A Complex I/O Workload Prediction Example

Let **A**, **B** and **C** represents workload directions.
For instance: **A** = 8 KB RR
B = 64 KB SW
C = 32 KB SW

Let the average load level be 5.6.
A sequence of such system states, where N is Load Level:

A	**A**	**A**	
B	**A**	**B**	
A	**C**	**C**	
C	**B**	**C**	
A	**A**	**B**	
B	**C**		
N=6	N=6	N=5	…

Counts: A = 7, B = 5, C = 5
Total: 17
Portions: Pa = .4117
Pb =.2941
Pc =.2941

For a sustained I/O workload state described by the above direction and magnitude the response time for each workload component when applied to a specific storage resource is the sum of the response times for each component present scaled by its portion by mean convergence of the stationary distribution.

Let the storage resource sustained response time at a load level 5.6 for each of A, B and C be respectively: Ra = .008
Rb = .012
Rc = .010

The resultant response time is a linear combination of the component response times. It is the sum of the products of the portions of response time present for each workload direction:

$$R = Pa * Ra + Pb * Rb + Pc * Rc$$
$$= .4117 * 0.008 + .2914 * .012 + .2941 * .010$$
$$= .0097$$

By Littles Law the throughput in operations per second is thus:
N/R = 5.6/.0097 = 577 IOPS @ 9.7 ms.

Example 2, using ORtera:

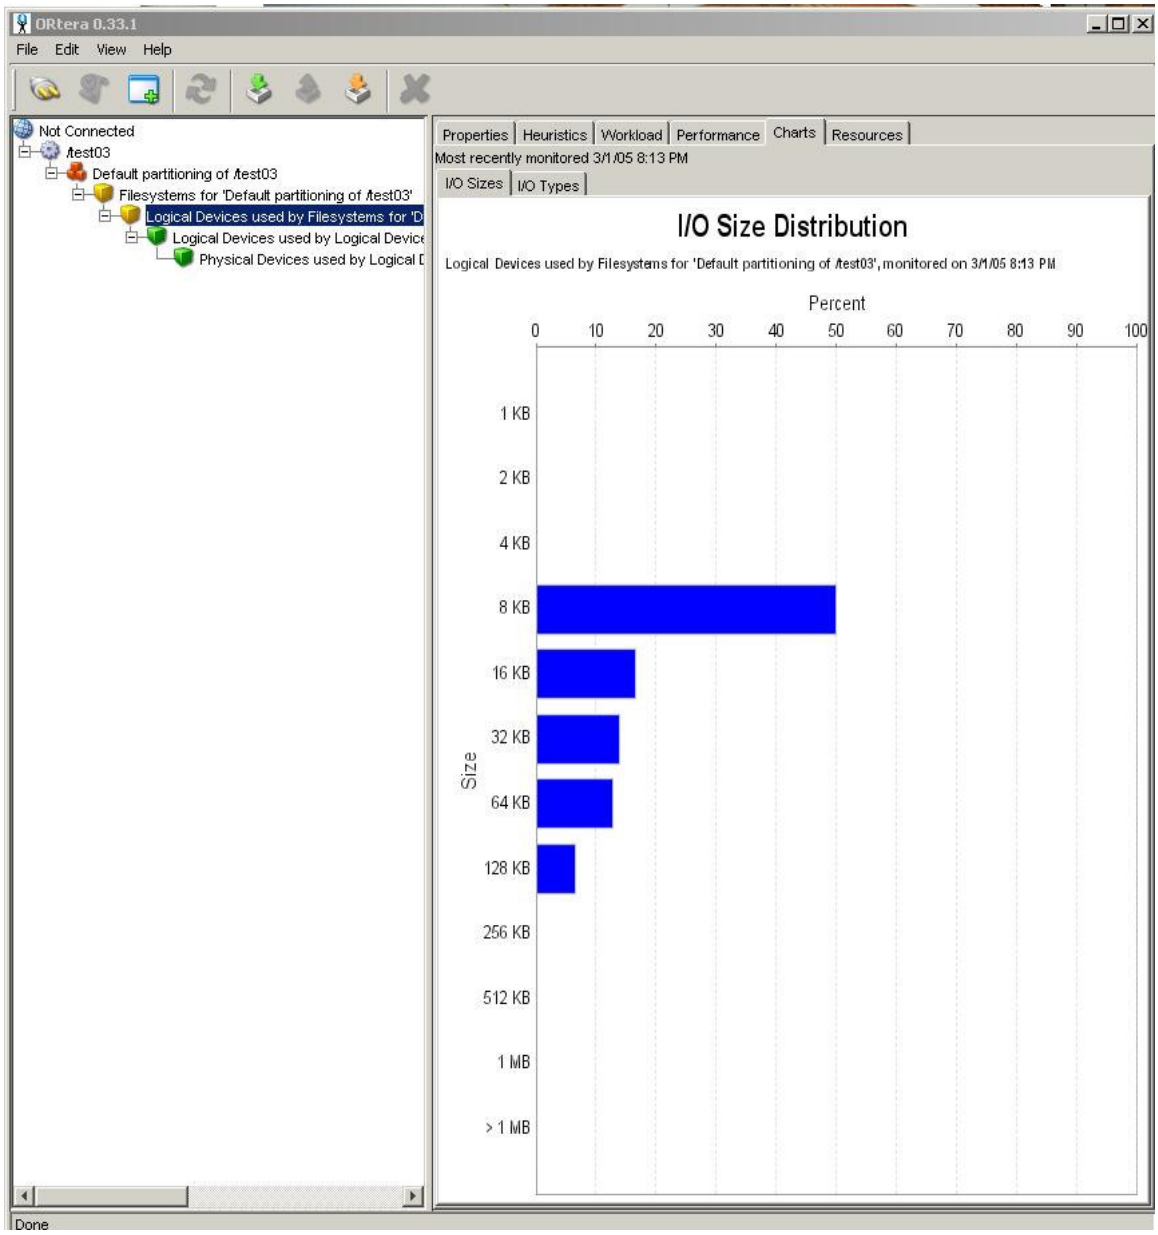

The above chart defines the I/O Size of the workload primary directions. Holding the cursor at the end of each bar gives the portion of each I/O size histogram.

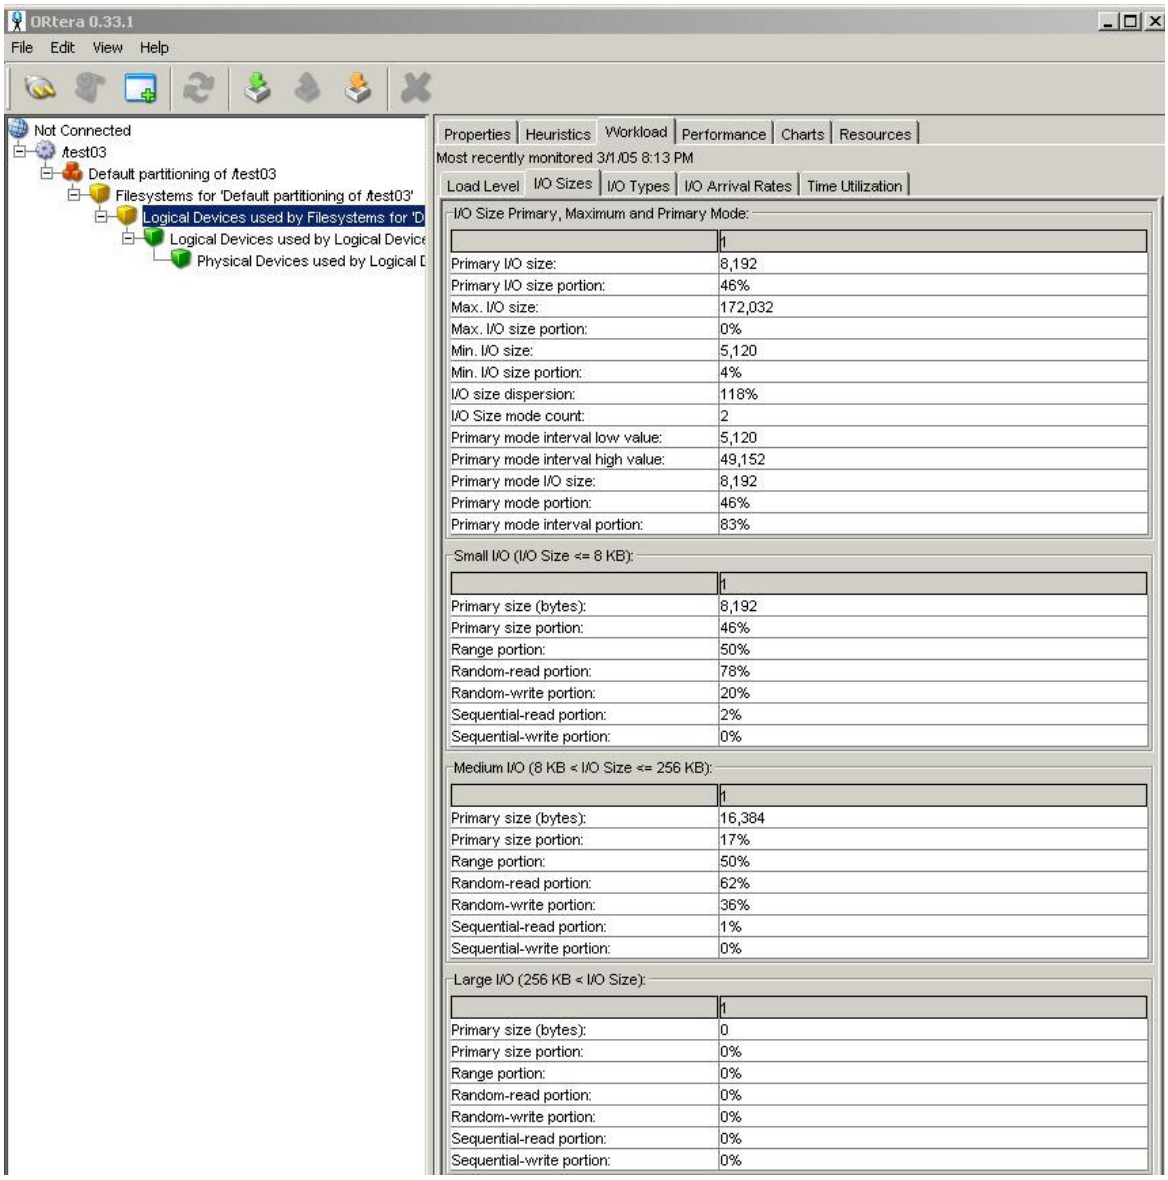

The above defines the Access Type per I/O Size categories, small, medium and large. Note we are approximating the primary workload directions based on empirical data describing the I/O Size and Access Types portions present in the workload.

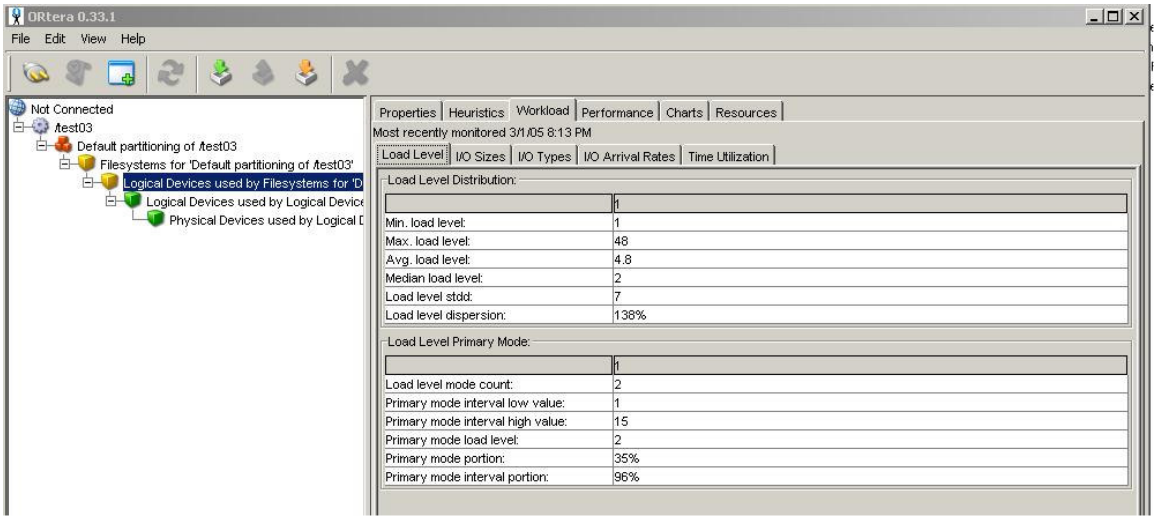

The above shows the load level (magnitude) of the workload.

Logical Device

Threads	IO Type	IO Size	Resp (ms.)	IOPS	IO Count	Total sec.
5	RR	8192	26.992371	182.7597	5506	30.127
5	RW	8192	42.200454	116.8896	3522	30.131
5	RR	16384	27.874296	176.5661	5330	30.187
5	RW	16384	44.628964	107.8065	3342	31
5	RR	32768	29.314555	168.2073	5077	30.183
5	RW	32768	48.622598	101.7089	3071	30.194
5	RR	65536	32.533391	151.5212	4567	30.141
5	RW	65536	56.928653	84.5484	2621	31
5	RR	131072	39.001048	126.6521	3814	30.114
5	RW	131072	72.760467	68.0561	2054	30.181

Portion	I/O Size	RR	RW	
0.5	8 KB	0.8	0.2	1
0.17	16 KB	0.64	0.36	1
0.14	32 KB			
0.13	64 KB			
0.06	128 KB			

8 KB	RR	0.400
8 KB	RW	0.100
16 KB	RR	0.109
16 KB	RW	0.061
32 KB	RR	0.090
32 KB	RW	0.050
64 KB	RR	0.083
64 KB	RW	0.047
128 KB	RR	0.038
128 KB	RW	0.022

		E[R]
8 KB	RR	0.01080
8 KB	RW	0.00422
16 KB	RR	0.00303
16 KB	RW	0.00273
32 KB	RR	0.00263
32 KB	RW	0.00245
64 KB	RR	0.00271
64 KB	RW	0.00266
128 KB	RR	0.00150
128 KB	RW	0.00157

Predict R	Accuracy %
0.0343	97

E[X]	Predict X	Accuracy %
137	132	95

The above shows the result of applying the workload magnitude and direction to the storage resource response time as measured with atlassp(1) the ORtera storage profiler. The profiler provides a response time for the given load level, I/O size, and access type. The I/O Size portions were obtained from the I/O Size chart. The RR, RW portions were obtained from the I/O Size table. The estimate is a linear combination as shown in example (1) above. The result in this case, is a response time prediction that is within 3% of the actual measured response time, and within 5% of the actual measured throughput.