

# ORtera Summary Report for 2005-07-20 14:49 Default partition for [1516] oracleDWHS ...


Monitored on 7/20/05 11:53 AM

## Filesystem /stage for '2005-07-20 14:49 Default partition for [1516] oracleDWHS (LOCAL=NO)'

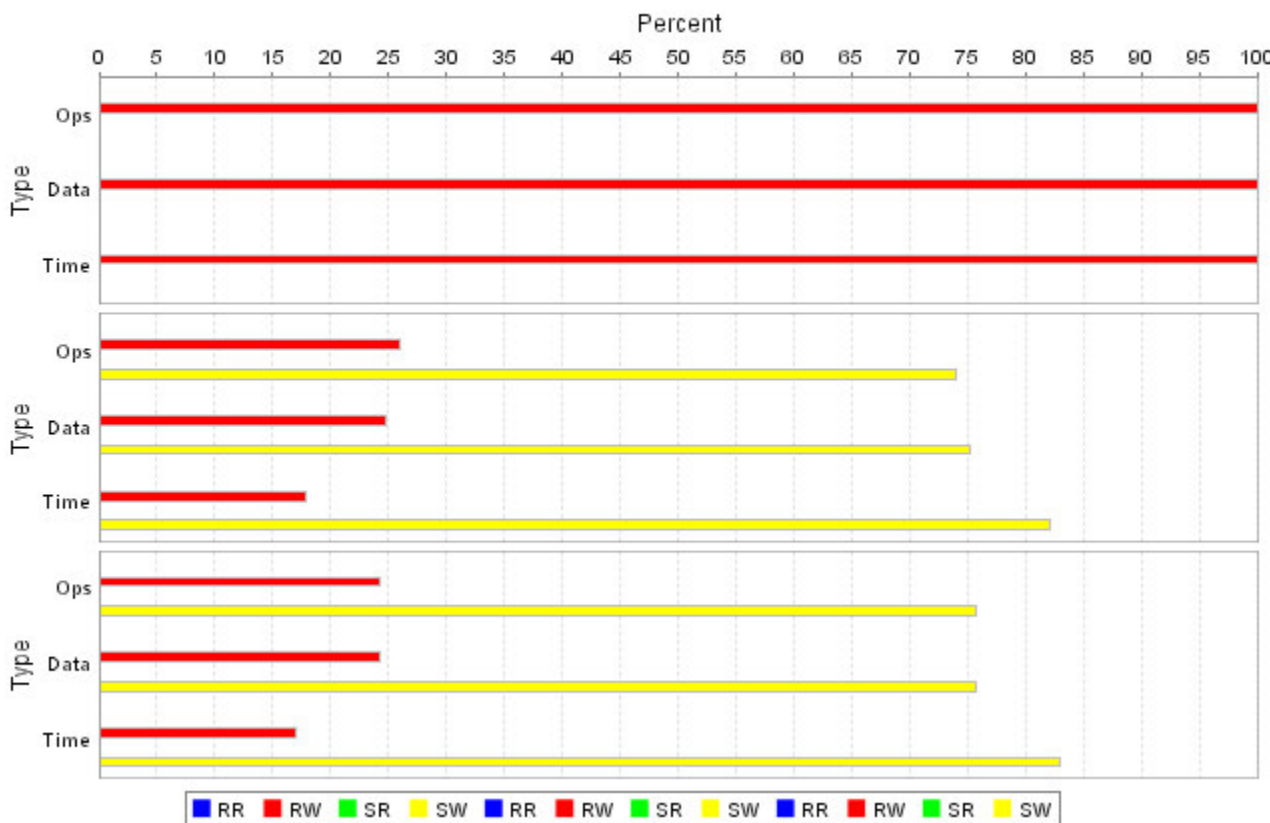
The charts show data for each level of this resource hierarchy in this order:

1. Filesystem /stage for '2005-07-20 14:49 Default partition for [1516] oracleDWHS (LOCAL=NO)'
2. Logical Devices used by Filesystem /stage for '2005-07-20 14:49 Default partition for [1516] orac...
3. Physical Devices used by Logical Devices used by Filesystem /stage for '2005-07-20 14:49 Default ...

### I/O Size Distribution



### I/O Type Composition



### Resource Hierarchy Heuristics

<b>Filesystem I/O Fragmentation</b>	<b>Severe issue</b>
<p><b>Condition:</b>                  The average I/O size at level 0 type Filesystem is 253.658 KB, and is larger than the average I/O size at level 1 type Logical Device, which is 62.711 KB. This indicates I/O fragmentation. The likely cause is suboptimal configuration of the filesystem read-ahead or write-behind buffer, or other filesystem-specific tunable for I/O size.</p> <p>At the filesystem level, 95% of the I/O is between 224 KB and 256 KB, with a primary I/O size of 256 KB and a standard deviation of 8.333 KB. At the supporting device level, 95% of the I/O is between 64 KB and 64 KB, with a primary I/O size of 64 KB and a standard deviation of 7.965 KB.</p>	
<p><b>Recommendation:</b>                  The filesystem I/O size should be larger. We suggest 512 KB. This suggestion is based on three times the standard deviation of the application I/O size plus the primary I/O size of 256 KB, rounded up to a power of two.</p>	

<b>Unused Physical Resources</b>	<b>Severe issue</b>
<p><b>Condition:</b>                  There is/are 14 unused device(s) out of 16 allocated. Only 12% of the devices allocated are used.</p>	
<p><b>Recommendation:</b>                  The maximum load at the physical level is 7 thread(s). More than 7 device(s) cannot be leveraged from a performance perspective at this load level. This is a good candidate to handle more work. Alternately, some resources could be freed up to be used elsewhere.</p>	

<b>CPU Deficit</b>	<b>Severe issue</b>
<p><b>Condition:</b>                  The top-level filesystem or logical device has a capability busy utilization level of 100.7%, while the bottom-level supporting devices have a capability busy utilization level of 0.7%. This indicates that the upper layer is running out of CPU to drive the bottom layer. Alternatively, there can be other indirect causes, such as insufficient number of threads, lock contention, or scheduling delays.</p> <p>The portion of CPU granted is only -15,178.1% of that desired. Negative values indicate an overloaded condition. Hence, some amount of the time attributed to device response time is process-scheduling delay, lock wait time, or caused by too few threads to be able to take advantage of the available CPUs, all of which cause behavior consistent with being out of CPU power.</p>	
<p><b>Recommendation:</b>                  More or faster CPUs are needed to drive the I/O sub-system to its full capability. The use of shorter time quanta in the system scheduling classes may improve I/O latency under these conditions. If this heuristic is reported for a filesystem, the kernel tunable autoup may need to be adjusted to give fsflush more time to manage memory. If your filesystem or application can be configured to use more threads, that may also help. In addition, a larger page cache may be able to offset saturation.</p>	

**Filesystem /stage for '2005-07-20 14:49 Default partition for [1516] oracleDWHS (LOCAL=NO)' Layer**

### Resource Layer Heuristics

<b>Unacceptable Response Times</b>	<b>Severe issue</b>
<p><b>Condition:</b>                  The average response time of 1,326.8 ms., with a standard deviation of 855.4, is unacceptable for most workloads. An example of files or devices with excessive response time:                  /stage/.ino_13:264_7000                  Response Time Avg. ms. 1,326.82                  Stdd. ms. 855.43                  RR Avg. (ms.) 0 RR Max. (ms.) 0 RR Stdd. (ms.) 0                  RW Avg. (ms.) 1,326.82 RW Max. (ms.) 2,093.55 RW Stdd. (ms.) 855.43                  SR Avg. (ms.) 0 SR Max. (ms.) 0 SR Stdd. (ms.) 0                  SW Avg. (ms.) 0 SW Max. (ms.) 0 SW Stdd. (ms.) 0</p>	
<p><b>Recommendation:</b>                  To maintain response time averages in the 10.0 ms. range, more physical resource is required here. Alternatively, some of the workload may be relocated to reduce the load level. Also, see other heuristics for load imbalance and/or I/O size consistency issues; correcting those issues may improve the current response time condition.</p>	

### Workload and Response Data Summary

Capability Busy:	101%			
Load Level:	1 (min)	2 (max)	1.3 (avg)	0.5 (stdd)
I/O Size (KB):	224 (min)	256 (max)	253.659 (avg)	8.333 (stdd)
I/O Per Second:	1 (min)	14 (max)	2.7 (avg)	3 (stdd)
Bandwidth (MB/sec):	0.6 (min)	0.68 (max)	0.34 (avg)	0.76 (stdd)
Response Time (ms):	2.48 (min)	2,093.55 (max)	1,326.82 (avg)	855.43 (stdd)
I/O Ops:	0% (RR)	0% (SR)	100% (RW)	0% (SW)
I/O Data:	0% (RR)	0% (SR)	100% (RW)	0% (SW)
I/O Time:	0% (RR)	0% (SR)	100% (RW)	0% (SW)

**Logical Devices used by Filesystem /stage for '2005-07-20 14:49 Default partition for [1516] orac... Layer**

### Resource Layer Heuristics

Consistent Resource Level	No issues
<p><b>Condition:</b>                  The collection of filesystems or devices is reasonably consistent with respect to I/O size, load distribution and load levels, and response times are at generally acceptable levels.</p>	
<p><b>Recommendation:</b>                  Perform ORtera analysis at regular intervals to detect any change of condition.</p>	

### Workload and Response Data Summary

Capability Busy:	1%			
Load Level:	1 (min)	7 (max)	2.7 (avg)	1.3 (stddev)
I/O Size (KB):	0.5 (min)	64 (max)	62.711 (avg)	7.966 (stddev)
I/O Per Second:	1 (min)	55 (max)	9.1 (avg)	11.1 (stddev)
Bandwidth (MB/sec):	0 (min)	0.57 (max)	0.32 (avg)	0.72 (stddev)
Response Time (ms):	0.62 (min)	12.8 (max)	3.45 (avg)	1.96 (stddev)
I/O Ops:	0% (RR)	0% (SR)	26% (RW)	74% (SW)
I/O Data:	0% (RR)	0% (SR)	25% (RW)	75% (SW)
I/O Time:	0% (RR)	0% (SR)	18% (RW)	82% (SW)

**Physical Devices used by Logical Devices used by Filesystem /stage for '2005-07-20 14:49 Default ... Layer**

## Resource Layer Heuristics

Unbalanced Device Data	Minor issue
<p><b>Condition:</b> A group of devices appear to be over-utilized relative to other devices in the same set. The distribution of 10,849 KB over 2 device(s) is not balanced.</p> <p>100.0% of the data is to/from only 50.0% of the devices. The first of these devices is /dev/rdisk/c5t14d1s2 with 10,848 KB, the last device is /dev/rdisk/c5t14d1s2 with 10,848 KB. See the Workload -&gt; I/O Per Resource tab (sorted on Data) for the complete list of devices in this group.</p>	
<p><b>Recommendation:</b> The 1 device(s) in this group is/are in a different class with respect to other devices in the set. This implies an imbalance of data per device. If these devices are supporting a logical device, the RAID stripe width may be too large, or a kernel I/O size limit (such as maxphys in Solaris) may be forcing I/O to be split. There could also be a hot address range, such as metadata for a filesystem or database table, causing uneven load. Check for and correct these conditions if they exist.</p>	
Unbalanced Device Response Times	Minor issue
<p><b>Condition:</b> A group of devices appear to be underutilized relative to other devices in the set. The distribution of 0.6 seconds total response time over 2 device(s) is not balanced.</p> <p>Only 0.1% of the response time is spent accessing 50.0% of the devices. The first of these devices is /dev/rdisk/c5t15d0s2 with 0.001 seconds, the last device is /dev/rdisk/c5t15d0s2 with 0.001 seconds. See the Workload -&gt; I/O Per Resource tab (sorted on Time) for the complete list of devices in this group.</p>	
<p><b>Recommendation:</b> The 1 device(s) in this group is/are in a different class with respect to other devices in the set. This implies an imbalance in response-time per device. These devices are good candidates to handle more work.</p>	

Unbalanced Device I/O Operations	Minor issue
<p><b>Condition:</b> A group of devices appear to be underutilized relative to other devices in the same set. The distribution of 173 I/O operations over 2 device(s) is not balanced.</p> <p>Only 1.2% of all operations are to 50.0% of the devices. The first of these devices is /dev/rdisk/c5t15d0s2 with 2 OPS, the last device is /dev/rdisk/c5t15d0s2 with 2 OPS. See the Workload -&gt; I/O Per Resource tab (sorted on Ops) for the complete list of devices in this group.</p>	
<p><b>Recommendation:</b> The 1 device(s) in this group is/are in a different class with respect to other devices in the set. This implies an imbalance of operations per device. These devices are good candidates to handle more work.</p>	
Unbalanced Device Response Times	Minor issue
<p><b>Condition:</b> A group of devices appear to be over-utilized relative to other devices in the same set. The distribution of 0.6 seconds total response time over 2 device(s) is not balanced.</p> <p>99.9% of the response time is spent accessing only 50.0% of the devices. The first of these devices is /dev/rdisk/c5t14d1s2 with 0.579 seconds, the last device is /dev/rdisk/c5t14d1s2 with 0.579 seconds. See the Workload -&gt; I/O Per Resource tab (sorted on Time) for the complete list of devices in this group.</p>	
<p><b>Recommendation:</b> The 1 device(s) in this group is/are in a different class with respect to other devices in the set. This implies an imbalance in response time per device. If these devices are supporting a logical device, the RAID stripe width may be too large, or a kernel I/O size limit (such as maxphys in Solaris) may be forcing I/O to be split. There could also be a hot address range, such as metadata for a filesystem or database table, causing uneven load. Check for and correct these conditions if they exist.</p>	

Unbalanced Device I/O Operations	Minor issue
<p><b>Condition:</b> A group of devices appear to be over-utilized relative to other devices in this set. The distribution of 173 I/O operations over 2 device(s) is not balanced.</p> <p>98.8% of all operations are to only 50.0% of the devices. The first of these devices is /dev/rdisk/c5t14d1s2 with 171 OPS, the last device is /dev/rdisk/c5t14d1s2 with 171 OPS. See the Workload -&gt; I/O Per Resource tab (sorted on Ops) for the complete list of devices in this group.</p>	
<p><b>Recommendation:</b> The 1 device(s) in this group is/are in a different class with respect to other devices in this set. This implies an imbalance of operations per device. If these devices are supporting a logical device, the RAID stripe width may be too large, or a kernel I/O size limit (such as maxphys in Solaris) may be forcing I/O to be split. There could also be a hot address range, such as metadata for a filesystem or database table, causing uneven load. Check for and correct these conditions if they exist.</p>	
Unbalanced Device Data	Minor issue
<p><b>Condition:</b> A group of devices appear to be underutilized relative to other devices in the same set. The distribution of 10,849 KB over 2 device(s) is not balanced.</p> <p>Only 0.0% of the data is to/from 50.0% of the devices. The first of these devices is /dev/rdisk/c5t15d0s2 with 1 KB, the last device is /dev/rdisk/c5t15d0s2 with 1 KB. See the Workload -&gt; I/O Per Resource tab (sorted on Data) for the complete list of devices in this group.</p>	
<p><b>Recommendation:</b> The 1 device(s) in this group is/are in a different class with respect to other devices in the set. This implies an imbalance of data per device. These devices are good candidates to handle more work.</p>	

## Workload and Response Data Summary

Capability Busy:	1%			
Load Level:	1 (min)	7 (max)	2.7 (avg)	1.3 (stdd)
I/O Size (KB):	0.5 (min)	64 (max)	62.711 (avg)	7.966 (stdd)
I/O Per Second:	1 (min)	55 (max)	9.1 (avg)	11.1 (stdd)
Bandwidth (MB/sec):	0 (min)	0.57 (max)	0.32 (avg)	0.72 (stdd)
Response Time (ms):	0.8 (min)	12.58 (max)	3.35 (avg)	1.95 (stdd)
I/O Ops:	0% (RR)	0% (SR)	24% (RW)	76% (SW)
I/O Data:	0% (RR)	0% (SR)	24% (RW)	76% (SW)
I/O Time:	0% (RR)	0% (SR)	17% (RW)	83% (SW)