CPUs. Dell offers the 6650 4-way server with Xeon processors and the 7250 4-way server with up to four Itanium 2 (64-bit) processors.
- Datacenter – up to 32 CPUs. There are only a few implementations of large Windows SMP servers on the market. Dell has elected not to offer this type of server.

Industry experience has been that large Symmetric Multi-Processor implementations (above 8 processors) on Windows have not proven to have good scalability or cost-

efficiency characteristics. In fact, this is a general problem with SMP implementations on all O/S architectures. As you add more and more processors on a single system bus, you get diminishing returns in terms of scalability.

Fortunately, Oracle has developed an alternate solution for Oracle databases, starting with Oracle9i. By grouping multiple single processor or multi-processor servers into a Real Applications Cluster, near linear scaling has been documented. If your database needs the processing power of multiple processors, but you do not want to suffer the limitations and cost of a large SMP server, consider RAC.

CPU cache

One of the ways that Intel processors achieve high performance is by utilizing large CPU caches. CPU caches allow data to be stored in memory, but physically resident on the chip. This saves significant time in retrieving information from main memory that will be utilized during multiple CPU cycles.

Two types of cache are offered: Level 1 and Level 2. L1 cache is located on the CPU core and offers the fastest memory access. L1 cache is commonly 8-16 KB in size. 16 KB is the limit on how much information can be cached at the CPU core on an Intel processor. To access more memory, L2 cache is used. L2 cache is located on the chip, but it is accessed through the memory bus. While this is slower access than L1 cache, it is significantly faster than accessing main memory from the CPU. With L2 cache, cache sizes from 128 KB to 4 MB can be achieved. The efficiency of caching can be measured by the “hit rate” for information found within the cache during each cycle vs. information that must be retrieved from main memory. L2 caches can achieve a higher cache hit rate than L1 cache, due to their larger size.

For Oracle databases, always use chips with the largest available L1 and L2 caches. Large cache sizes are actually more important to Oracle performance than processor speed. This is why high-end Dell servers use Intel Xeon processors instead of Pentium 4 processors. Pentium 4 processors have a faster clock speed, but the larger cache sizes on the Xeon processors are generally more beneficial to Oracle performance.

Intel x86 32-bit processors vs. Intel Itanium 2 64-bit processors

If you wish to take your Oracle database to the “next level” of performance, there is no better way than to migrate to Dell PowerEdge servers with Itanium 2 processors. Oracle database takes full advantage of the 64-bit Itanium 2 processor architecture

An important aspect of this architecture is memory addressing. While 32-bit processors are limited to addressing a maximum of 4 GB of RAM, 64-bit processors can address up to 16 TB of RAM. In addition, the 64-bit bus is twice as wide, allowing more data to be transferred from I/O to CPU and CPU to memory. 64-bit math can be used to calculate larger numbers with greater accuracy.

Of course, one factor holding back widespread adoption of Itanium 2 processors has been a lack of 64-bit software. That is no longer the case for Oracle databases on Windows. Windows Server 2003 offers full 64-bit OS support. Oracle has already provided a 64-bit version of Oracle9i and 10g for Windows 2003.

Dell offers the PowerEdge 7250 with up to four Itanium 2 processors. For many Oracle database applications, this server is an alternative that bears careful consideration.

CPU Resource Utilization

In general, high CPU resource utilization is a good thing. If CPU resource utilization is mostly high, then your applications are placing the load where it belongs: on the CPUs. If a high percentage of the time is spent handling I/O or network calls while the CPUs sit relatively idle, then performance will suffer.

However, CPU utilization should not consistently be pegged at 100%. If it is, you should of course first check the Performance Monitor to see what processes are causing the high CPU utilization. If the utilization is caused by unnecessary processes not related to your application or database, consider moving the processing to another server. If your database or application processes are causing the load, you should next do normal database or application tuning. Assuming your database is well tuned and CPU utilization is high, you may need to add CPUs or move to a server offering more and/or faster CPUs.

There are several factors that increase CPU resource utilization. These include:

- Large sorts done either in memory or on disk
- Heavy amounts of parsing
- Heavy navigation of complex forms
- Mathematical manipulation of row/column

To establish a baseline, measure the CPU workload at different periods during the business day. By definition, CPU workload = Average CPU capacity utilization * time interval. Consider the impact of OLTP workloads vs. batch workloads at different times during the day.

Sizing memory requirements

The component of the Oracle database architecture that has the greatest impact on performance is how Oracle handles memory usage. It is a well known fact that manipulating data in memory is several orders of magnitude faster than manipulating data on disk. Therefore, Oracle does everything possible to cache as much data as possible in memory, and to use that memory efficiently.

In terms of sizing Oracle databases, it is important to understand how Oracle uses memory, and to properly size the memory available to Oracle. The following section gives an overview of Oracle memory usage, and some guidelines for sizing Oracle memory.

Oracle memory usage

Oracle allocates two areas within system memory: the System Global Area (SGA) and the Program Global Area (PGA). The SGA contains several components, including the database buffer cache, the redo log buffers, the large pool, and the shared pool.

The database buffer cache is the largest single component, and the most critical for performance. It is used to store the most recently used data blocks. By caching the most frequently accessed data in the buffer cache, disk I/O is minimized. Starting with Oracle9i, multiple block sizes may be used within a single database, so that different data types can be stored efficiently.

The redo log buffer is a cache for redo log entries before writing to the redo log. The size of this buffer can be tuned to influence the frequency of disk writes (checkpoints) and write performance.

The large pool is an optional cache that can be used for several purposes. It can be used for session memory for multithreaded server processes. It can also be used for I/O server processes and backup and recovery operations. Use of the large pool is encouraged to prevent excessive fragmentation in the main buffer cache.

The shared pool consists of the library cache and the data dictionary cache. The library cache is used to store shared SQL. Re-use of parsed SQL can have a favorable impact on performance. The data dictionary stores “meta-data” about the structure of Oracle tables and other data types. Tuning the shared pool properly can be critical for performance.

The Program Global Area (PGA) contains several data types. This included stack space for session variables, arrays, etc. In addition, session information is stored in the PGA if dedicated Oracle Net server processes are used. (If the multi-threaded server option is used, session information is stored in the SGA.) The PGA also contains the Private SQL Area, including binding variables and runtime buffers.

Minimum Memory Requirements

The memory that Oracle allocates comes directly from the RAM that you install on the Dell PowerEdge server. The following general guidelines may be used to determine minimum amounts of memory. You will need:

- Enough memory for the SGA to fit in main memory
- Enough memory for dedicated or shared server processes to fit in main memory

- Enough memory for Oracle operations such as archiving, loading, backing up online, etc.
- Enough memory for OS activities
- Enough memory to allow for all the activities listed above, and to still avoid paging and swapping.

You should also be aware that large amounts of memory may be used for sorts and hash joins by shared server processes. Two Oracle parameters should be tuned for efficient sorts and hash joins. The `SORT_AREA_SIZE` parameter is 64 KB per user process by default, and may be reset as high as 10 MB. The `HASH_AREA_SIZE` parameter is 128 KB per user process by default, and may be reset as high as 10 MB.

Memory Usage Estimates

The following are some “rules of thumb” for sizing memory usage. The guidelines differ based on whether your database may be classified as dominantly transactional (OLTP), primarily a Data Warehouse (heavy reporting and query usage), or is a mixed-use database. Err on the high side, since the performance penalties of having inadequate memory can be severe. Significantly more RAM can be used than 8 GB, depending on the number of users and the workload.

- **OLTP**
 - 10 users: 5 MB user memory, 512 MB total system memory
 - 100 users: 50 MB user memory, 1024 MB total system memory
 - 250 users: 125 MB user memory, 2048 MB total system memory
 - 500 users: 250 MB user memory, 4096 MB total system memory
 - 1000 users: 500 MB user memory, >8192 MB total system memory
- **Data Warehouse**
 - 10 users: 50 MB user memory, 512 MB total system memory
 - 100 users: 500 MB user memory, 1024 MB total system memory
 - 250 users: 1250 MB user memory, 2048 MB total system memory
 - 500 users: 2500 MB user memory, 4096 MB total system memory
- **Multipurpose**
 - 10 users: 10 MB user memory, 512 MB total system memory
 - 100 users: 100 MB user memory, 1024 MB total system memory
 - 250 users: 250 MB user memory, 2048 MB total system memory
 - 500 users: 500 MB user memory, 4096 MB total system memory
 - 1000 users: 1000 MB user memory, >8192 MB total system memory

Memory Implications of 32-bit Processors and OS

The 32-bit memory model allows only 4 GB of memory to be accessed directly. Memory greater than 4 GB must be addressed through some form of Large Memory mapping. For 32-bit versions of Windows Server, Oracle provides access to larger amounts of memory through Very Large Memory implementations. Under Windows NT, Oracle utilizes the

Intel PSE36 device driver, which provides up to 16 GB of memory for Oracle. Under Windows 2000, up to 64 GB of memory access is provided through Microsoft's Address Windowing Extensions (AWE). With Windows Server 2003 (32-bit), the Standard Edition supports up to 4 GB of memory. The Enterprise Edition allows up to 32 GB, and the Datacenter Edition allows the maximum of 64 GB of memory.

In all cases, all components of the SGA except the database buffers must fit within 3 GB of memory. The amount of AWE memory allocated above 3 GB is equal to the `db_block_size` parameter multiplied by the `db_block_buffers` parameter.

The problem with Large Memory Mapping is that access times are much slower than for memory below the 4 GB limit. The overhead for managing Large Memory should not be underestimated. Test results show that Large Memory overhead causes SGA sizes between 4 and 8 GB to perform little better than a 4 GB SGA. If you want to use an SGA size greater than 4 GB on a 32-bit system, you should use a minimum of 8 GB of RAM.

As an alternative to large amounts of RAM per server, consider Real Application Clusters to achieve greater memory availability. Near linear scaling can often be obtained for shared memory usage. Four cluster nodes with 4 GB RAM each offer the equivalent of 16 GB of shared memory. Each server can use up to 3 GB of RAM for SGA, yielding a "virtual" SGA size of 12 GB. In this example, there is no Large Memory penalty for 32-bit servers, since each node does not exceed 4 GB of memory.

Sizing storage requirements

Storage is a necessary evil for Oracle databases. Not that there is anything wrong with storage hardware or software. In fact, Dell storage utilizes state-of-the-art hardware and software. However, when compared to Oracle data access speed in memory, disk access is many orders of magnitude slower. The simple fact is that data storage in memory is not designed to be persistent for long periods, and to this date you cannot cache large databases entirely in memory. Therefore, persistent data storage must be placed on hard disk.

Modern storage arrays offer very innovative ways to boost performance and many features that increase usability. The variety of options offered can be a little bewildering at first. The following sections are a brief overview of major data storage types.

Direct Attached Storage (DAS)

Direct attached storage comes in two flavors: internal storage and external storage. Internal storage is usually offered via disks using the SCSI protocol. Some servers utilize either ATA or SATA (serial ATA) disks. On Dell PowerEdge servers, internal disks should be configured with a RAID controller. External direct attached storage is usually a SCSI RAID array. Dell offers the PowerVault series of external SCSI arrays. Each

PowerVault array may contain up to 14 disks. In addition, Dell/EMC AX100 and CX series fibre channel storage arrays can be configured in direct attached configurations. This offers increased scalability for a single server or a two-node cluster.

Direct Attached Storage is relatively inexpensive. It is commonly used for small to mid-size Oracle databases. At most, four servers may be attached in a cluster configuration to an external array via a single connection, or two servers with redundant connections. (This is a common configuration for two-node Oracle RAC systems.) Performance can be good for Ultra320 SCSI or Fibre Channel, but flexibility and scalability are relatively low when compared to networked storage.

Storage Area Network (SAN)

Storage Area Networks feature connectivity to multiple servers via a network of cables. Most commonly, fiber optic cables are used, but copper cables may also be used. In any case, the Fibre Channel protocol is used to connect from the disk arrays to special adapters on the hosts, called Host Bus Adapters (HBAs). To connect multiple servers to the same storage array, special fibre channel switches are used. Internally, the disks use either the Fibre Channel protocol or the SATA protocol. Fibre Channel disks offer higher performance, while SATA disks are more economical.

One of the benefits of a SAN is great flexibility on the number and type of hosts attached to storage array. They offer the highest storage performance on the market today. SANs are commonly used for Oracle databases of all types.

On the negative side, SANs are relatively expensive compared to DAS. They also require specialized training in storage management.

Dell offers the entry level AX100 SAN with SATA disks (up to 12 disks). Most production systems use a CX series SAN, which can use Fibre Channel or SATA disks. The CX300 is shipped with 15 disks, and can be expanded to 60 disks. It is commonly used for connecting a small number of servers. The CX500 can be expanded to 120 drives and the CX700 can be expanded to 240 drives. These storage systems offer the opportunity for enterprise consolidation of storage for multiple servers.

Network Attached Storage (NAS)

Network Attached Storage offers a way to simplify network storage, often at a lower cost. Instead of relatively expensive Fibre Channel cables and HBAs, NAS uses standard Ethernet cables and standard NICs. In terms of the Ethernet network, a Gig-E LAN or VLAN is preferred. NAS offers high flexibility, getting closer to true “plug-in” network storage. There is a low training threshold for NAS implementations.

The basic concept of providing shared storage is similar in concept to Windows Drive Shares or UNIX NFS drives. However, the original versions of these protocols did not

offer the performance and reliability that Oracle requires. NAS devices utilize a stripped-down OS that is designed to offer dedicated, high performance NAS support. For example, Windows 2003 Server offers native NAS support. An example of a Windows 2003-based NAS device is the Dell PowerVault 745N NAS server.

Despite its flexibility and ease of use, NAS tends to suffer from a performance penalty. This is due to the requirement for conversion from native SCSI or ATA block format on the disks to TCP/IP for transmission. High end NAS devices can compensate for performance overhead to achieve decent performance. One of the ways that this is done is to place a NAS gateway in front of a SAN storage device. This leverages the back-end performance of the SAN, with the flexibility of the NAS gateway. This is the approach used for the Dell PowerVault 770N and 775N NAS servers. They may be used in combination with any Dell external RAID array or SAN (of course, this adds some expense, depending on the SAN device chosen).

NAS devices are commonly used for small to moderate size databases or non-mission critical databases. They are less commonly used for mission-critical databases, due to the performance overhead. Dell PowerVault 770N or 775N NAS gateways may be used for small to moderate sized Oracle databases.

Commonly Used RAID Types

RAID stands for Redundant Array of Inexpensive Disks. Since data storage is so important for Oracle databases, it is important to understand the available RAID levels and their performance and availability implications. In general, hardware based RAID is preferable to software based RAID, due to lower performance overhead on the Oracle server.

RAID 0

With RAID 0, data is striped across disks, but no data protection is offered. This is not truly RAID, since there is no redundancy. RAID 0 offers high performance, since there is only one Physical I/O per logical read or write. However, any error on any disk causes the whole RAID array to fail. RAID 0 can be appropriate for temp or scratch space.

RAID 1

RAID 1 consists of pairs of mirrored disks. The two disks together offer a single logical disk. If one disk fails, the other disk takes over until the disk is replaced. RAID 1 is relatively expensive since for every two disks purchased, only the storage capacity of one disk is usable.

RAID 1 offers good sequential read and write performance. However, it offers only moderate random read and write performance, due to lower spindle count than multi-disk striped configurations. RAID 1 is commonly used for redo logs and archive logs, due to the sequential reads and writes performed by logging processes. RAID 1 redo and

archive logs are particularly good for OLTP performance. Note that each redo thread should be separated on separate disk pairs to maintain sequential access. Since redo logs rarely fill an entire disk, using a RAID 1 disk pair for redo logs balances the need for higher sequential I/O performance with the less efficient storage utilization that can be afforded to transient files.

RAID 5

RAID 5 utilizes a different algorithm for redundancy. Instead of mirroring, parity-based error correction is striped across all disks in the array. If data on any disk (or a whole disk) is lost, the data can be recalculated from the remaining disks. The net effect is the number of accessible disks is one less than the total disks in the RAID group. This is a relatively cost effective form of RAID.

After a disk failure, a replacement disk can be automatically synched with the remaining disks. However, please note that rebuild speeds may be slow. Performance will be degraded during rebuilds, although the RAID group continues to function.

RAID 5 tends to suffer from a write performance penalty. Each logical write requires four I/Os: read the disk stripe, read the parity, write the disk stripe, calculate and write the new parity. Hardware cache can be used to mitigate the write penalty to some extent. All Dell storage offers write cache to optimize RAID 5 write performance.

RAID 5 offers good performance for read-only databases such as Data Warehouses, due to data striping across multiple disks (reads are performed in parallel). RAID 5 is generally considered poor for OLTP databases, due to the write penalty. However, it is the least expensive form of protected RAID storage, and cost considerations often trump performance considerations. If you have to use RAID 5 for OLTP databases, make sure you maximize the amount of write cache.

RAID 10

RAID 10 (often called RAID 1 + 0, or RAID 0 + 1) offers a combination of mirroring and striping techniques. Each individual disk is mirrored, but the data is striped across multiple disks. A minimum of 4 disks is required for a RAID 10 group. There is a hardware specific maximum for each array (e.g., 16 disks for the Dell/EMC CX series). RAID 10 groups are commonly implemented as a 10 disk RAID group.

RAID 10 offers the highest write performance for any form of protected disks. It is commonly used for Oracle data files. RAID 10 is especially good for OLTP databases.

Hard Drive Performance

Despite all of the performance improvements that you get with disk striping, disk read and write caches, and other techniques, I/O performance still comes down to the

performance on each individual disk. If I/O performance is poor for each disk, no amount of technological “trickery” will hide the fact that I/O is generally slow.

In this context, it is important to note that performance for a single disk is sensitive to percent utilization. High percent utilization causes excessive queuing, increasing response times (latency) exponentially. The “Knee” of the curve is at 80 % utilization. Above 80 % utilization per disk, I/O performance gets exponentially worse.

To avoid this problem, do not try to drive too much I/O through each disk. A typical 10,000 rpm drive can handle 85 I/Os per second at an 80 % utilization rate. A typical 15,000 rpm drive can handle up to 100 IOPS at an 80 % utilization rate. Given these guidelines, it is easy to predict how many disks you will require to meet a given target for I/Os per second (remember to double the final disk count for mirrored drives).

Hard disks required to meet IOPS requirements

- 100 IOPS: 10K – 2 drives, 15K – 1 drive
- 1000 IOPS: 10K – 12 drives, 15K – 10 drives
- 5000 IOPS: 10K – 59 drives, 15K – 50 drives
- 10,000 IOPS: 10K – 118 drives, 15K – 100 drives
- 20,000 IOPS: 10K – 236 drives, 15K – 200 drives
- 25,000 IOPS: 10K – 295 drives, 15K – 250 drives

Recommendations for Physical Oracle Data Layout

There are many published recommendations for how to design the physical layout of an Oracle database for optimum performance. Unfortunately, if you read all of the available literature, you will find that every recommendation is contradicted somewhere else. The problem is that storage technology has changed greatly with time. What was once good advice is now a guarantee for sub-optimal performance.

But old storage advice never dies. It is not uncommon to hear Database Administrators quoting storage performance tips that were current at about the time that bell-bottomed pants were the rage. For example, “everybody knows” that Oracle tables and indexes should be placed on separate disks. The basic concept is that data and indexes may be searched simultaneously. If they are on the same disk, there could be contention. This may have been sound advice in the pre-RAID era. On any modern storage array, read cache and disk queuing mechanisms allow several processes to access the same disks efficiently. Even Oracle has stated that separating data and index does not guarantee improved performance.

In addition, this advice completely ignores the whole concept of disk striping. The “separate all data types” philosophy tends to place all data on a great number of RAID 1 disk pairs. As noted above, you should not try to drive too much I/O through a single disk, or even a RAID 1 disk pair. Disk striping is a way to avoid this problem.

Oracle now recommends the SAME strategy (Stripe and Mirror Everywhere). In this approach, you seek to stripe all data types across large RAID 10 arrays. The higher the spindle count in the RAID 10 array, the greater the supported I/Os per second. This is due to the opportunity to read data in parallel across the stripe set. I/O performance is directly related to the stripe width.

One thing has not changed through time. The old advice was that Oracle database performance benefits from multiple small disks. It is still better to use multiple relatively small disks rather than a few large disks. Given the extremely large capacities of individual disk drives these days, it is easy to ignore this advice. However, eight 36 GB disks in a RAID 10 configuration will offer greater performance than a mirrored pair of 142 GB drives in an Oracle database, “hands down”.

To summarize, below are a few up-to-date guidelines for Oracle database physical layout that, if followed, can have a substantial impact on Oracle database performance:

- **Place all schema objects on one or more large RAID 10 arrays**
 - The performance advantage of hardware striping far outweighs any disadvantage caused by combining data and index, or any other schema types.
- **Place redo logs on RAID 1 or RAID 10 arrays**
 - Place each log thread on a separate RAID 1 pair. This allows sequential read optimization algorithms on the storage array to be activated.
 - An alternative is to include redo log files with the main RAID 10 disk set. This is useful where log write activity is low, or where there are a limited number of disks.
- **Place archive log files on RAID 1, RAID 10, or RAID 5 disks**
 - Both for performance and availability reasons, archive logs should be physically separated from other data types.

Oracle 10g and ASM

With Oracle Database 10g, Oracle introduced Automated Storage Management, or ASM. With ASM, Oracle software takes on the responsibility for managing disk mirroring and striping. You can literally give Oracle a group of raw disks without any hardware RAID. Oracle software will handle mirroring and striping for you.

You simply specify how many disks that you want to include in the ASM group and the desired level of disk mirroring (you can even have three-way disk mirroring, or elect to use hardware mirroring). Striping is handled on a per file basis, so that each table, redo log, etc. may be striped across a different number of disks. Oracle analyzes the most efficient striping technique for each individual object. ASM is designed to make the DBA's life easier. You no longer have to worry about how to lay out Oracle data. Oracle software can take care of it for you, automatically.

There are a variety of ways to take advantage of Oracle ASM, including blending ASM with storage hardware. One good strategy is to allow ASM to handle striping, but to implement disk mirrors with hardware. This is desirable because hardware mirrors are generally more efficient than software mirrors. In this scenario, you would give ASM a group of RAID 1 pairs to manage, and turn off the ASM mirrors. ASM would then stripe data across all of the RAID 1 disks, in a similar fashion to RAID 10.

Some DBAs will want to maintain greater control over their storage, not completely trusting a software algorithm to manage all of their storage. As a compromise, instead of providing ASM with RAID 1 pairs, it is also possible to provide larger RAID 10 or RAID 5 groups as the base for ASM to manage. ASM will then stripe data across all of the RAID groups it is given. Experienced DBAs that don't feel a need for automated storage management assistance may elect not to use ASM at all. In theory, pure hardware striping should have a performance edge over software striping.

You should use your own judgment over which strategy to follow. If possible, test the real-world performance of your strategy. Thanks to Oracle10g, you now have more storage management options than ever before.

Example PowerEdge Server Configurations for Oracle on Windows

The following are some example PowerEdge server and Dell storage configurations for different Oracle database sizes and uses. When in doubt, size larger than these examples. The storage configurations should be considered as minimum configurations for capacity and performance purposes.

As an alternative, if you use Oracle10g with ASM, you can use the same number of disks as in the example, configured as hardware RAID 1 groups with ASM striping.

- **Small database for testing or development**
 - PowerEdge 750 or 1750
 - 1 or 2 CPUs
 - 512 KB - 1 GB RAM
 - < 1 GB SGA
 - Windows 2000 Server or 2003 Server (standard edition)
 - Oracle9i (9.2.0.1) Standard Edition or Oracle10g Standard Edition One
 - 10 to 100 users
 - 2 - 4 internal 36 or 73 GB SCSI disks – RAID 1
- **Moderate size OLTP database**
 - Dell PowerEdge 2650 or 6650
 - 2 or 4 CPUs
 - 4 GB RAM
 - ~ 3 GB SGA

- Windows 2000 Server or 2003 Server (Standard or Enterprise Edition)
- Oracle9i (9.2.0.1) Standard Edition or Oracle10g Standard Edition One (2 CPUs) or Standard Edition (4 CPUs)
- 100 to 500 users
- OS and Oracle binaries on 2 internal 36 GB SCSI disks- RAID 1
- External RAID array – PowerVault SCSI, CX300 SAN, or PowerVault NAS
 - 10 disk (36 GB, 15K rpm) RAID 10 for Oracle data
 - RAID 1 pair of 36 GB disks for redo logs
 - RAID 1 pair of 36 GB disks for archive logs
 - One hot spare disk (except for PowerVault SCSI array)
- **Data Warehouse server**
 - Dell PowerEdge 7250
 - 4 Itanium 2 CPUs
 - 12 GB+ RAM
 - 8 GB+ SGA
 - Windows 2003 Server (Enterprise Edition, 64-bit)
 - Oracle9i (9.2.0.2) Enterprise Edition (64-bit) or Oracle10g Enterprise Edition (64-bit)
 - 500 to 1000 users
 - OS and Oracle binaries on 2 internal 36 GB SCSI disks- RAID 1
 - External SAN – CX300 with one DAE2 expansion unit (20 disks)
 - Two 10 disk RAID 5 groups with 72 GB disks - data and redo
 - Two 4 disk RAID 10 groups with 72 GB disks for ETL and staging
 - One hot spare disk
- **ERP system on RAC**
 - Four Dell PowerEdge 6650 or 7250 servers in a RAC cluster
 - 4 Xeon or Itanium 2 CPUs on each node
 - 4 GB RAM on each node
 - Windows 2000 Server or 2003 Server (Standard Edition, 32-bit) or 2003 Server (Enterprise Edition, 64-bit)
 - Oracle9i (9.2.0.1) Enterprise Edition or Oracle10g Enterprise Edition (32-bit), Oracle9i (9.2.0.2) Enterprise Edition or Oracle10g Enterprise Edition (64-bit)
 - Gig-E network with switch(es) for private interconnect
 - OS and Oracle binaries on 2 internal 36 GB SCSI disks- RAID 1
 - External SAN – CX300 with two DAE2 expansion units (30 disks)
 - Two 8 disk RAID 10 groups with 36 GB disks - data on Oracle Cluster File System (OCFS)
 - 4 disk RAID 10 group with 36 GB disks for redo logs
 - 4 disk RAID 10 group with 36 GB disks for archive logs
 - One 5 disk RAID 5 group for ERP data (NTFS)
 - One hot spare disk

Deliverables

There is a payoff to going to the trouble to define and rigorously follow a formal sizing methodology. You will be able to promise and produce specific deliverables for your sizing project. These deliverables will contain estimates and projections, but you will also be able to justify your sizing predictions with well documented data. These deliverables include the following:

Justified Server Sizing Recommendation

This deliverable report should include the number of servers and number of Oracle instances. The server classes should be identified, if not the specific servers. The number of CPUs and the memory requirements for each server should also be listed. Documentation justifying the server selections and sizing parameters should also be included in the report.

It is also appropriate to include in this report any recommendations for high availability and scalability options. The impact on Service Level Agreements should also be noted.

Justified Storage Sizing Recommendation

The choice of storage is critical to the performance of the Oracle database. Therefore, a separate report detailing the storage recommendation should always be prepared, even if the storage is as simple as local disks on the server (Direct Attached Storage). The justification for the storage recommendation should be derived from the customer storage capacity requirements, and your study of I/O patterns on comparison systems.

The report should specify the type of storage: DAS, SAN, or NAS. For Direct Attached Storage, you should further distinguish between local SCSI or (S)ATA disks and a dedicated SCSI, (S)ATA, or Fibre RAID array in a separate enclosure. For a Storage Area Network, any switching components should be identified. A map relating servers to storage arrays should also be included, if possible down to the individual Logical Unit device (LUN) level. For Network Attached Storage, a map of the file system types and network directory or mount points should be included.

Within each storage array, the size and number of disks should be specified. The rotational speed of the disks should also be noted, along with documented latency times. The RAID type of each RAID group, the number of disks included, and the resulting storage capacity should be documented.

Budgetary Sizing Recommendation

It is often useful to prepare a separate report that spells everything out in budgetary terms. This report will be most useful for Accounting and Management. In this report, not only the total price, but the timing of hardware and software acquisition should be considered. Other useful factors to include are the trade-in value of legacy equipment, lease vs. buy options, cost of support contracts, and depreciation through time.

Best Practices

If you have followed the previous recommendations, you started the project by collating known best practices. You then modified and added to this document as the project proceeded. Your last deliverable is this Best Practices document. This document should be “evergreen”, with plans in place to periodically refresh the published documentation.

Summary and Best Practices

Proper sizing is the key to a successful implementation of an Oracle database. In theory, sizing is an iterative process. You should have the freedom to experiment with different hardware and software configurations until you achieve the perfect setup. However, in the real world, we all know that you basically have one chance to size your system properly. Any changes other than minor growth can lead to expensive un-budgeted purchases of additional hardware, or can even lead to project failure.

You can avoid sizing catastrophes by using a systematic approach to sizing. The key steps outlined in this paper are:

- Define customer requirements
- Collect customer information
- Analyze system data
- System determination
 - Predict total CPU use
 - Predict total memory use
 - Estimate total I/O load
 - Estimate storage resources required for capacity and performance
 - Use general guidelines for server and storage sizing
 - Compare system requirements to relevant benchmarks
 - Compare system requirements to similar implementations
- System test (optional)
- Implement

With disciplined application of this systematic approach, along with the use of the specific sizing guidelines listed in this paper, you should be able to properly size your Oracle database system. You will be able to recommend specific Dell PowerEdge servers with the correct Intel processors, the correct amount of RAM, and an appropriate storage capacity with a high-performance configuration. You will be able to choose the right version of Microsoft Windows Server, along with the Oracle database edition that meets your needs. Finally, you will be able to fully justify your recommendations to your peers and management.

References

“Dell and Oracle”; <http://www.dell.com/oracle>

“Dell Server Solutions”; <http://www.dell.com/servers>

“Oracle7 Sizing and Capacity Planning Guide for Alpha AXP OSF/1 Systems”; Author: Scott Foote, Contributors: Lakshmi Arya-Bakshi, Judy Clifford, Jennifer Stave; May 1994, Oracle Corporation, Part A19178; www.allenhayden.com/cgi-bin/getdoc.pl?file=vaxperf.pdf

“Oracle9i Release 2 Database Architecture on Windows”; Author: David Colello, Oracle Corporation, Contributing Author: Alex Keh, Oracle Corporation; April 2003, Oracle Corporation; http://otn.oracle.com/tech/windows/rdbms/9i2_windb.pdf

“Oracle10g: Leveraging Windows Server 2003”, Authors: David Colello & Beldalker Anand, Server Technologies Division, Oracle Corporation; June 2004, Oracle Corporation;
http://www.oracle.com/technology/tech/windows/wp/Oracle10g_on_Windows2003-TechEd.zip

“Oracle Database 10g Architecture on Windows”; Author: David Colello, Oracle Corporation, Contributing Author: Alex Keh, Oracle Corporation; December 2003, Oracle Corporation;
http://www.oracle.com/technology/tech/windows/rdbms/oracle_db_10g_architecture_on_windows_twp.pdf

“Oracle Database 10g on Windows: Architecture for Performance”; Authors: David Colello, Oracle Corporation, Steven Fluge, Intel Corporation; September, 2003;
http://download.oracle.com/owsf_2003/40171_colello.ppt

“Oracle Database on Windows FAQ”; 2004, Oracle Corporation;
<http://otn.oracle.com/tech/windows/faq.html>

“Oracle Performance Tuning”; Authors: Edward Whalen, Mitchell Schroeter, Performance Tuning Corporation; Addison-Wesley Pub Co; 1st edition (April 16, 2002); ISBN 0672321467

“Sizing Oracle 8i on PowerEdge Servers”; Author: Scott Reichmanis, Dell Corporation; Dell Power Solutions Magazine, Issue 1 2001;
http://www1.us.dell.com/content/topics/global.aspx/power/en/ps1q01_oracle?c=us&cs=555&l=en&s=biz

Larry Pedigo (lpedigo@perftuningl.com) is a Performance Consultant for Performance Tuning Corporation. He has been involved with Oracle RAC deployment and configuration for the past year. He is currently working on Oracle 9i and 10g RAC deployments, Oracle 10g Application Server deployments, as well as performance optimization on all versions of Oracle. Larry has a Bachelors degree in Geology from Texas A&M University, and over 15 years experience in the IT field.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© Copyright 2004 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell. November 2004

Dell is a trademark of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks or the names of their products. Dell disclaims proprietary interest in the marks and names of others.

Copyright © 2004, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

© 2004 Microsoft Corporation. All rights reserved.

Microsoft, Microsoft Windows Server, and Microsoft Windows Storage Server 2003 are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.