Oracle Application Server Discoverer 10g (9.0.4) Capacity Planning Guide

An Oracle White Paper
January 2005

ORACLE

Oracle Application Server Discoverer 10g (9.0.4) Capacity Planning Guide

Oracle Application Server Discoverer 10g (9.0.4) Capacity Planning Guide

OVERVIEW

This paper demonstrates the exceptional scalability of the Oracle Application Server Discoverer 10g (9.0.4) solution. Discoverer achieves near-linear scalability with the CPU processing power and memory of the server.

This paper provides sizing guidelines to help consultants and system administrators successfully plan and implement their system to leverage the scalability of Discoverer.

Servers may be added to a clustered Discoverer deployment to accommodate any number of users. For example, if a user population supported on 2 servers grows by 50%, an additional server can be added to accommodate the additional user population while maintaining the current performance. This flexible approach provides insurance in case of sudden shifts in business needs – servers may be re-tasked quickly without interrupting the availability of Discoverer.

INTRODUCTION

Capacity planning of the middle tier for OracleAS Discoverer is based on:

- the number of users concurrently accessing the system
- the particular mix of Discoverer Plus and Discoverer Viewer they are using to access their data
- the size of the data set and complexity of the layout of the worksheets being executed.

The testing methodology used to determine the sizing formula was to run various scenarios against systems that were limited in either memory or CPU capacity. The scenarios were repeated with increasing numbers of users until a significant drop off in response time was detected. The test results were then used to determine the incremental memory and CPU required per user.

The capacity planning formula provided in this white paper serves only as general guidance for your Discoverer implementation. It is strongly recommended that you follow the methodology explained in this white paper and conduct your own

scalability tests using your own worksheets and data to more accurately achieve the capacity planning goals for your business.

CAPACITY PLANNING GOALS

The goal of capacity planning for Discoverer is to determine the hardware needed in the middle tier to support a given number of concurrent users using a particular mix of Discoverer Plus and Viewer and a range of data set size and worksheet complexity while maintaining a reasonable response time to user actions.

Number of concurrent users

Concurrent users for Discoverer are those users who are actively running queries and analyzing data.

When planning your Discoverer system you first need to get a realistic estimate of how many users will be accessing the system at the same time. Be sure to use realistic estimates of the number of concurrent users when using these sizing guidelines. In most cases the number of concurrent users is a fraction of the entire user population. Experience with customers has shown that the number of concurrent users is typically 10%, and often less, of the entire Discoverer population.

Second, you need to determine how many of those users will be using Discoverer Plus and how many will be using Discoverer Viewer. Determining whether users should use Plus or Viewer should be based on the requirements each user has and the functionality provided in each interface. It is expected that users who need to create and modify worksheets will use Plus. Users who only need to view and analyze data are expected to use Viewer. However, you may also want to consider resource requirements when determining how many users will use Plus and how many will use Viewer. Discoverer Plus is a Java based client that provides a high degree of interactive manipulation, query modification, and report building. Plus tends to consume more network bandwidth but less middle-tier resource since the rendering and display of the user interface is on the client machine. Discoverer Viewer is a pure HTML client designed to view reports and do some analysis. Viewer tends to consume less network bandwidth and more mid-tier resources since the rendering and display of the user interface runs on the middle tier.

Expected response time

See the *Oracle Database Data Warehousing Guide* for guidance on how to design your database for Business Intelligence queries.

Since Discoverer is an end user, ad-hoc query tool it is difficult to predict response times for the initial querying of data. Query response times can be greatly impacted by database design, the amount of data being returned, or the layout of data in the worksheet. Having said that, with a well designed database, reasonable data set and worksheet layout, users can expect query response times anywhere from a second or two to less than 10 seconds even against large data sets. Once the data is queried, near instantaneous response time is expected as users drill and pivot through the data.

Data set size and worksheet layout

You will need to understand the relative size of the data set that users will be querying and the complexity of the layout in their worksheets. The larger the data set and the more complex the worksheet layout, the greater the demand on middle tier resources. Keep in mind that Discoverer is an ad-hoc data analysis tool for end users so it is expected that users will create worksheets that display a reasonable, humanly readable result.

By following the "Optimizing OracleAS Discoverer performance and scalability" guidelines in the *Oracle Application Server Discoverer Configuration Guide 10g (9.0.4),* users should be able to create worksheets that perform well, display data in a reasonable layout, and also put only reasonable loads on the middle tier.

It is recommended that users follow the guidelines in chapter 10 "Optimizing OracleAS Discoverer performance and scalability" of the *Oracle Application Server Discoverer Configuration Guide 10g (9.0.4).* By following these guidelines users should be able to create worksheets that not only perform well and display data in a reasonable layout, but also put only reasonable loads on the middle tier.

RESOURCES

See the *Oracle Application Server Discoverer Configuration Guide 10g (9.0.4)* for details on how to use Discoverer with Web Cache. See the *Oracle Discoverer Administrator Administration Guide (9.0.4)* for details on using summary folders and enabling scheduling. See the *OracleAS Discoverer Plus User's Guide 10g (9.0.4)* for details on scheduling workbooks.

The outcome of capacity planning will be a recommendation that covers the physical resources required to meet your capacity planning goals. The recommendation will be the number of processors and amount of physical memory (RAM) needed to support the number and mix of concurrent users you expect to have on the system. If you are able to use Web Cache, you may be able to reduce the middle tier requirements for your user population. If you are able to use summary folders or scheduled workbook results, you may be able to achieve better response times using your Discoverer system.

Number of processors

The sizing recommendation for number of processors is based on the class and number of processors used in the chosen benchmark environment.

There is no universally accepted industry standard for comparing the "power" of different CPUs. However, many major vendors do publish popular benchmarks, such as SPEC CPU 2000. For more information visit: http://www.spec.org. The SPEC website provides benchmarks, results and information regarding the correct use and interpretation of those results.

In the absence of data for other processors and operating systems you may make an estimation based on no less than the total MHz determined by this sizing guideline.

Memory requirements

The sizing recommendation for amount of memory (RAM) is based on the chosen benchmark environment.

Like processor comparisons, there is no accepted industry standard for comparing memory of different machines. In the absence of comparison data, you may make an estimate of no less than the total RAM determined by this sizing guideline.

SIZING

Methodology

Users should follow the methodology explained in this white paper and conduct scalability tests with their own data and environment to more accurately achieve the capacity planning goals for their business.

The testing methodology used to determine the sizing formula was to run various scenarios against systems that were limited in either memory or CPU capacity. There are two scenarios – Simple and Complex – based on the complexity of the tasks completed during the scenario. The Simple scenario represents a casual user who typically views existing reports and performs limited additional analysis. The Complex scenario simulates a more sophisticated user who performs deeper analysis and uses more product features. (see Appendix A for details). The scenarios were repeated using Plus or Viewer with increasing numbers of concurrent users until a significant drop off in response time was detected. The test results are then used to determine the incremental memory and CPU required per user.

For example, to determine the incremental memory requirements per user of Discoverer Plus running the Complex Scenario, a machine is physically configured to have 4 CPUs but only 1 GB of memory. Configured this way the machine has plenty of processing power but is memory bound. Then, using Plus, the Complex scenario is run seven times with a range of users starting with 10 and going up to 80. For each run, average response time is recorded and memory usage is measured. We observe for the runs of 10 through 60 users, the average response time remains relatively flat, only increasing ~ 1 second over the entire range from 10 to 60 users. However, average response time dramatically increased for the runs with 70 and 80 users by 10 and 15 seconds respectively. So we look at the average memory consumption for the run with 60 users, when response time was still acceptable, and observe that incremental memory consumption per user is 10.14 MB. Additional tests are run on a machine with 4 CPUs and 2 GB memory and these tests confirm that the factor is accurate. This number becomes the factor for incremental memory needed to support each user running the Complex Scenario using Plus and still maintain acceptable performance.

In a similar way the CPU requirements for incremental users are determined. Tests are run with a range of users for all four combinations of Viewer and Plus for the Simple and Complex scenarios.

Architecture

Architecture used in the testing is the standard architecture recommended in the product documentation – consisting of a middle tier with one machine running OracleAS Discoverer and another machine running OracleAS Infrastructure.

There was one machine running Oracle Database Enterprise Edition and a PC client machine simulating multiple users.

Hardware and Software Employed

Hardware

Host	CPU class	CPUs	MHz/CPU	RAM
Application Sever	UltraSPARC II	2 or 4	450	1, 2, or 4 GB
Infrastructure	Intel P4	1	2400	2 GB
Database	UltraSPARC II	2	450	4 GB
Client	Intel P4	1	1700	512 MB

Software

Host	Product version	JVM	OS
Application Sever	Oracle AS 10g Discoverer (9.0.4)	1.4.1_03	Sun OS 5.9
Infrastructure	Oracle AS 10g Infrastructure (9.0.4)	1.4.1_03	Red Hat Linux AS 2.1
Database	Oracle Enterprise Edition 9.2	N/A	Sun OS 5.8
Client	Mercury LoadRunner 7.6	JInitiator 1.3.1	Windows XP

The network between the client, application server, infrastructure server, and database server is a Fast Ethernet, Full-duplex, switched 100Mbps network. All machines are physically close to each other, within the same subnet and building.

Configuration or Tuning Parameters

The Application Server Discoverer machine was optimally configured to run Discoverer. For example, the only OC4J process running was the OC4J_BI_Forms instance. Application Server Control was not running. Oracle Application Server and Discoverer configuration parameters were set according to the guidelines in chapter 10 "Optimizing OracleAS Discoverer performance and scalability" of the *Oracle Application Server Discoverer Configuration Guide 10g (9.0.4)*. For additional details see the appendix.

Derived Formula

Sizing Calculator for Sun Solaris on UltraSPARC II 450 MHz

	User Supplied	Incremental MHz factor	Incremental MHz req'd	Incremental RAM factor	Incremental RAM MB req'd
# of Simple Plus Users		* (4.1) =		* (8.1) =	
# of Complex Plus Users		* (8.3) =		* (10.2) =	
# of Simple Viewer Users		* (18.6) =		* (10.1) =	
# of Complex Viewer Users		* (28.1) =		* (20.3) =	
Total Concurrent Users		Total MHz		Total RAM	

Note that an additional 260 MB RAM is required for the base Oracle Application Server processes – HTTP Server, OC4J, etc.

Sizing Example

Example Sizing Calculator for Sun Solaris on UltraSPARC II 450 MHz

	User Supplied	Incremental MHz factor	Incremental MHz req'd	Incremental RAM factor	Incremental RAM MB req'd
# of Simple Plus Users	20	* (4.1) =	82	* (8.1) =	162
# of Complex Plus Users	10	* (8.3) =	83	* (10.2) =	102
# of Simple Viewer Users	50	* (18.6) =	930	* (10.1) =	505
# of Complex Viewer Users	20	* (28.1) =	562	* (20.3) =	406
Total Concurrent Users	100	Total MHz	1657	Total RAM	1175

Note that an additional 260 MB RAM is required for the base Oracle Application Server processes – HTTP Server, OC4J, etc.

In this example, there are 20 Simple Plus users, 10 Complex Plus users, 50 Simple Viewer users, and 20 Complex Viewer users – a total of 100 users.

The calculator estimates this mix of users requires 1657 MHz of UltraSPARC II CPU processing power and 1435 MB of RAM (1175 MB for the users + 260 MB for the base Application Server) above whatever processor and memory is needed for the OS and any other applications. This user population could be supported on a Sun E420R with four 450 MHz processors (`CEILING (1657/450)`) and 2 GB RAM (`CEILING (1435/1024)`).

Limitations

The factors in this calculator are based on the specified test scenarios that are intended to simulate real world workloads. Results may vary with different application databases, different hardware, and different workloads. In particular,

differences in CPU architecture and speed may affect the results. Within the same CPU architecture, it may be reasonable to equate CPU power to the number of CPUs multiplied by their speed (# of CPUs x clock speed). For example, two 450MHz UltraSPARC CPUs may provide roughly the same power as one 1050MHz UltraSPARC CPU. However, power comparisons between different CPU architectures may not be valid.

Given the degree to which database design, worksheet implementation, network, operating system, and processor architecture differences can impact performance and scalability, there is no substitute for running your own scalability tests in your environment.

Having said that, these sizing guidelines can be used as a starting point for planning your hardware needs for your Discoverer environment. To create a more conservative estimate you could weight your user population to include more Complex Viewer users than actually anticipated. To create a more aggressive estimate, you could weight your population to include more the Simple Plus users than actually anticipated.

In either case you can depend on the near linear scalability of Discoverer and add more low cost servers as your needs grow.

Upgrading Versus Adding Servers

These tests show that OracleAS Discoverer (9.0.4) is primarily CPU-bound. If a server's performance dramatically decreases or cannot support additional users, CPU resources should be checked first, memory resources second, and the database performance last.

1. Monitor the CPU usage, free memory available, disk space available, and disk I/O of the server.

2. If CPU usage appears high (almost always above 90%) while memory and disk are not heavily used, then add CPUs by upgrading the server or adding another server. Use the calculators to determine the hardware required for the additional server(s).

3. Otherwise, if CPU usage appears moderate and memory usage and disk I/O appear very high, the server has insufficient memory (real and virtual). Virtual memory operations also account for some of the CPU usage. Add more real memory to the system and continue to observe the system.

4. However, if CPU usage, memory usage, and disk I/O all appear high, the server is completely undersized. Use the calculators for a sanity check on the server configuration and plan to either upgrade the server or add additional servers.

Refer to the Oracle Database documentation library for a detailed description of how to monitor database performance.

5. If none of the server's resources appear to be the limiting factor, monitor the database server(s) that users are querying. Discoverer must retrieve data from the database to respond to query requests - the application server may be waiting for data while the database server is saturated by query activity.

CONCLUSION

The results from the benchmark tests show that Oracle Application Server Discoverer 10g is a robust BI platform with near-linear scalability.

This white paper presents the methodology for conducting capacity tests and an illustrative capacity planning calculator for OracleAS Discoverer 10g. Since the benchmark systems and test scenarios may not reflect your environment, you should follow the methodology and conduct your own capacity planning tests with your systems and data to derive the most accurate resource requirements for your business needs. You can use the calculator to derive a preliminary capacity planning estimate for your implementation.

Because of its proven scalability, customers know with confidence that Oracle Application Server Discoverer 10g will meet their future demands, and business decisions concerning future growth - such as purchasing hardware or upgrading existing systems - can be made safely.

APPENDIX

Appendix A – Test Scenario Details

The test scenarios were recorded using an industry standard load-testing tool to simulate a multi-user environment. The scenarios are recorded in the exact sequence of steps defined in the scenario flows below. A randomized think time of 45-75 seconds is provided between various steps in the workload. Think time randomizes the tasks being performed at a snapshot in time to simulate real world usage. Think time is not included in the transaction response time results.

Users are launched simultaneously in groups. Each group consists of up to 25 users. Within a group, 5 users are initialized every 30 seconds (this consists of login and opening a default worksheet). Simultaneous login and opening of a workbook by all users does not represent a real world environment, thus, for the first iteration through the scenario, users in a group do not start running tasks until all the users in that group are initialized. This ensures that login and opening the workbook, which are part of the test initialization, do not simultaneously impact performance. Login and the opening of a workbook *are* included in the results of all subsequent iterations though. After initialization, one user begins the full test scenario every 30 seconds.

Once all users finish executing the scenario, the transaction response times reported by the tool for the steady state are analyzed. The steady state is defined as the time during which all the users are running through the scenario - no users are

still initializing and no users have finished the scenario. Hence, maximum usage of resources will only be seen during the steady state.

The business area used for the test scenarios is the Video Stores Tutorial. It is based on a small star schema with 87,489 rows in the fact table with 5 facts. There are three dimension tables with 20, 141, and 912 unique values.

Two scenarios were used for the scalability tests, based on the complexity of the tasks done during the scenario.

Simple Scenario flow

The simple scenario consists of tasks performed on a workbook with two worksheets:

- Table – table worksheet with 125 rows display, aggregated from 358 rows of data in the database.
- Crosstab – crosstab worksheet with 1 page item (6 values), 2 items on the top axis (12 values), and 3 items on the side axis (20 values). 240 cells displayed, aggregated from 1,181 rows of data in the database.

The following flow of tasks are performed in this scenario:

  - Connect
  - Open workbook (executes Table sheet by default)
  - Page through results of the table (5 times)
  - Change page axis item
  - Switch to Crosstab sheet
  - Pivot an item
  - Change page axis item
  - Drill once
  - Disconnect

The scenario uses a think time of 45-75 seconds between each task. This scenario represents a casual user who typically views existing reports and performs limited additional analysis. The worksheets involved do not return large amounts of data. This is appropriate as the casual user views reports that should already be focused and refined for frequent use. This scenario represents the "simplest" set of tasks that a real user performs.

Complex Scenario flow

The complex scenario consists of tasks performed on a workbook with four worksheets:

- Big Crosstab – crosstab worksheet with 2 page items (4 and 3 values each), 2 items on the top axis (10 values), and 5 items on the side axis (280 values). 2800 cells displayed, aggregated from 4200 database rows.
- Avg Crosstab – crosstab worksheet with 1 page item (3 values), 2 items on the top axis (12 values), and 4 items on the side axis (18 values). 216 cells displayed, aggregated from 211 database rows. Some cells show no value.
- Big Table – 100 rows displayed per page 1573 rows total, 1573 rows from the database (no aggregation)
- Small Table – 53 rows displayed based on 53 rows from the database (no aggregation)

The following flow of tasks are performed in this scenario:

  - Connect
  - Open workbook (executes Big Crosstab by default, returns 1050 rows)
  - Drill on the side axis from Product Category to Product Description for all values (now 4200 rows total)
  - Drill on the top axis from Calendar Quarter to Calendar Month for all values (now 12,436 rows total)
  - Change the page axis item
  - Switch to Avg Crosstab
  - Export to Excel
  - Display graph below data (a graph already exists but it-s display option is set to "hidden")
  - Pivot Data Points to the page axis
  - Change the page axis items
  - Switch to Big Table
  - Skip to page 2 of the results
  - Skip to the last page of the results (at 25 rows per page, 63 should be the last page)
  - Go to Small Table
  - Print the worksheet with both the table and the graph
  - Disconnect

The scenario uses a random think time of 45-75 seconds between each task. This scenario simulates a more sophisticated user who performs deeper analysis and uses more product features.

Appendix B – Oracle AS 10g (9.0.4) Configuration Details

The Application Server Discoverer machine was optimally configured to run Discoverer.

The Oracle HTTP server httpd.conf was configured as follows:

- KeepAlive Off
- MaxClients 400

J2EE/Oracle Containers for Java was configured as follows:

- OC4J_BI_Forms was the only OC4J process running
- OC4J_BI_Forms instance Server settings of one VM process per CPU (note that the increase in performance of having one VM process per CPU over one VM process per 2 CPUs is marginal so that if memory usage is a concern, 1 VM process per 2 CPUs might be better for real world applications)
- Java command line configured with 512 MB heap:

  -Xmx512M

- Java command line configured mod_oc4j timeout set to 1 minute:

  -Doracle.j2ee.http.socket.timeout=60000

WebCache process was disabled.

See the appropriate Oracle Application Server documentation for details on configuring your Oracle Application Server.

Discoverer Query Prediction was disabled. See chapter 20 "Discoverer registry settings" in the *Oracle Discoverer Administrator Administration Guide 10g (9.0.4)* for details on how to disable Query Prediction.

Appendix C – Database Configuration Details

The following parameters were set for the Database.

```
db_block_size 8192
open_cursors 400
db_block_buffers 24000
shared_pool_size 104857600
large_pool_size 614400
java_pool_size 20971520
log_checkpoint_interval 100000
log_checkpoint_timeout 18000
log_buffer 1638400
query_rewrite_integrity STALE_TOLERATED
_query_rewrite_vop_cleanup TRUE
job_queue_processes 2
job_queue_interval 60
timed_statistics true
```

```
sort_area_size 1310720
sort_multiblock_read_count 16
sort_area_retained_size 0
optimizer_mode all_rows
query_rewrite_enabled true
processes 400
```

Appendix D – References

Oracle Application Server Discoverer Configuration Guide 10g (9.0.4)

Oracle Discoverer Administrator Administration Guide 10g (9.0.4)

Oracle Application Server Discoverer Plus User's Guide 10g (9.0.4)

Oracle Database Data Warehousing Guide

Oracle Application Server Administrator's Guide 10g (9.0.4)

Oracle Application Server Performance Guide 10g (9.0.4)

Oracle Application Server Web Cache Administrator's Guide 10g (9.0.4)

Capacity Planning White Paper
January 2005
Author: Mike Donohue

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
www.oracle.com

Oracle Corporation provides the software that powers the internet.