

# ***IBM Insight for Oracle Database***

## ***Results of Analysis for XXX***

*Insight Data Collected On*

*February 9, 2001*

**(12:00am to 11:59pm)**

Document Prepared By IBM America's Techline  
On  
**February 11, 2001**

**<http://www.ibm.com/erp/oracle/insight>**

***IMPORTANT: Supplying correct hardware information*** during data collection process with IBM Insight for Oracle Database (Server Model Numbers, # of CPUs, Speed in MHz) ***is the sole responsibility of the customer.***

---

## **Control Section**

### **CONTACTING IBM**

You can contact IBM in one of the following ways:

- Send a note to [IBMERP@US.IBM.COM](mailto:IBMERP@US.IBM.COM) and include the word "INSIGHT" in your subject heading.
- Call us at 1-800-IBM-0222
- For a more detailed analysis than provided here, please contact Rudy Waldner of IBM Global Services at 919-301-4162. Rudy is responsible for providing service offering for client/server capacity analysis and modeling.



## Client Site Description

<b>Company Name</b>	RSD
<b>Industry</b>	IBM
<b>Address1</b>	sdfs
<b>Address2</b>	sdf
<b>City, State ZIP</b>	sdfsgs
<b>Country</b>	US
<b>Contact Name</b>	RSD
<b>Telephone</b>	123456
<b>email</b>	xyz@sfgasg
<b>IBM Sales Rep or Business Partner</b>	asfdgasg
<b>IBM Address1</b>	asdg
<b>IBM Address2</b>	ads
<b>IBM City, State ZIP</b>	dasfgfd
<b>IBM Country</b>	US
<b>IBM Telephone</b>	056256
<b>IBM email</b>	ter@sfdgds
<b>Reason for analysis</b>	test
<b>Special sizing considerations</b>	none
<b>Manufacturer of DB server</b>	IBM
<b>Model of DB server</b>	7009
<b>Number of CPUs</b>	1
<b>Processor Speed (MHz)</b>	222
<b>Physical Real Memory</b>	256
<b>Physical Disk (GB)</b>	20
<b>OS</b>	AIX
<b>Hostname</b>	rsdux
<b>Oracle SID</b>	rdm91
<b>Run Date</b>	Fri Feb 9 00:00:02 2001
<b>DB Size</b>	746586112
<b>DB Free Space</b>	320036864




## System Performance

**DISCLAIMER:** *Supplying correct hardware information during data collection process with IBM Insight for Oracle Database (Server Model Numbers, # of CPUs, Speed in MHz) is the sole responsibility of the customer.*

Server Name	Model	Number of CPUs	CPU Speed (MHz)	Memory	Physical Disk	Avg CPU Utilization	Peak Hour CPU Utilization
DB Server (rsdux)	RS/6000 7009	1	222	256 MB	20GB	24%	58%

## CPU Utilization




### ANALYSIS

This chart contains two main elements: The CPU utilization (system, user, idle, and wait caused by I/O wait): and the number of processes available in the run queue. The I/O wait depicted is only important if a run queue exists that is larger then the number of processors on the server being monitored. The CPU utilization elements are scaled using the y-axis on the left side while the run queue is indexed with the y-axis on the right side. Additionally, the CPU utilization is part of the graphic is what is called a stacked area chart, therefore total CPU utilization is the top of the system area which represents user plus system.

In general, the user consumption should be three to four times that of system. There are many various opinions on what the total CPU utilization should be. In general, the more processors installed, the more consistent the performance of Oracle will be during peaks. Basically, small spikes do not affect performance as much on CPU's with more processors.

If the I/O wait component comes down dramatically from the top of the graph, and the run queue is above the number of processors, it is important to look closely at the next two charts ("Demand Paging" and "Disk Activity")

## Demand System Paging




### ANALYSIS

This chart was designed to provide an understanding of the relationship between paging and Oracle memory utilization. It also consists of two graphical elements. The first is the stacked area chart of page ins and page outs plotted against the left y-axis, and the second being the total Oracle memory being used plotted against the y-axis on the right.

The two, in conjunction, should be informative. If there are spikes in paging without an associated spike in Oracle memory consumption, other applications running on the system should be investigated. If there is a common spike, an evaluation of the nature should be evaluated. For example, a scheduled batch Oracle job, starting everyday at 11:00, increases memory and results in more paging. Can it be moved to another time slot or should more memory be procured.

In general, the paging represented here should be avoided. This is real paging, causes system overhead, and potential delays in dispatching work. To say it another way, this does not represent the benign paging that takes place with today's advanced operating systems as they implement advanced algorithms. This chart represents the type of paging to be avoided.

## Disk Activity




### ANALYSIS

This compound chart shows the number of disk over 15% busy against the left axis and the average service time of all disks on the right axis. It has been seen, over time, that when the number of disk over 15% busy start to rise, it affects the total performance of the disk farm. The reason for this is that a higher percentage of I/O's are going to disks that are potentially over committed.

In the RDBMS world, it is not unusual for one or two disks to become busy for a limited period as reports and queries run. These tend to have a limited affect on the overall performance of the disk farm. However, if the database consistently has a set of very active disk, a potential performance problem is in the making. Normally, average service times on modern disk should be running under 10ms.

## CPU Consumption Types By TOD

**Break Out Of CPU Consumption Types By TOD**  
 (Grouped By Major Application, Always Adds To 100%)



### ANALYSIS

This chart is a simple stacked area chart. It represents the CPU seconds consumed by each of the various groups defined. The stack was designed to exemplify many problems. In this context, the “user” consumption being reported are users doing “user” things, and system is those system tasks required to run the system (not the system component of other non-system activities).


Hopefully, on a production database, there will be a preponderance of activity because of database activity. To determine the percent of the system taken up by these various elements, use the following formulae:  $\text{CPU seconds} / (300 * \text{number of engines})$ . However, that was not the purpose of the chart, it was to provide a quick visual understanding of where CPU resources were being spent.



---

## CPU Consumption Distribution By Type

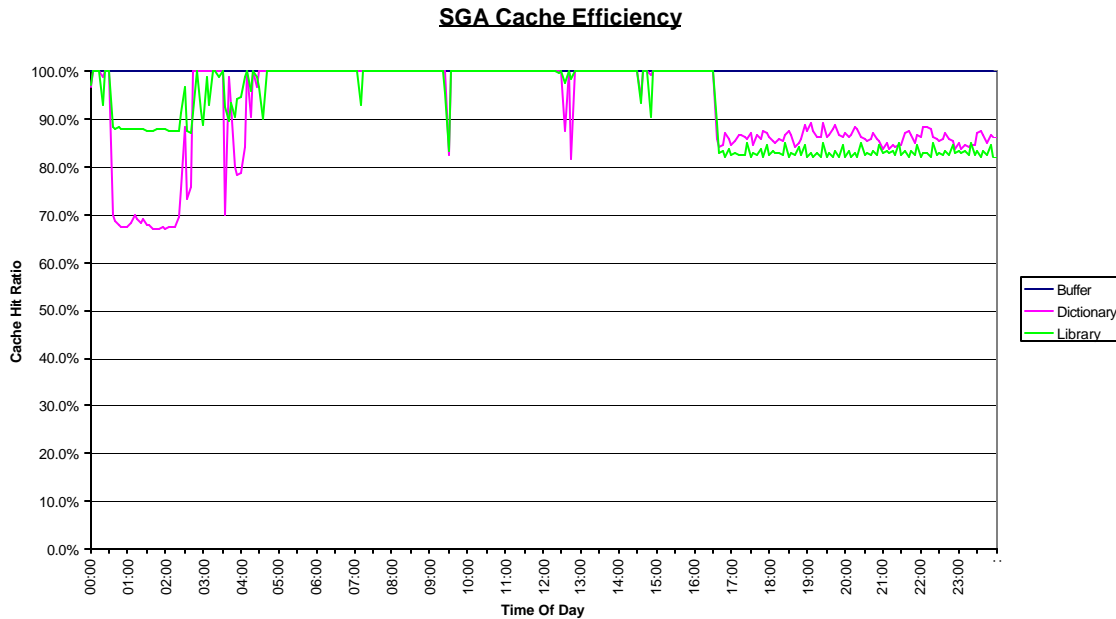
**CPU Consumption Distribution By Type**  
**(During Measurement Period)**



### ANALYSIS

This chart simply summarizes the details of the previous chart “CPU Consumption Types By TOD”. It represents the ratios that exist between the various seconds consumed by the various application segments during the collection period. Depending on the primary purpose and use of the CPU, it is up to management to decide if it is appropriate. However, it should be noted that if this is the primary serving database, and there is significant other activity, the database may provide erratic responsiveness to database applications.

## SGA Cache Efficiency




### ANALYSIS

This chart portrays the Cache hit ratios for various buffering elements within Oracle. In general, higher is better. A 100% value on the chart simply means that all requests for that type of data were retrieved from the buffers with no I/O required. Different types of applications have different acceptable values for these three main ratios. It should be noted that these ratios represent the value for each 5-minute interval, not the ratios normally referenced in books that represent the hit ratios from the time Oracle was started. But, as is true for all buffer hit ratios, the system should have been running long enough to populate the buffers to some level of steady state before measuring.

To use this chart, check with appropriate sources for acceptable and good values for these ratios, and compare. Another element would to watch for any significant changes during the collection period. This could imply some less then optimal code running that is lowering the rate, or incorrect setup.

## SGA Memory Allocations

**SGA Memory Allocations**



### ANALYSIS

This pie chart shows the major memory components within the Oracle System Global Area (SGA). The SGA primarily consists of the database buffer cache, the redo log buffers and one or more “pool” areas. All Oracle database instances have a “shared pool” which is used to cache important Oracle structures in memory. For example, the dictionary cache and the library cache are both contained in the shared pool. In later Oracle releases, it is also possible to have other pools within the Oracle SGA. For example, if a Java pool is specified in the Oracle ini.ora parameter file, memory for this pool will be allocated in the SGA.

In many Oracle installations, the database buffer cache (identified by the `db_block_buffers` component) is the largest component of the SGA. If this is not true on your system, it may be appropriate to investigate whether or not the database buffer cache is adequately sized for your environment. For each of the pool areas, Oracle reports statistics on the memory components within that pool as well as the available free memory in that pool. Typically in a well tuned environment, free memory will be around 10-20% of the pool size. If free memory is less than this, it may be an indication that the pool is undersized and Oracle is doing additional work (I/Os, object parses, etc.) that could be avoided if the pool were adequately sized. If there is a large amount of free memory available, this could indicate that the pool is fragmented (that might be avoided by pinning objects in the pool) or the pool is oversized. Cache hit ratios (shown on the Buffer Hit Ratios chart) for the major SGA components can be used to help determine if the SGA components are properly sized.



**END OF INSIGHT FOR ORACLE REPORT**