

Collecting IO requirements.

- Requirements are typically expressed in MB/s or IOPS.
- It is important to know both the read and write requirements as they can be very different.
- Notice there is always a unit of work and a unit of time.
- The unit of time is nearly always seconds.
- Oracle v\$ views provide data for both MB/s and IOPS, plus read and write details.
- The data source changes with Oracle 10g.
- OraPub has a number of IO requirement reports.

OraPub IO requirement reports.

```

SQL> @iosumx
Remember: This report must be run twice so both the initial and
final values are available. If no output, press ENTER a few times.

Database: prod3                               27-AUG-09 06:25pm
Report:   iosumx.sql                           OSM by OraPub, Inc.      Page 1
          Oracle IO Interval (v$sysstat) Summary

IOP/s and IOP
Total Read           :           197.043           45714
Total Write          :           469.539           108933
Total R+W            :           666.582           154647
MB/s and MB
Total Read           :             3.131             726.477
Total Write          :             7.029           1630.771
Total R+W            :            10.161           2357.248
Detailed Component Data
Interval (s)        :                232
SRVR Read IOP/s, IOP :           197.043           45714
SRVR Read MB/s, MB  :             3.131             726.477
DBWR+SRVR Write IOP/s, IOP :           460.866           106921
DBWR+SRVR Write MB/s, MB  :             3.741             867.945
LGWR Write IOP/s, IOP  :             8.672             2012
LGWR Write MB/s, MB   :             3.288             762.826
    
```

10g+ SQL for IO requirements.

```

Server process
  read IO operations, total
    v$sysstat.physical read IO [total] requests
  read MBs, total
    v$sysstat.physical reads X block size / (1024X1024)

Database writer and server process
  write IO operations, total
    v$sysstat.physical write IO [total] requests
  write MBs, total
    v$sysstat.physical writes X block size / (1024X1024)

Log writer process
  write IO operations, total
    v$sysstat.redo writes
  write MBs, total
    v$sysstat.redo size / (1024X1024)

```

IO capacity data sources.

- Only the IO administrator can tell you the IO capacity...if you're lucky.
- However,
 - We know Oracle's requirements, both in IOPS and MB/s, plus the read and write details.
 - One Oracle IO call can result in multiple IO actions by the IO subsystem.
 - The Oracle IO wait events provides the "truer" IO response time.
 - Only requests that result in a wait *occurrence* are recorded by the wait interface. When there are lots of direct and async reads and there is a clear IO issue, the ORTA response time will typically be much less than the average IO wait time or the IO subsystem's recorded IO response time.
 - If the ORTA IO RT or the wait interface event time is greater than your rule of thumb or SLA, then you know the requirements are exceeding the available capacity.
 - As presented later in this course, since we know the IO arrival rate, response time, and queue time, we can infer the IO subsystem's capacity! But be careful because once you hit the elbow in modern IO subsystems, it skyrockets... sometimes without warning!

Introducing OraPub's FF diagnostic template.

- The free template takes input statistics from standard performance data sources, but is formatted to naturally enable a completely quantified response time analysis.
- The numerics can then be used for 3-Circle ORTA firefighting work, RT graph creation, and also for anticipating a performance solution.
- There are four integrated worksheets in the workbook covering the Oracle RTA, the Oracle workload, the operating system, and application SQL.

“Ora RT” worksheet

Firefighting Diag												
	A	B	C	D	E	F	G	H	I	J	K	
1	Oracle Response Time Diagnosis											
2										Total (s)	Average (ms)	Avg Type
3	Response Time				347.636							
4	Service Time					343.916						
5	Srvr Proc								325.706			
6	BG Proc								18.210			
7												
8												
9												
10	Queue Time					3.720						
11	IO						3.480					
12	Read							2.420			3.010	wa
13											3.769	wa
14	1	na							2.100		3.875	straight
15	2	na							0.320		3.077	straight
16	3	na							0.000		0.000	straight
17	Write							1.060			1.277	wa
18	1	control file par wr							0.840		1.400	straight
19	2	log file par wr							0.220		0.809	straight
20	3	na							0.000		0.000	straight
21	Other						0.240				8.571	wa
22	1	latch free							0.240		8.571	straight
23	2	na							0.000		0.000	straight
24	3	na							0.000		0.000	straight

This should come as no surprise as it is standard Oracle RTA information.

Notice there are **two** average calculations.

The use of Excel makes it much easier to quickly quantify the situation, use more precise data and math, and reduce mistakes.

“Ora WL” worksheet

Metric	Workload (total)	ST (ms/wl)	QT (ms/wl)	RT (ms/wl)	Workload (wl/ms)	Workload (wl/sec)
Block Changes	1178	291.9493735	3.1578947	295.1072683	0.000654	0.654212
Phys Blk Gets	2 171958.1810000	1860.0000000	173818.1810000	0.000001	0.001111	
Log Blk Gets	36954234	0.0093065	0.0001007	0.0094072	20.522833	20522.832993
Executions	53123	6.4739635	0.0700262	6.5439896	0.029502	29.502288
User Calls	103618	3.3190793	0.0359011	3.3549804	0.057545	57.545095
Interval	1800.640 sec					

The sample interval can be entered with high precision.

Make sure the time unit is in seconds!

The workload metrics are common instance statistics (`v$sysstat`) and one will likely be the best RT curve arrival rate.

With the workload, CPU consumption, and wait time known, the Oracle arrival rate, service time, queue time, and RT can be calculated!

“Appl” Worksheet

SQL_ID	Elapsed Time (sec)	PIO (tot)	LIO (tot)	CPU (tot sec)	Exec (tot)	Rows (tot)	Sorts (tot)	SQL Type	Elapsed Time (ms/exec)	CPU Time (sec/exe)	PIO (pio/exe)	LIO (lio/exe)
7pkuzfhjsgp0f	51.842	0	8,409,060	51.842	395	396	0	SELEC	0.131	0.131	0.000	21288.759
40k7ynchcigt	25.201	0	4,077,120	25.201	191	192	0	SELEC	0.132	0.132	0.000	21346.178
fmy27wpg8x0ya	3.261	0	318,525	3.261	14	15	0	SELEC	0.233	0.233	0.000	22751.786
cwf9j5kmx4h7	1.518	0	148,645	1.518	6	7	0	SELEC	0.253	0.253	0.000	24774.167
bhd50g0yh9b2y	3.247	0	318,525	3.247	14	15	0	SELEC	0.232	0.232	0.000	22751.786
65a7ucrtg226p	3.890	0	382,230	3.890	17	18	0	SELEC	0.229	0.229	0.000	22484.118
g4y6nw3tts7cc	0.158	0	0	0.158	1021	1021	0	PL/SQ	0.000	0.000	0.000	0.000
4dxxk6cgg0jx	0.506	0	1,866	0.506	120	12	0	INSER	0.004	0.004	0.000	15.550

Data entry can be a hassle, so being able to copy/cut/paste using a text editor and Excel is very helpful.

Ensure the units of time and units of work are correct!

All data are deltas ($T_1 - T_0$) and can be gathered from `v$sql`, Statspack, or AWR.

The data in this worksheet is **not** linked to the other worksheets. This means you do NOT need to enter ALL the cells and can limit the SQL entered.

Operating Systems Diagnosis

“OS” Worksheet

The CPU cores and v \$osstat .busy details must be entered as they are tightly linked to other statistics.

This is what actually occurred during the sample interval.

Row	Column A	Column B	Column C	Column D	Column E	Column F
1						
2						
3	Capacity		21019.2	Core se		
4	CPU Cores		4			
5	Interval		5254.8	sec		
6	Requirements					
7	OS vsosstat.busy		21,013	sec		
8						
9	Utilization					
10	OS vsosstat		99.97%			
11	OS vmstat		100.00%			
12	Oracle		98.75%			
13						
14						
15	Response Time (from Oracle)					
16	Read and Write		0.00000	ms/lop		
17	Read		0.00000	ms/lop		
18	Write		0.00000	ms/lop		
19						
20	Workload					
21	IOPS Read + Write		0.2	IOPS		
22	IOPS Read		0.0	IOPS R (svr prc)	99	
23	IOPS Write		0.2	IOPS W (dbwr+lgwr)	918	75
24	MB/s Read + Write		0.002	MB/s		
25	MB/s Read		0.000	MB/s R (svr prc)	1	
26	MB/s Write		0.002	MB/s W (dbwr+lgwr)	7	2
27						
28						
29						
30	Network Subsystem					
31	tnsping					
32	client 1	na		avg ms		
33	client 2	na		avg ms		
34	client 3	na		avg ms		
35	SQL*Net total round trips		0.1427	occurs/s		
36	SQL*Net total bytes transferred		0.1105	KB/s		
37						
38	SQL*Net roundtrips to/from client		0.1427	trips/s	750	trips
39	SQL*Net roundtrips to/from dblink		0.0000	trips/s	0	trips
40	bytes received via SQL*Net from client		0.0541	KB/s	291,162	bytes
41	bytes received via SQL*Net from dblink		0.0000	KB/s	0	bytes
42	bytes sent via SQL*Net from client		0.0564	KB/s	303,349	bytes
43	bytes sent via SQL*Net from dblink		0.0000	KB/s	0	bytes
44						
45						
46						
47	Memory Subsystem					

pg:9 v2f ©2011 OraPub, Inc. OS Situational Analysis