

## Re: I/O performance troubleshooting

**Source:** <http://unix.derkeiler.com/Mailing-Lists/AIX-L/2005-04/0128.html>

---

*Holger.VanKoll\_at\_SWISSCOM.COM*

**Date:** 04/11/05

Date: Mon, 11 Apr 2005 23:30:28 +0200  
To: aix-l@Princeton.EDU

what I can see:

- your fr:sr ratio is very low, sometimes even 1. indicates poor memory-usage
- oracle tries constantly to read and write

the first thing I'd check is if SGA isnt too low. if its a dedicated db-server, you can start with 1/3 up to 1/2 (or more) of physical memory

then, see what changes

from what you posted I doubt that another adapter will help. your disks are constantly overloaded with random io. maybe oracle hasnt enough sga to optimize what it sends to disk.

same goes for vmtune - you are not paging to paging-space -> filecache reduction probably wont help

but your poor memory-usage (aix only needs to look at one page to get one page that can be freed...) could be improved by increasing read-ahead.

first, however, I'd experiment with sga

---

From: IBM AIX Discussion List [mailto:aix-l@Princeton.EDU] On Behalf Of Mark Schlechte  
Sent: Monday, April 11, 2005 10:50 PM  
To: aix-l@Princeton.EDU  
Subject: I/O performance troubleshooting

I know this is a common problem and I've even posted an earlier message on this.  
Just wanted a little more feedback.

It seems pretty obvious to me my performance problem is with the disks I have. I had to turn what was our backup server into an oracle database server and the disks just seem to get hammered.

Filemon output:

Most Active Physical Volumes

---

util	#rblk	#wblk	KB/s	volume	description
------	-------	-------	------	--------	-------------

---

1.00	12248	2524	697.8	/dev/hdisk2	N/A
------	-------	------	-------	-------------	-----

## AIX-L: Re: I/O performance troubleshooting

```
1.00 6064 5148 529.6 /dev/hdisk4 N/A
1.00 11808 7024 889.6 /dev/hdisk0 N/A
1.00 15720 3052 886.8 /dev/hdisk1 N/A
0.99 6088 2544 407.8 /dev/hdisk5 N/A
0.00 56 0 2.6 /dev/hdisk3 N/A
```

This is a 7029-6C3 using the onboard ultra320 scsi controller to access 5 x 36GB disks (hdisk3 is rootvg). The oracle database is only about 10GB in size.

This is just a JBOD config (I know, I know) and we have spread the logical volumes holding the oracle datafiles across all of the disks to try and spread the load.

```
xx:ROOT> lvmstat -v datavg
```

```
Logical Volume iocnt Kb_read Kb_wrtn Kbps
hdu02 302586 5256881 510548 115.43
hdu01 249668 1796601 662886 49.22
hdu00 164084 252128 1439293 33.85
hdu04 150442 3776832 159044 78.77
hdu06 104773 3666164 434822 82.07
hdu08 88403 3509632 147029 73.18
hdu07 77660 367444 487839 17.12
hdu03 36489 90088 249764 6.80
hdu05 28255 110024 144824 5.10
loglv00 11887 0 47548 0.95
```

```
xx:ROOT > lslv -l hdu02
```

```
hdu02:/u02
```

### PV COPIES IN BAND DISTRIBUTION

```
hdisk1 012:000:000 100% 000:012:000:000:000
hdisk2 012:000:000 100% 000:000:012:000:000
hdisk4 012:000:000 100% 000:012:000:000:000
hdisk5 012:000:000 100% 000:012:000:000:000
hdisk0 011:000:000 100% 011:000:000:000:000
```

Just wondering if the group consensus would be to invest some money in a RAID adapter (I think the 5703 is the adapter I am looking for) as I'm thinking that might be my best bet but I don't know for sure.

All my other servers use SSA and whether it's the non-arbitrated loop technology that makes it work or the on-board FW cache I'm not sure, but I've never seen the disk get smacked like this before.

Should I look at vmtune settings etc.? Can that make much of a difference? I also have about 300 aioservers while on my prod server it is still using the default settings and it is only using 10 (go figure).

Here's some vmstat output also if that helps. Testing starts at 9:06.

```
kthr memory page faults cpu time
```

```
-----
r b avm fre re pi po fr sr cy in sy cs us sy id wa hr mi se
1 0 567068 7880 0 0 0 0 0 440 4102 531 6 1 85 8 09:05:33
0 0 567061 7885 0 0 0 0 0 559 15328 731 3 1 80 16 09:05:38
0 1 567063 7864 0 0 0 0 0 562 15756 747 2 1 89 8 09:05:43
0 0 567087 7839 0 0 0 0 0 565 2539 723 2 1 82 16 09:05:48
0 0 567087 7839 0 0 0 0 0 544 2406 683 1 1 95 3 09:05:53
```

AIX-L: Re: I/O performance troubleshooting

0 0 567105 7821 0 0 0 0 0 532 3713 679 4 0 92 3 09:05:58  
0 0 567111 7815 0 0 0 0 0 537 2343 712 4 1 82 13 09:06:03  
1 0 567393 7473 0 0 0 0 0 615 2343 872 2 1 78 18 09:06:08  
0 1 567399 7466 0 0 0 0 0 550 2591 718 3 1 92 4 09:06:13  
0 1 567414 7450 0 0 0 0 0 609 1962 797 3 2 73 22 09:06:18  
0 0 567414 7450 0 0 0 0 0 539 2765 692 1 1 94 4 09:06:23  
0 0 567135 7731 0 0 0 0 0 555 3016 720 2 1 84 12 09:06:28  
0 0 567191 7672 0 0 0 0 0 567 2789 740 3 1 88 8 09:06:33  
2 1 572755 277 0 0 3 5 7 0 767 5738 1153 7 4 61 28 09:06:38  
1 0 574859 122 0 0 46 399 5692 0 920 7274 1177 9 4 73 14 09:06:43  
1 0 575087 128 0 0 0 277 1442 0 905 4380 1278 48 2 27 22 09:06:48  
2 1 575111 172 0 0 0 373 381 0 1288 3607 1697 22 3 39 37 09:06:53  
1 0 574939 412 0 0 0 29 30 0 995 5027 977 9 2 79 10 09:06:58  
2 0 575341 265 0 0 0 51 52 0 598 3257 690 4 2 89 5 09:07:03  
1 0 575336 220 0 0 0 47 50 0 728 12856 862 17 1 69 12 09:07:08  
kthr memory page faults cpu time

---

r b avm fre re pi po fr sr cy in sy cs us sy id wa hr mi se  
1 1 575903 128 0 1 0 220 236 0 663 18783 872 8 2 54 36 09:07:13  
1 1 576150 123 0 0 0 323 328 0 818 5047 1307 5 3 17 75 09:07:18  
1 1 576249 1166 0 0 0 324 325 0 818 3975 1498 6 3 25 67 09:07:23  
1 1 576250 840 0 0 0 0 0 882 4673 1479 15 2 26 57 09:07:28  
1 1 576337 651 0 0 0 0 0 769 43422 1430 13 3 29 55 09:07:33  
1 1 576382 419 0 0 0 0 0 835 4252 1530 13 2 30 56 09:07:38  
1 1 581778 188 0 0 0 1049 1073 0 899 5472 1671 24 6 25 45 09:07:43  
1 1 581035 778 0 0 0 10 11 0 951 3484 1416 15 2 38 45 09:07:48  
2 1 580579 885 0 0 0 0 0 1052 4126 1577 34 2 21 43 09:07:53  
1 1 581258 165 0 0 0 17 17 0 930 3715 1892 46 2 9 43 09:07:58  
1 1 581442 124 0 0 0 52 52 0 1014 3215 1374 14 2 34 51 09:08:03  
1 0 581412 128 0 0 0 21 21 0 1075 4704 1345 18 2 39 41 09:08:08  
1 1 581399 120 0 0 0 17 17 0 885 3523 1306 17 2 44 37 09:08:13  
1 0 581129 130 0 0 0 25 25 0 778 2569 1208 13 1 36 50 09:08:18  
0 1 581134 125 0 0 0 29 29 0 729 2692 1137 11 2 54 33 09:08:23  
2 1 580906 365 0 0 0 42 42 0 820 3865 1331 17 1 40 41 09:08:28  
1 1 580646 638 0 0 0 41 41 0 959 3986 1465 17 2 33 48 09:08:33  
1 1 580902 125 0 0 0 88 90 0 1013 4332 1797 14 3 18 66 09:08:38  
3 2 580908 131 0 0 0 111 113 0 1822 10575 3454 13 4 19 64 09:08:43  
1 1 580925 124 0 0 0 107 107 0 910 3172 1712 19 2 9 69 09:08:48  
kthr memory page faults cpu time

---

r b avm fre re pi po fr sr cy in sy cs us sy id wa hr mi se  
1 1 580928 120 0 0 0 139 142 0 980 5233 1768 24 2 14 61 09:08:53  
1 1 580960 124 0 0 0 205 206 0 1080 4114 1942 19 2 10 69 09:08:58  
1 1 580968 124 0 0 0 125 126 0 954 3076 1781 14 2 11 73 09:09:03  
1 1 580785 123 0 0 0 64 66 0 1080 4921 1946 23 3 13 61 09:09:08  
1 1 580783 140 0 0 0 46 47 0 968 4737 1879 23 2 17 59 09:09:13  
1 1 580820 193 0 0 0 71 72 0 1064 4776 2009 27 2 10 61 09:09:18  
2 1 580819 133 0 0 0 45 51 0 985 3727 1773 23 2 20 55 09:09:23  
1 1 580833 133 0 0 0 34 37 0 864 3868 1520 19 2 34 45 09:09:28  
1 1 580833 123 0 0 0 38 42 0 929 3934 1766 20 2 18 60 09:09:33  
1 1 580915 117 0 0 0 118 130 0 994 4521 1849 18 3 14 65 09:09:38

AIX-L: Re: I/O performance troubleshooting

1 1 582014 196 0 0 0 259 264 0 940 5895 1769 18 3 31 48 09:09:43  
1 1 582046 128 0 0 0 23 24 0 929 5225 1830 16 2 26 56 09:09:48  
2 1 584348 2150 0 0 0 965 1054 0 1040 4497 1928 18 5 12 64 09:09:53  
1 1 586386 168 0 0 0 74 74 0 1023 4148 1740 20 3 23 54 09:09:58  
1 0 590267 1230 0 0 0 1002 1027 0 732 8614 1297 20 5 41 33 09:10:03  
1 1 591018 293 0 0 0 36 36 0 1171 4345 1885 23 2 19 55 09:10:08  
3 1 593669 1082 0 0 0 728 799 0 986 5118 1618 23 5 36 36 09:10:13  
2 1 599543 140 0 0 0 1018 1106 0 1140 5670 1986 62 6 2 30 09:10:18  
1 0 599403 453 0 0 0 49 49 0 1143 4542 2191 30 4 8 58 09:10:23  
2 1 598815 917 0 0 0 0 0 1052 17280 1618 46 3 14 36 09:10:28  
kthr memory page faults cpu time

---

r b avm fre re pi po fr sr cy in sy cs us sy id wa hr mi se  
2 1 598977 657 0 0 0 0 0 1158 4022 1736 35 3 26 36 09:10:33  
3 0 598950 541 0 0 0 0 0 1158 17688 1755 50 3 10 37 09:10:38  
1 1 599185 212 0 0 0 0 0 1130 17483 1802 47 3 8 42 09:10:43  
2 1 600324 255 0 0 0 256 256 0 1149 4873 1674 47 2 14 37 09:10:48  
1 1 604557 245 0 0 0 871 1303 0 1253 7394 2053 32 5 10 52 09:10:53  
1 0 605361 573 0 0 0 237 393 0 973 3893 1417 22 3 42 33 09:10:58  
1 1 605204 698 0 0 0 0 1013 3405 1206 23 2 48 27 09:11:03  
2 1 604745 1092 0 0 0 0 1011 3470 1328 19 2 44 35 09:11:08  
1 1 604953 703 0 0 0 0 1130 4272 1628 27 3 26 44 09:11:13  
2 0 605161 384 0 0 0 0 980 4294 1344 34 3 27 36 09:11:18  
2 1 604781 679 0 0 0 0 985 3558 1981 39 2 11 48 09:11:23  
1 1 605984 539 0 0 0 219 271 0 1006 5064 1705 42 4 13 42 09:11:28  
1 0 606073 409 0 0 0 0 971 3534 1565 22 2 33 44 09:11:33  
1 1 608239 198 0 0 0 401 768 0 973 4229 1466 26 3 33 38 09:11:38  
1 0 608186 200 0 0 0 0 920 3759 1168 21 2 51 26 09:11:43  
1 1 607810 507 0 0 0 0 977 3510 1439 26 2 30 43 09:11:48  
1 1 608073 120 0 0 0 14 52 0 995 4238 1514 20 2 32 46 09:11:53  
1 0 608118 128 0 0 0 19 62 0 785 3933 1123 20 2 55 23 09:11:58  
1 1 608130 155 0 0 0 12 34 0 817 3642 1253 21 1 49 29 09:12:03  
2 0 608134 128 0 0 0 12 58 0 809 2901 1107 21 2 37 41 09:12:08

Hope the information I've provided is usefull. Sorry for the long post.

Mark

DISCLAIMER: The information transmitted is intended only for the addressee and may contain confidential, proprietary and/or privileged material. Any unauthorized review, distribution or other use of or the taking of any action in reliance upon this information is prohibited. If you received this in error, please contact the sender and delete or destroy this message and any copies.