| **Readers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 8Way LVM 30 Q Depth | 45944 | 76798 | 97727 |
| CFQ 12Way LVM 30 Q Depth | 46228 | 76049 | 97845 |
| CFQ 8Way Meta 30 Q Depth | 39831 | 68172 | 96147 |
| CFQ 8Way LVM 256 Q Depth | 46682 | 76272 | 97658 |
| CFQ 12Way LVM 256 Q Depth | 46468 | 75856 | 97818 |
| CFQ 12Way Filer 256 Q Depth | 67804 | 95504 | 107153 |
| CFQ 12Way Filer 256T Q Depth | 74118 | 96838 | 102163 |
| CFQ 12Way Filer 30 Q Depth | 73993 | 97338 | 105422 |
| CFQ 12Way NAS 920 | 90111 | 95791 | 104344 |
| CFQ 12Way NAS 960 | 83969 | 95789 | 108245 |

| **Random Readers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 8Way LVM 30 Q Depth | 46576 | 77091 | 98127 |
| CFQ 12Way LVM 30 Q Depth | 46895 | 76938 | 98187 |
| CFQ 8Way Meta 30 Q Depth | 40428 | 68958 | 96365 |
| CFQ 8Way LVM 256 Q Depth | 46503 | 76814 | 98125 |
| CFQ 12Way LVM 256 Q Depth | 47027 | 77108 | 98231 |
| CFQ 12Way Filer 256 Q Depth | 3466 | 6894 | 13704 |
| CFQ 12Way Filer 256T Q Depth | 3802 | 7120 | 13115 |
| CFQ 12Way Filer 30 Q Depth | 3644 | 6953 | 13948 |
| CFQ 12Way NAS 920 | 3562 | 6874 | 10664 |
| CFQ 12Way NAS 960 | 5542 | 14149 | 33564 |

| **Writers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 8Way LVM 30 Q Depth | 26793 | 45288 | 70207 |
| CFQ 12Way LVM 30 Q Depth | 27478 | 46829 | 71461 |
| CFQ 8Way Meta 30 Q Depth | 23470 | 36635 | 59758 |
| CFQ 8Way LVM 256 Q Depth | 26811 | 45620 | 67512 |
| CFQ 12Way LVM 256 Q Depth | 27465 | 46655 | 71622 |
| CFQ 12Way Filer 256 Q Depth | 34540 | 48864 | 62560 |
| CFQ 12Way Filer 256T Q Depth | 34696 | 49236 | 60110 |
| CFQ 12Way Filer 30 Q Depth | 42450 | 52634 | 67859 |
| CFQ 12Way NAS 920 | 37313 | 40530 | 42824 |
| CFQ 12Way NAS 960 | 51026 | 45361 | 48543 |

| **Random Writers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 8Way LVM 30 Q Depth | 18825 | 35114 | 63008 |
| CFQ 12Way LVM 30 Q Depth | 25775 | 44699 | 71292 |
| CFQ 8Way Meta 30 Q Depth | 18907 | 34041 | 57274 |
| CFQ 8Way LVM 256 Q Depth | 19469 | 34482 | 63085 |
| CFQ 12Way LVM 256 Q Depth | 25762 | 44837 | 71731 |
| CFQ 12Way Filer 256 Q Depth | 35113 | 45933 | 53900 |
| CFQ 12Way Filer 256T Q Depth | 35332 | 46238 | 58521 |
| CFQ 12Way Filer 30 Q Depth | 35554 | 49271 | 59931 |
| CFQ 12Way NAS 920 | 36339 | 41374 | 42581 |
| CFQ 12Way NAS 960 | 47384 | 51432 | 54864 |

Random Readers - CFQ Scheduler
Queue Depth Comparison

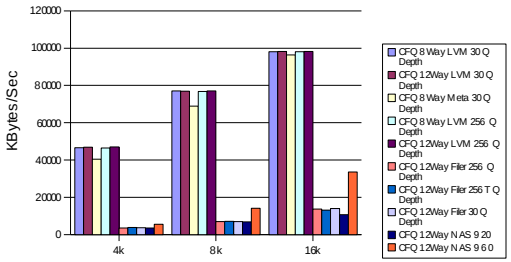


READER - CFQ Scheduler
Queue Depth Comparison

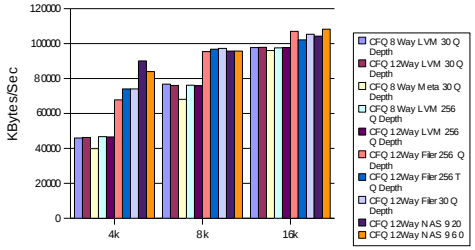


Writer - CFQ Scheduler
Queue Depth Comparison

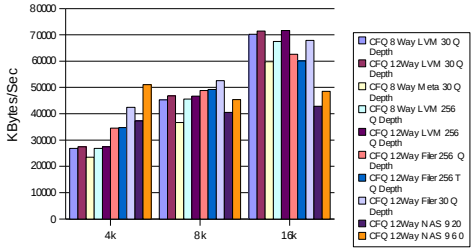


Writer - CFQ Scheduler
Queue Depth Comparison

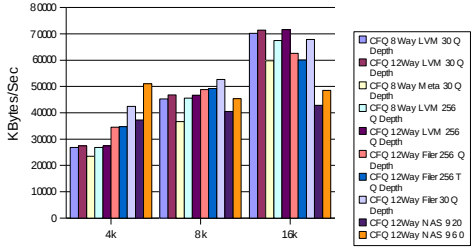


This is a comparison of disk configurations using the CFQ scheduler with 8 threads in 4,8, and 16k block file size transfer rates.

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

L12 = LVM over 12 spindles.
256 = Queue Depth.
Tuned = /etc/sysctl.conf see appendix.
NAS 920 = NetApp 920 Filer 7.23
NAS 960 = NetApp 960 Filer 7.23

| Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L8 Q Depth | 68163 | 75683 | 76261 | 65880 | 55926 |
| Noop 8K 256 L8 Q Depth | 68705 | 75981 | 76717 | 73216 | 53713 |
| AS 8K 256 L8 Q Depth | 66999 | 69743 | 69941 | 59393 | 46922 |
| CFQ 8K 256 L8 Q Depth | 68996 | 75871 | 76272 | 72943 | 52161 |
| Deadline 8K 30 L8 Q Depth | 68451 | 75804 | 76017 | 64896 | 58304 |
| Noop 8K 30 L8 Q Depth | 68752 | 75955 | 76464 | 70183 | 52794 |
| AS 8K 30 L8 Q Depth | 66598 | 69745 | 69823 | 66807 | 52286 |
| CFQ 8K 30 L8 Q Depth | 69064 | 75988 | 76798 | 70898 | 59725 |
| NAS CFQ 8K | 77707 | 93254 | 95791 | 100925 | 101298 |

| Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L8 Q Depth | 41255 | 45973 | 45093 | 45508 | 44397 |
| Noop 8K 256 L8 Q Depth | 47479 | 46126 | 45386 | 45306 | 44309 |
| AS 8K 256 L8 Q Depth | 44565 | 46339 | 44801 | 44972 | 44392 |
| CFQ 8K 256 L8 Q Depth | 42161 | 45791 | 45620 | 44954 | 44226 |
| Deadline 8K 30 L8 Q Depth | 45620 | 45173 | 45108 | 44802 | 44173 |
| Noop 8K 30 L8 Q Depth | 46484 | 45712 | 45187 | 44643 | 44259 |
| AS 8K 30 L8 Q Depth | 45738 | 45025 | 44939 | 44363 | 44215 |
| CFQ 8K 30 L8 Q Depth | 42281 | 45995 | 45288 | 45030 | 44100 |
| NAS CFQ 8K | 40966 | 41600 | 40530 | 43686 | 43462 |

| Random Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L8 Q Depth | 69848 | 76244 | 76778 | 67894 | 36277 |
| Noop 8K 256 L8 Q Depth | 69949 | 76030 | 77173 | 68007 | 33747 |
| AS 8K 256 L8 Q Depth | 45917 | 51938 | 57822 | 49052 | 22847 |
| CFQ 8K 256 L8 Q Depth | 69710 | 76321 | 76814 | 67885 | 36318 |
| Deadline 8K 30 L8 Q Depth | 69919 | 76155 | 77060 | 68295 | 36584 |
| Noop 8K 30 L8 Q Depth | 69765 | 76222 | 77262 | 67945 | 36237 |
| AS 8K 30 L8 Q Depth | 45073 | 51918 | 57875 | 51379 | 23661 |
| CFQ 8K 30 L8 Q Depth | 69718 | 76116 | 77091 | 67657 | 32414 |
| NAS CFQ 8K | 5315 | 6475 | 6874 | 7687 | 7605 |

| Random Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L8 Q Depth | 42384 | 43749 | 33836 | 28714 | 23898 |
| Noop 8K 256 L8 Q Depth | 45044 | 42887 | 38024 | 28532 | 23496 |
| AS 8K 256 L8 Q Depth | 43089 | 43493 | 35093 | 29594 | 24308 |
| CFQ 8K 256 L8 Q Depth | 44621 | 43938 | 34482 | 28885 | 24176 |
| Deadline 8K 30 L8 Q Depth | 42164 | 44011 | 35052 | 27897 | 24958 |
| Noop 8K 30 L8 Q Depth | 41759 | 42622 | 36056 | 28088 | 24459 |
| AS 8K 30 L8 Q Depth | 43928 | 43839 | 37090 | 27889 | 24950 |
| CFQ 8K 30 L8 Q Depth | 44630 | 44230 | 35114 | 27836 | 24883 |
| NAS CFQ 8K | 41178 | 40472 | 41374 | 38218 | 39554 |

Using Direct I/O no caching 8Hypers LVM Tuned.

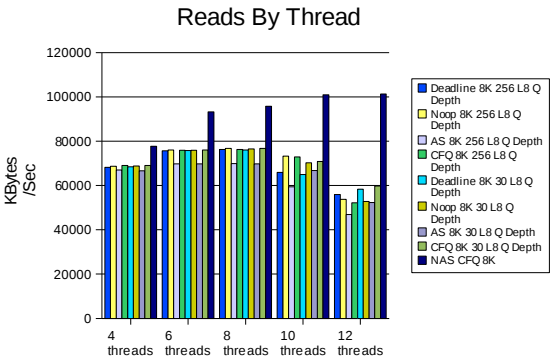


LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

L12 = LVM over 12 spindles.
256 = Queue Depth.
Tuned = /etc/sysctl.conf see appendix.

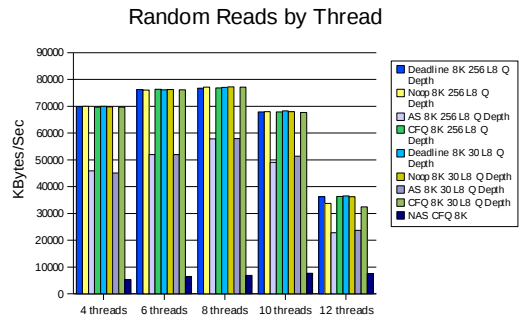


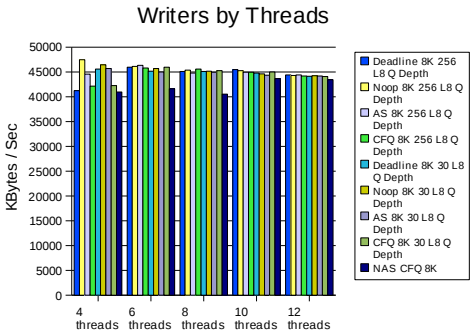


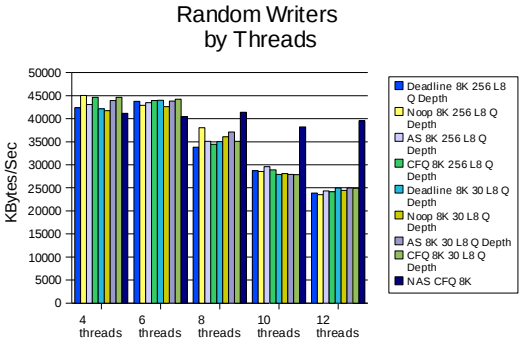

| Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 68138 | 74754 | 75963 | 73702 | 63434 |
| Noop 8K 256 L12 Q Depth | 67973 | 74828 | 75939 | 74831 | 63422 |
| AS 8K 256 L12 Q Depth | 65711 | 73083 | 76558 | 76069 | 62588 |
| CFQ 8K 256 L12 Q Depth | 67807 | 74654 | 75856 | 74547 | 63069 |
| Deadline 8K 30 L12 Q Depth | 67605 | 74976 | 76182 | 74353 | 63880 |
| Noop 8K 30 L12 Q Depth | 68182 | 74737 | 76196 | 70475 | 64484 |
| AS 8K 30 L12 Q Depth | 65806 | 72963 | 76762 | 71383 | 62623 |
| CFQ 8K 30 L12 Q Depth | 68006 | 74478 | 76049 | 75414 | 63629 |
| NAS CFQ 8K | 77707 | 93254 | 95791 | 100925 | 101298 |

| Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 47217 | 46855 | 46032 | 46111 | 45253 |
| Noop 8K 256 L12 Q Depth | 46521 | 46581 | 46681 | 46368 | 44913 |
| AS 8K 256 L12 Q Depth | 46836 | 46772 | 46415 | 46157 | 45139 |
| CFQ 8K 256 L12 Q Depth | 46734 | 46588 | 46655 | 46404 | 45935 |
| Deadline 8K 30 L12 Q Depth | 46867 | 46621 | 46788 | 46474 | 44935 |
| Noop 8K 30 L12 Q Depth | 46680 | 46915 | 46111 | 46018 | 45731 |
| AS 8K 30 L12 Q Depth | 46816 | 46731 | 46386 | 46392 | 45363 |
| CFQ 8K 30 L12 Q Depth | 46529 | 46858 | 46829 | 45705 | 45312 |
| NAS CFQ 8K | 40966 | 41600 | 40530 | 43686 | 43462 |

| Random Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 68834 | 75812 | 77356 | 71576 | 36804 |
| Noop 8K 256 L12 Q Depth | 69001 | 75488 | 77129 | 69613 | 39679 |
| AS 8K 256 L12 Q Depth | 45006 | 50465 | 55639 | 50744 | 25567 |
| CFQ 8K 256 L12 Q Depth | 68641 | 75444 | 77108 | 69732 | 35020 |
| Deadline 8K 30 L12 Q Depth | 68494 | 75903 | 76950 | 69646 | 35923 |
| Noop 8K 30 L12 Q Depth | 68888 | 75713 | 77149 | 69989 | 35870 |
| AS 8K 30 L12 Q Depth | 44939 | 50390 | 55885 | 48750 | 25137 |
| CFQ 8K 30 L12 Q Depth | 68814 | 75804 | 76938 | 70077 | 37063 |
| NAS CFQ 8K | 5315 | 6475 | 6874 | 7687 | 7605 |

| Random Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 45889 | 45592 | 44867 | 43382 | 39905 |
| Noop 8K 256 L12 Q Depth | 46081 | 45554 | 44818 | 42769 | 39872 |
| AS 8K 256 L12 Q Depth | 46068 | 45816 | 45028 | 43795 | 39198 |
| CFQ 8K 256 L12 Q Depth | 45638 | 45314 | 44837 | 43501 | 39489 |
| Deadline 8K 30 L12 Q Depth | 45856 | 45655 | 44957 | 43490 | 40199 |
| Noop 8K 30 L12 Q Depth | 46081 | 45619 | 45054 | 43740 | 39912 |
| AS 8K 30 L12 Q Depth | 46065 | 45593 | 45150 | 43998 | 40242 |
| CFQ 8K 30 L12 Q Depth | 45679 | 45798 | 44699 | 43627 | 39490 |
| NAS CFQ 8K | 41178 | 40472 | 41374 | 38218 | 39554 |

Using Direct I/O no caching LVM over 12 hypers no OS tuning.

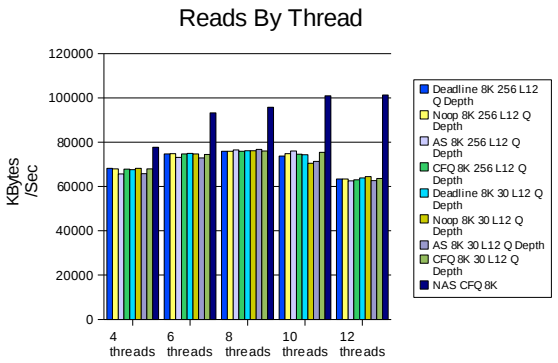


LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

L12 = LVM over 12 spindles.
256 = Queue Depth.
Tuned = /etc/sysctl.conf see appendix.

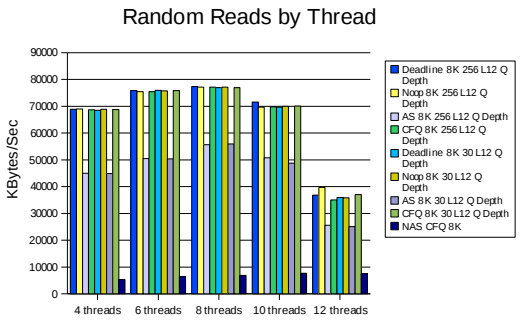


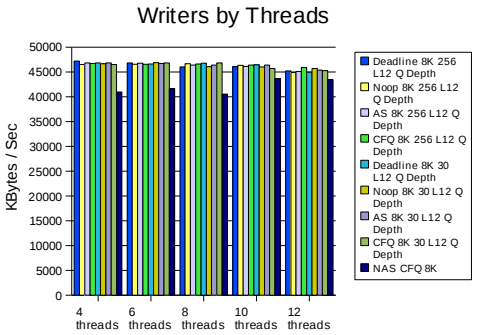


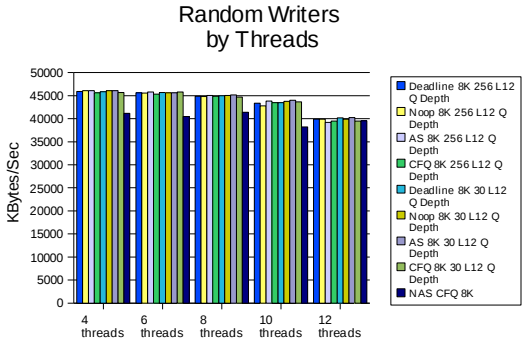

| Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 68138 | 74754 | 75963 | 73702 | 63434 |
| Noop 8K 256 L12 Q Depth | 67973 | 74828 | 75939 | 74831 | 63422 |
| AS 8K 256 L12 Q Depth | 65711 | 73083 | 76558 | 76069 | 62588 |
| CFQ 8K 256 L12 Q Depth | 67807 | 74654 | 75856 | 74547 | 63069 |
| Deadline 8K 256 L12 Tuned Q Dep | 68285 | 74616 | 76534 | 68812 | 63915 |
| Noop 8K 256 L12 Tuned Q Depth | 68484 | 75229 | 76758 | 75594 | 63760 |
| AS 8K 256 L12 Tuned Q Depth | 66688 | 73659 | 76892 | 69595 | 62867 |
| CFQ 8K 256 L12 Tuned Q Depth | 68047 | 75310 | 76277 | 74442 | 64016 |
| NAS CFQ Tuned 8K | 77707 | 93254 | 95791 | 100925 | 101298 |

| Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 47217 | 46855 | 46032 | 46111 | 45253 |
| Noop 8K 256 L12 Q Depth | 46521 | 46581 | 46681 | 46368 | 44913 |
| AS 8K 256 L12 Q Depth | 46836 | 46772 | 46415 | 46157 | 45139 |
| CFQ 8K 256 L12 Q Depth | 46734 | 46588 | 46655 | 46404 | 45935 |
| Deadline 8K 256 L12 Tuned Q Dep | 46694 | 46877 | 46575 | 46262 | 45469 |
| Noop 8K 256 L12 Tuned Q Depth | 46866 | 47149 | 46516 | 46093 | 45420 |
| AS 8K 256 L12 Tuned Q Depth | 46830 | 47040 | 46885 | 46578 | 45344 |
| CFQ 8K 256 L12 Tuned Q Depth | 46856 | 46945 | 46746 | 45974 | 45215 |
| NAS CFQ Tuned 8K | 40966 | 41600 | 40530 | 43686 | 43462 |

| Random Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 68834 | 75812 | 77356 | 71576 | 36804 |
| Noop 8K 256 L12 Q Depth | 69001 | 75488 | 77129 | 69613 | 39679 |
| AS 8K 256 L12 Q Depth | 45006 | 50465 | 55639 | 50744 | 25567 |
| CFQ 8K 256 L12 Q Depth | 68641 | 75444 | 77108 | 69732 | 35020 |
| Deadline 8K 256 L12 Tuned Q Dep | 69358 | 76144 | 77493 | 69388 | 36389 |
| Noop 8K 256 L12 Tuned Q Depth | 69287 | 76022 | 77455 | 70159 | 39677 |
| AS 8K 256 L12 Tuned Q Depth | 45089 | 50678 | 55810 | 48906 | 25335 |
| CFQ 8K 256 L12 Tuned Q Depth | 69378 | 76359 | 77504 | 69569 | 39827 |
| NAS CFQ Tuned 8K | 5315 | 6475 | 6874 | 7687 | 7605 |

| Random Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 45889 | 45592 | 44867 | 43382 | 39905 |
| Noop 8K 256 L12 Q Depth | 46081 | 45554 | 44818 | 42769 | 39872 |
| AS 8K 256 L12 Q Depth | 46068 | 45816 | 45028 | 43795 | 39198 |
| CFQ 8K 256 L12 Q Depth | 45638 | 45314 | 44837 | 43501 | 39489 |
| Deadline 8K 256 L12 Tuned Q Dep | 46640 | 46011 | 44911 | 43077 | 39343 |
| Noop 8K 256 L12 Tuned Q Depth | 46463 | 45842 | 45110 | 43785 | 40092 |
| AS 8K 256 L12 Tuned Q Depth | 46145 | 45945 | 45106 | 43669 | 40489 |
| CFQ 8K 256 L12 Tuned Q Depth | 46519 | 45656 | 44944 | 43531 | 39491 |
| NAS CFQ Tuned 8K | 41178 | 40472 | 41374 | 38218 | 39554 |

Using Direct I/O no caching LVM over 12 hypers, on the odd chance that OS NAS tuning affected
The output, we put in the tuning just to compare with little variants.

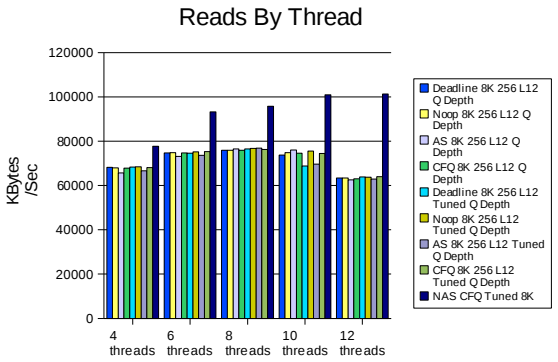


LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

L12 = LVM over 12 spindles.
256 = Queue Depth.
Tuned = /etc/sysctl.conf see appendix.

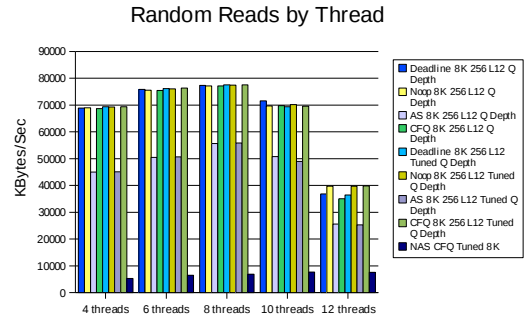


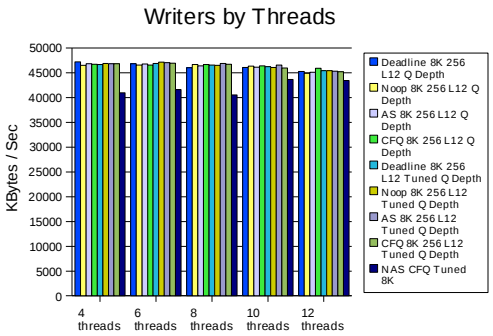


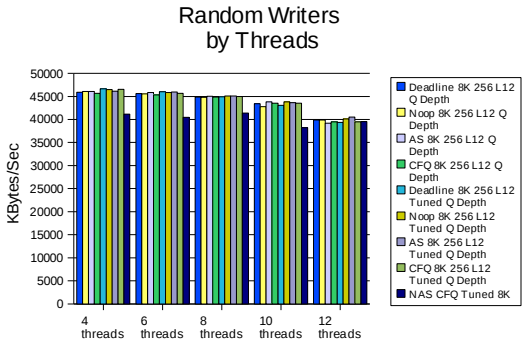

| Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 68138 | 74754 | 75963 | 73702 | 63434 |
| Noop 8K 256 L12 Q Depth | 67973 | 74828 | 75939 | 74831 | 63422 |
| AS 8K 256 L12 Q Depth | 65711 | 73083 | 76558 | 76069 | 62588 |
| CFQ 8K 256 L12 Q Depth | 67807 | 74654 | 75856 | 74547 | 63069 |
| Deadline 8K 256 L12 Filer Q Depth | 79534 | 93280 | 90232 | 95188 | 94567 |
| Noop 8K 256 L12 Filer Q Depth | 75613 | 87082 | 98914 | 97871 | 94531 |
| AS 8K 256 L12 Filer Q Depth | 79005 | 92484 | 96250 | 95440 | 96120 |
| CFQ 8K 256 L12 Filer Q Depth | 80751 | 94555 | 95504 | 98130 | 95167 |
| NAS CFQ 8K | 77707 | 93254 | 95791 | 100925 | 101298 |

| Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 47217 | 46855 | 46032 | 46111 | 45253 |
| Noop 8K 256 L12 Q Depth | 46521 | 46581 | 46681 | 46368 | 44913 |
| AS 8K 256 L12 Q Depth | 46836 | 46772 | 46415 | 46157 | 45139 |
| CFQ 8K 256 L12 Q Depth | 46734 | 46588 | 46655 | 46404 | 45935 |
| Deadline 8K 256 L12 Filer Q Depth | 41466 | 44209 | 44310 | 52087 | 53595 |
| Noop 8K 256 L12 Filer Q Depth | 40114 | 42708 | 46015 | 51534 | 51986 |
| AS 8K 256 L12 Filer Q Depth | 41194 | 44325 | 46497 | 54247 | 54925 |
| CFQ 8K 256 L12 Filer Q Depth | 41565 | 44254 | 48864 | 54283 | 53324 |
| NAS CFQ 8K | 40966 | 41600 | 40530 | 43686 | 43462 |

| Random Reads | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 68834 | 75812 | 77356 | 71576 | 36804 |
| Noop 8K 256 L12 Q Depth | 69001 | 75488 | 77129 | 69613 | 39679 |
| AS 8K 256 L12 Q Depth | 45006 | 50465 | 55639 | 50744 | 25567 |
| CFQ 8K 256 L12 Q Depth | 68641 | 75444 | 77108 | 69732 | 35020 |
| Deadline 8K 256 L12 Filer Q Depth | 7100 | 6690 | 6440 | 7479 | 8064 |
| Noop 8K 256 L12 Filer Q Depth | 5980 | 6005 | 6376 | 7522 | 7806 |
| AS 8K 256 L12 Filer Q Depth | 6678 | 6852 | 7045 | 7552 | 8253 |
| CFQ 8K 256 L12 Filer Q Depth | 6759 | 6461 | 6894 | 7446 | 7400 |
| NAS CFQ 8K | 5315 | 6475 | 6874 | 7687 | 7605 |

| Random Writers | 4 threads | 6 threads | 8 threads | 10 threads | 12 threads |
|---|---|---|---|---|---|
| Deadline 8K 256 L12 Q Depth | 45889 | 45592 | 44867 | 43382 | 39905 |
| Noop 8K 256 L12 Q Depth | 46081 | 45554 | 44818 | 42769 | 39872 |
| AS 8K 256 L12 Q Depth | 46068 | 45816 | 45028 | 43795 | 39198 |
| CFQ 8K 256 L12 Q Depth | 45638 | 45314 | 44837 | 43501 | 39489 |
| Deadline 8K 256 L12 Filer Q Depth | 40344 | 42034 | 45831 | 48298 | 49051 |
| Noop 8K 256 L12 Filer Q Depth | 38851 | 43311 | 44802 | 48072 | 48894 |
| AS 8K 256 L12 Filer Q Depth | 39333 | 45540 | 46532 | 48479 | 49559 |
| CFQ 8K 256 L12 Filer Q Depth | 40275 | 42142 | 45933 | 48089 | 49136 |
| NAS CFQ 8K | 41178 | 40472 | 41374 | 38218 | 39554 |

Using Direct I/O no caching LVM over 12 hypers no OS tuning.

## Reads By Thread

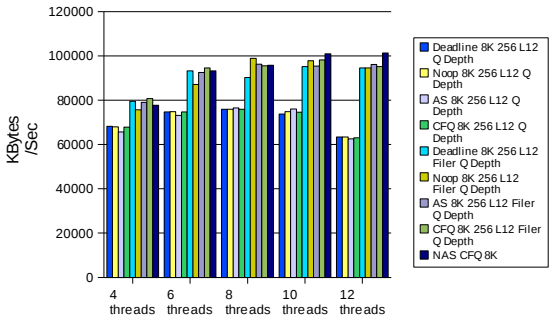


## Random Reads by Thread

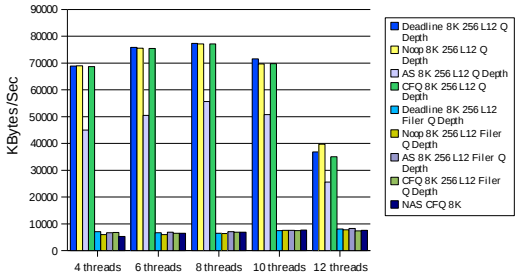


## Writers by Threads

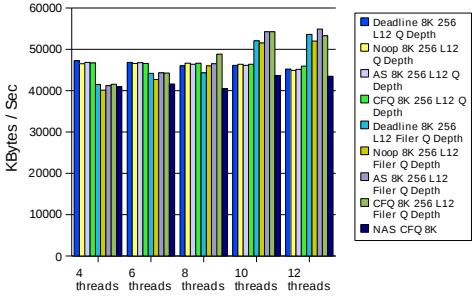


## Random Writers by Threads

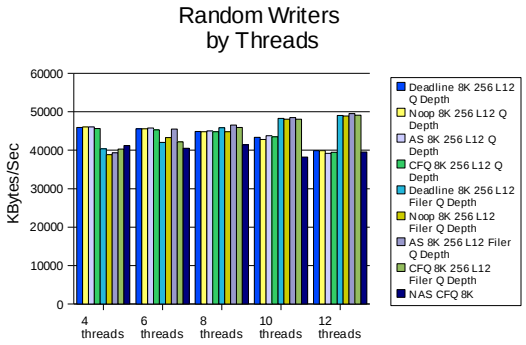


LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

L12 = LVM over 12 spindles.
256 = Queue Depth.
Tuned = /etc/sysctl.conf see appendix.

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 26826 | 45108 | 69657 |
| CFQ 30 Queue Depth | 26793 | 45288 | 70207 |
| AS 30 Queue Depth | 27882 | 44939 | 71064 |
| NOOP 30 Queue Depth | 27462 | 45187 | 69728 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 27479 | 44913 | 67678 |
| CFQ 30 Queue Depth | 27782 | 44773 | 67040 |
| AS 30 Queue Depth | 27870 | 44934 | 67304 |
| NOOP 30 Queue Depth | 27824 | 44703 | 67247 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 46356 | 76017 | 97668 |
| CFQ 30 Queue Depth | 45944 | 76798 | 97727 |
| AS 30 Queue Depth | 45793 | 69823 | 87065 |
| NOOP 30 Queue Depth | 46582 | 76464 | 97751 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 46448 | 76183 | 97863 |
| CFQ 30 Queue Depth | 46446 | 76734 | 97926 |
| AS 30 Queue Depth | 45736 | 70083 | 87402 |
| NOOP 30 Queue Depth | 47145 | 76794 | 97797 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 46627 | 77060 | 97917 |
| CFQ 30 Queue Depth | 46576 | 77091 | 98127 |
| AS 30 Queue Depth | 35756 | 57875 | 80114 |
| NOOP 30 Queue Depth | 46674 | 77262 | 98142 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 19514 | 35052 | 62907 |
| CFQ 30 Queue Depth | 18825 | 35114 | 63008 |
| AS 30 Queue Depth | 18381 | 37090 | 65853 |
| NOOP 30 Queue Depth | 18611 | 36056 | 62901 |

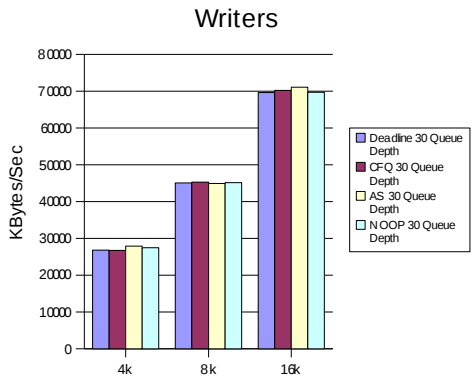


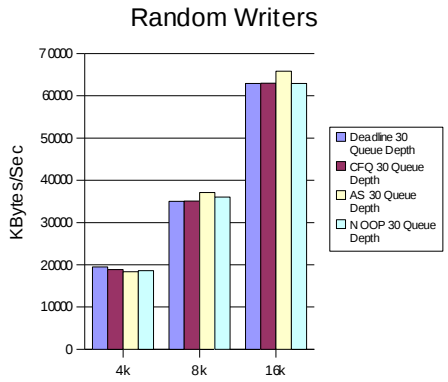


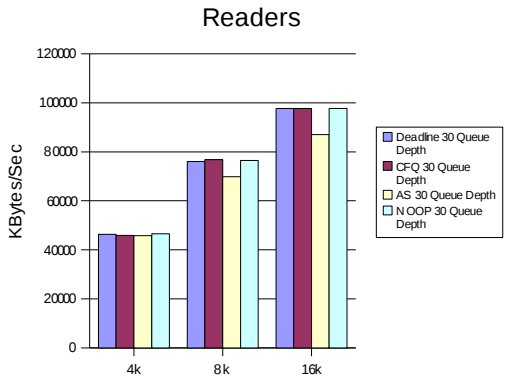


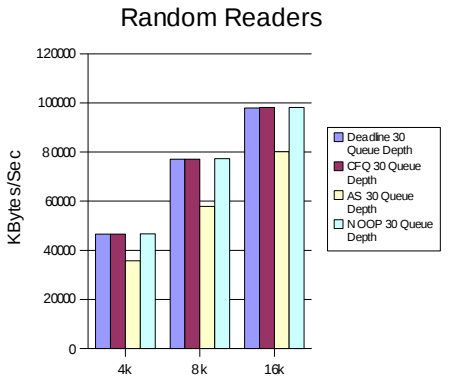


Using Direct I/O (no caching) on 8 partition LVM device with SAN card queue depth set to 30.
LVM on 8 Hypers stripped using 128K stripe size.

LVM – refers to EXT3 laid on an LVM volume made
up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA
'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer
'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning
in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 30 Queue Depth | | 27378 | 46788 | 70938 |
| CFQ 30 Queue Depth | | 27478 | 46829 | 71461 |
| AS 30 Queue Depth | | 27496 | 46386 | 72211 |
| NOOP 30 Queue Depth | | 27461 | 46111 | 71244 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 30 Queue Depth | | 28008 | 47508 | 69568 |
| CFQ 30 Queue Depth | | 28065 | 47287 | 69795 |
| AS 30 Queue Depth | | 28094 | 47693 | 70278 |
| NOOP 30 Queue Depth | | 28148 | 47528 | 69795 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 30 Queue Depth | | 46466 | 76182 | 97864 |
| CFQ 30 Queue Depth | | 46228 | 76049 | 97845 |
| AS 30 Queue Depth | | 47742 | 76762 | 96048 |
| NOOP 30 Queue Depth | | 46105 | 76196 | 97905 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 30 Queue Depth | | 46554 | 76294 | 98085 |
| CFQ 30 Queue Depth | | 46221 | 75801 | 98039 |
| AS 30 Queue Depth | | 48082 | 77160 | 96250 |
| NOOP 30 Queue Depth | | 46345 | 76138 | 98049 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 30 Queue Depth | | 47036 | 76950 | 98208 |
| CFQ 30 Queue Depth | | 46895 | 76938 | 98187 |
| AS 30 Queue Depth | | 34317 | 55885 | 79030 |
| NOOP 30 Queue Depth | | 46942 | 77149 | 98175 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 30 Queue Depth | | 25779 | 44957 | 71687 |
| CFQ 30 Queue Depth | | 25775 | 44699 | 71292 |
| AS 30 Queue Depth | | 25520 | 45150 | 72113 |
| NOOP 30 Queue Depth | | 25590 | 45054 | 71149 |

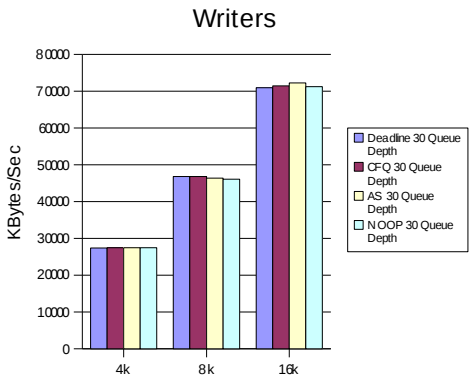


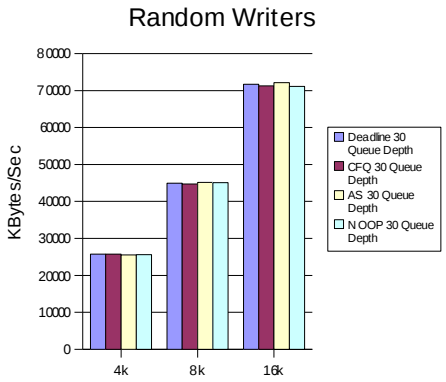


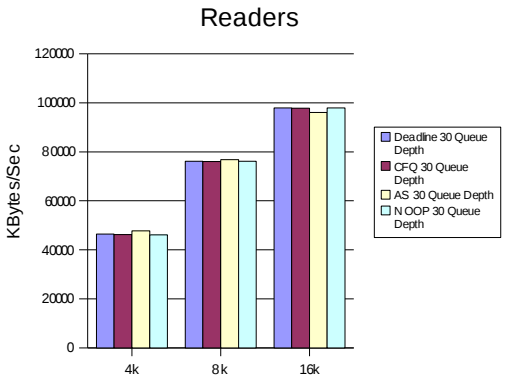


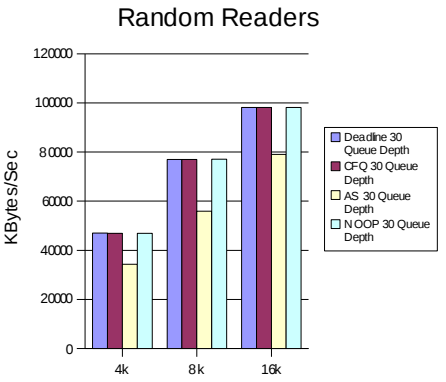


Using Direct I/O (no caching) on 8 partition LVM device with SAN card queue depth set to 30.
LVM on 12 Hypers stripped using 128K stripe size.

LVM – refers to EXT3 laid on an LVM volume made
up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA
'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer
'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning
in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 22585 | 36520 | 59125 |
| CFQ 30 Queue Depth | 23470 | 36635 | 59758 |
| AS 30 Queue Depth | 20544 | 36424 | 59982 |
| NOOP 30 Queue Depth | 20501 | 36449 | 59652 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 20642 | 36667 | 60312 |
| CFQ 30 Queue Depth | 22086 | 36753 | 60330 |
| AS 30 Queue Depth | 20636 | 36730 | 60163 |
| NOOP 30 Queue Depth | 20635 | 36689 | 60184 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 39723 | 68228 | 96337 |
| CFQ 30 Queue Depth | 39831 | 68172 | 96147 |
| AS 30 Queue Depth | 21557 | 35192 | 52095 |
| NOOP 30 Queue Depth | 39683 | 68403 | 96317 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 39776 | 68172 | 96367 |
| CFQ 30 Queue Depth | 39776 | 68103 | 96318 |
| AS 30 Queue Depth | 21549 | 35193 | 51888 |
| NOOP 30 Queue Depth | 39830 | 68189 | 96439 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 40436 | 69157 | 96614 |
| CFQ 30 Queue Depth | 40428 | 68958 | 96365 |
| AS 30 Queue Depth | 38689 | 65242 | 90260 |
| NOOP 30 Queue Depth | 40479 | 69162 | 96637 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 18874 | 33974 | 57284 |
| CFQ 30 Queue Depth | 18907 | 34041 | 57274 |
| AS 30 Queue Depth | 18872 | 34078 | 57348 |
| NOOP 30 Queue Depth | 18844 | 34110 | 57243 |

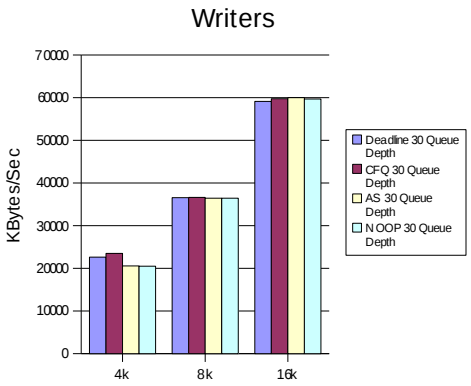


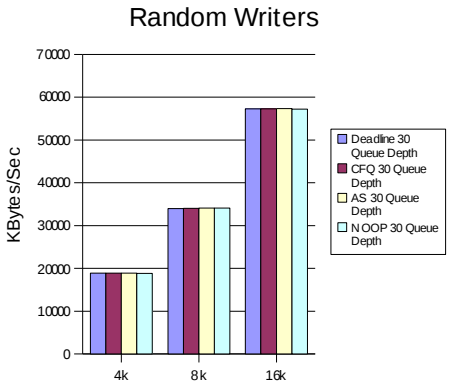


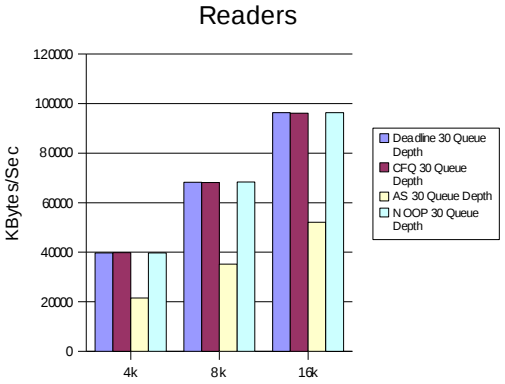


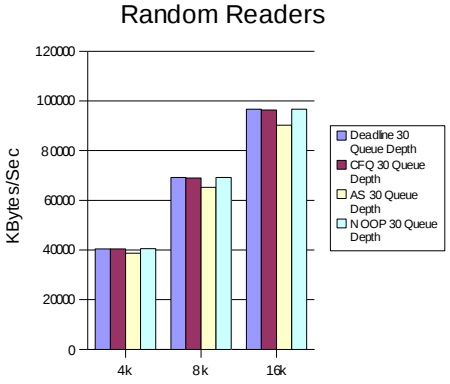


Using Direct I/O (no caching) on 8way Meta Device withSAN card queue depth set to 30

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 38135 | 50179 | 64005 |
| CFQ 30 Queue Depth | 42450 | 52634 | 67859 |
| AS 30 Queue Depth | 35933 | 46699 | 54526 |
| NOOP 30 Queue Depth | 36459 | 49208 | 62457 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 38569 | 55831 | 71138 |
| CFQ 30 Queue Depth | 40360 | 57465 | 72523 |
| AS 30 Queue Depth | 38382 | 48999 | 65368 |
| NOOP 30 Queue Depth | 37953 | 55459 | 71168 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 69671 | 94097 | 105258 |
| CFQ 30 Queue Depth | 73993 | 97338 | 105422 |
| AS 30 Queue Depth | 18954 | 26032 | 35184 |
| NOOP 30 Queue Depth | 72392 | 96501 | 104990 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 75602 | 109909 | 128820 |
| CFQ 30 Queue Depth | 77958 | 104474 | 124532 |
| AS 30 Queue Depth | 20372 | 29709 | 40282 |
| NOOP 30 Queue Depth | 77304 | 108206 | 127922 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 3539 | 6781 | 13251 |
| CFQ 30 Queue Depth | 3644 | 6953 | 13948 |
| AS 30 Queue Depth | 2723 | 4963 | 10496 |
| NOOP 30 Queue Depth | 3199 | 6943 | 13022 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 33843 | 46317 | 56023 |
| CFQ 30 Queue Depth | 35554 | 49271 | 59931 |
| AS 30 Queue Depth | 32611 | 48477 | 57511 |
| NOOP 30 Queue Depth | 35516 | 46456 | 55230 |

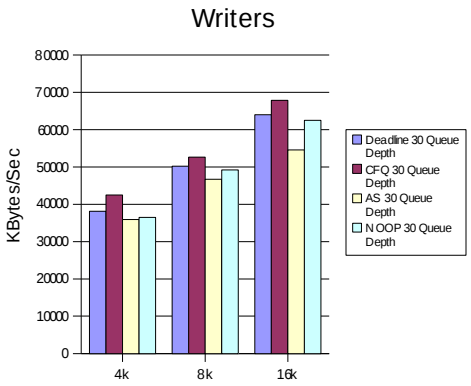


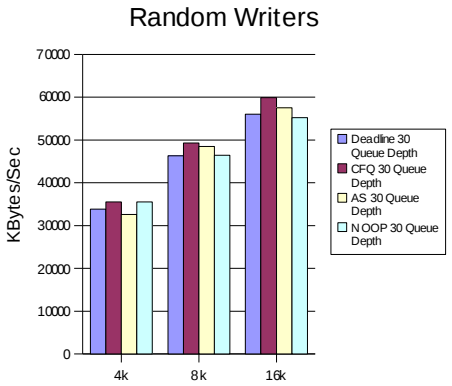


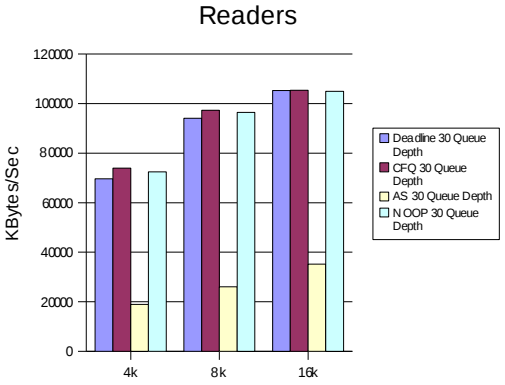


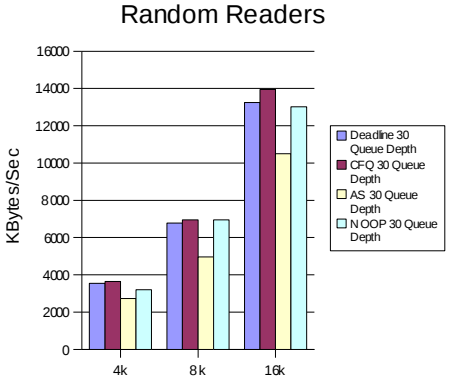


Using Direct I/O (no caching) on 12 way NetApp Filer LUN Device withSAN card queue depth set to 30

LVM – refers to EXT3 laid on an LVM volume made
up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA
'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer
'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning
in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 43126 | 64590 | 75184 |
| CFQ 30 Queue Depth | 43469 | 64179 | 89278 |
| AS 30 Queue Depth | 44032 | 65098 | 89620 |
| NOOP 30 Queue Depth | 44409 | 65012 | 86782 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 45815 | 71372 | 87368 |
| CFQ 30 Queue Depth | 44037 | 71296 | 96110 |
| AS 30 Queue Depth | 46816 | 72582 | 96805 |
| NOOP 30 Queue Depth | 46857 | 73006 | 98203 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 82019 | 125940 | 156688 |
| CFQ 30 Queue Depth | 83750 | 130634 | 156455 |
| AS 30 Queue Depth | 19927 | 34583 | 57327 |
| NOOP 30 Queue Depth | 84606 | 131118 | 158787 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 84284 | 127287 | 158804 |
| CFQ 30 Queue Depth | 83804 | 131762 | 159418 |
| AS 30 Queue Depth | 19992 | 34816 | 57295 |
| NOOP 30 Queue Depth | 84900 | 132350 | 159167 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 79501 | 112950 | 155903 |
| CFQ 30 Queue Depth | 84194 | 131867 | 159034 |
| AS 30 Queue Depth | 69311 | 108611 | 141780 |
| NOOP 30 Queue Depth | 84968 | 132687 | 159708 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 30 Queue Depth | 43450 | 66451 | 86227 |
| CFQ 30 Queue Depth | 42203 | 65639 | 86661 |
| AS 30 Queue Depth | 43588 | 69301 | 85869 |
| NOOP 30 Queue Depth | 44264 | 68412 | 91689 |

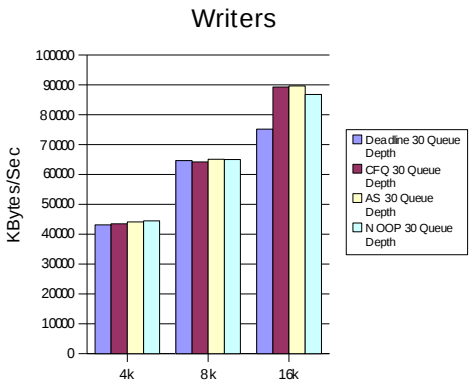


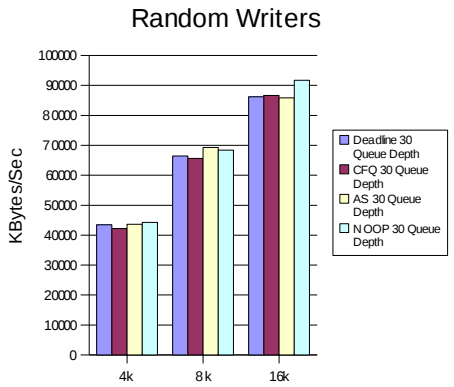


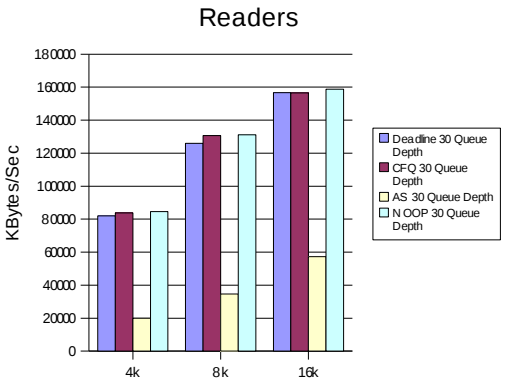


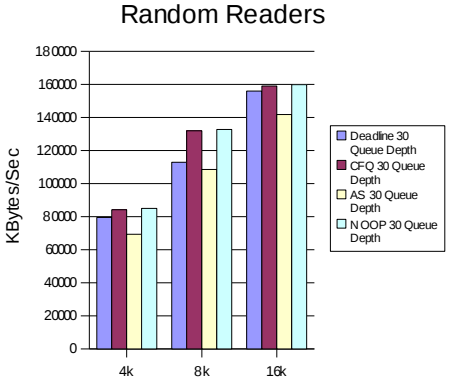


Using Direct I/O (no caching) on 12 way NetApp Filer LUN Device withSAN card queue depth set to 30

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

| **Readers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 12Way 920 Filer 256 Q Depth | 67804 | 95504 | 107153 |
| CFQ 12Way 920 Filer 256T Q Depth | 74118 | 96838 | 102163 |
| CFQ 12Way 920 Filer 30 Q Depth | 73993 | 97338 | 105422 |
| CFQ 12Way 960 Filer 30 Q Depth | 83750 | 130634 | 156455 |
| CFQ 12Way 960 Filer 30 Q Depth 1G | 75810 | 104127 | 124193 |
| CFQ 12Way NAS 920 | 90111 | 95791 | 104344 |
| CFQ 12Way NAS 960 | 83969 | 95789 | 108245 |

| **Random Readers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 12Way 920 Filer 256 Q Depth | 3466 | 6894 | 13704 |
| CFQ 12Way 920 Filer 256T Q Depth | 3802 | 7120 | 13115 |
| CFQ 12Way Filer 30 Q Depth | 3644 | 6953 | 13948 |
| CFQ 12Way 960 Filer 30 Q Depth | 84194 | 131867 | 159034 |
| CFQ 12Way 960 Filer 30 Q Depth 1G | 5495 | 10428 | 20573 |
| CFQ 12Way NAS 920 | 3562 | 6874 | 10664 |
| CFQ 12Way NAS 960 | 5542 | 14149 | 33564 |

| **Writers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 12Way 920 Filer 256 Q Depth | 34540 | 48864 | 62560 |
| CFQ 12Way 920 Filer 256T Q Depth | 34696 | 49236 | 60110 |
| CFQ 12Way 920 Filer 30 Q Depth | 42450 | 52634 | 67859 |
| CFQ 12Way 960 Filer 30 Q Depth | 43469 | 64179 | 89278 |
| CFQ 12Way 960 Filer 30 Q Depth 1G | 44391 | 61280 | 85344 |
| CFQ 12Way NAS 920 | 37313 | 40530 | 42824 |
| CFQ 12Way NAS 960 | 51026 | 45361 | 48543 |

| **Random Writers** | **4k** | **8k** | **16k** |
|---|---|---|---|
| CFQ 12Way 920 Filer 256 Q Depth | 35113 | 45933 | 53900 |
| CFQ 12Way 920 Filer 256T Q Depth | 35332 | 46238 | 58521 |
| CFQ 12Way 920 Filer 30 Q Depth | 35554 | 49271 | 59931 |
| CFQ 12Way 960 Filer 30 Q Depth | 42203 | 65639 | 86661 |
| CFQ 12Way 960 Filer 30 Q Depth 1G | 39991 | 60024 | 75157 |
| CFQ 12Way NAS 920 | 36339 | 41374 | 42581 |
| CFQ 12Way NAS 960 | 47384 | 51432 | 54864 |

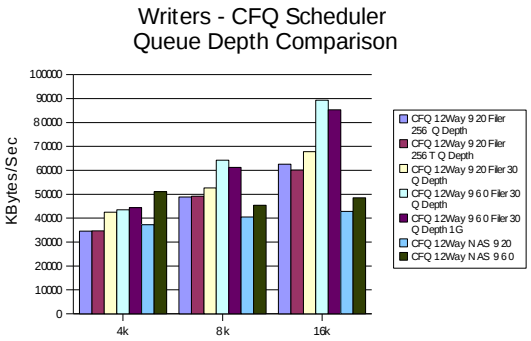


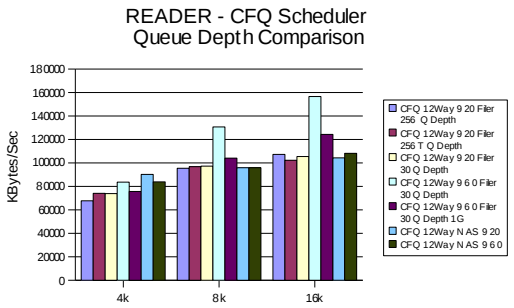


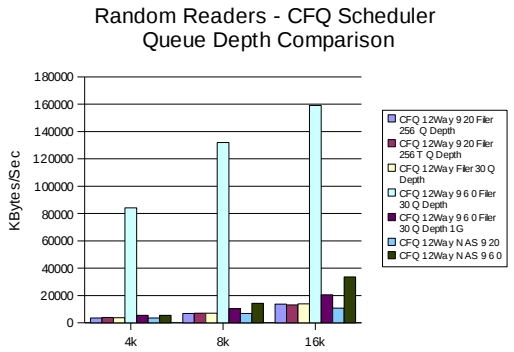


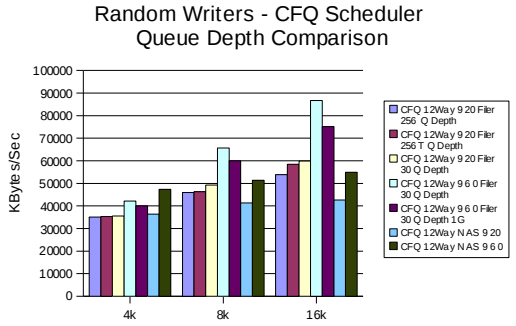


This is a comparison of disk configurations using up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA

'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

960 Filer – refers to the same 12 disk aggregate that the 920 filer used. An in place filer head replacement was done.

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 39488 | 38837 | 42803 |
| CFQ NAS | 37313 | 40530 | 42824 |
| AS NAS | 42780 | 38490 | 42385 |
| NOOP NAS | 41354 | 40677 | 42307 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 41674 | 43598 | 44401 |
| CFQ NAS | 41260 | 44943 | 45608 |
| AS NAS | 41751 | 43458 | 44414 |
| NOOP NAS | 38541 | 45204 | 43800 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 90854 | 94556 | 100046 |
| CFQ NAS | 90111 | 95791 | 104344 |
| AS NAS | 91657 | 94717 | 99395 |
| NOOP NAS | 89460 | 97552 | 102847 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 92317 | 95452 | 102360 |
| CFQ NAS | 93785 | 100079 | 108139 |
| AS NAS | 91340 | 95132 | 101769 |
| NOOP NAS | 90400 | 99674 | 106190 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 3312 | 6216 | 11678 |
| CFQ NAS | 3562 | 6874 | 10664 |
| AS NAS | 3680 | 6191 | 12031 |
| NOOP NAS | 3437 | 5873 | 13299 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 36018 | 39785 | 40964 |
| CFQ NAS | 36339 | 41374 | 42581 |
| AS NAS | 36274 | 40091 | 40830 |
| NOOP NAS | 36138 | 41943 | 42601 |

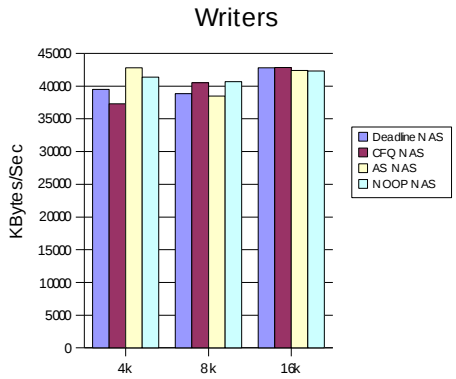


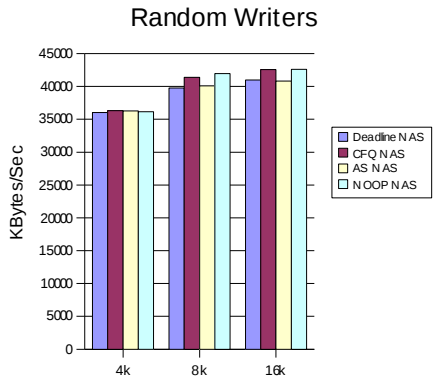


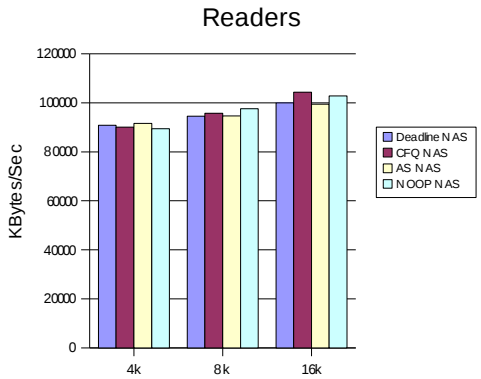


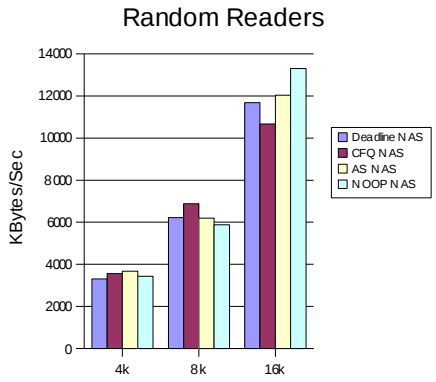


Using Direct I/O no caching.

NAS 920 Filer

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 47848 | 44553 | 48019 |
| CFQ NAS | 51026 | 45361 | 48543 |
| AS NAS | 51037 | 44519 | 47259 |
| NOOP NAS | 50147 | 42295 | 46810 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 49584 | 53604 | 54617 |
| CFQ NAS | 49942 | 53100 | 55919 |
| AS NAS | 51227 | 53787 | 55319 |
| NOOP NAS | 51168 | 52918 | 56542 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 84431 | 95959 | 99834 |
| CFQ NAS | 83969 | 95789 | 108245 |
| AS NAS | 85666 | 97045 | 107904 |
| NOOP NAS | 86902 | 93818 | 106386 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 87190 | 98912 | 109876 |
| CFQ NAS | 84770 | 98508 | 110126 |
| AS NAS | 86534 | 99065 | 110680 |
| NOOP NAS | 87034 | 96007 | 109156 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 4261 | 15539 | 25415 |
| CFQ NAS | 5542 | 14149 | 33564 |
| AS NAS | 6810 | 14521 | 35203 |
| NOOP NAS | 5908 | 12005 | 26622 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline NAS | 47445 | 51391 | 52941 |
| CFQ NAS | 47384 | 51432 | 54864 |
| AS NAS | 47899 | 51638 | 54628 |
| NOOP NAS | 50378 | 51165 | 55269 |

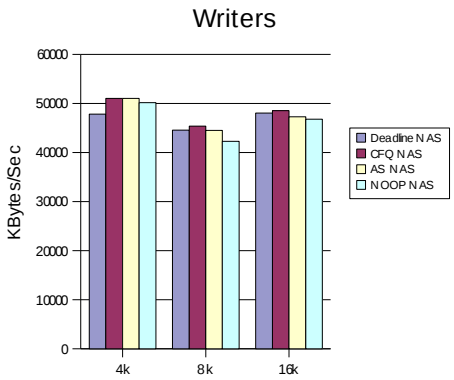


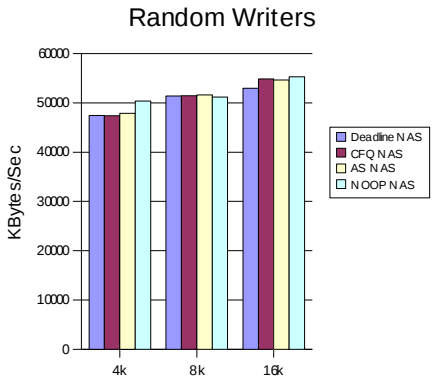


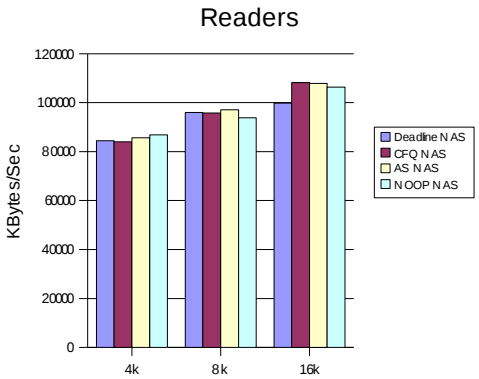


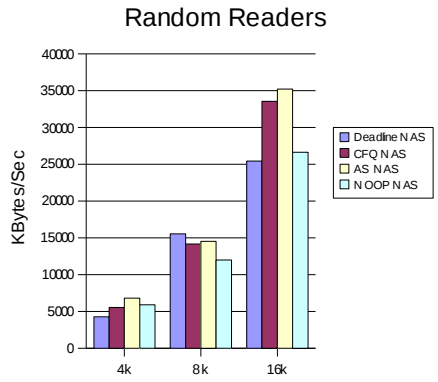


Using Direct I/O no caching.

NAS 960 Filer

| **Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 26813 | 45093 | 69000 |
| CFQ 256 Queue Depth | | 26811 | 45620 | 67512 |
| AS 256 Queue Depth | | 27091 | 44801 | 71026 |
| NOOP 256 Queue Depth | | 26942 | 45386 | 71376 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 27728 | 45003 | 68620 |
| CFQ 256 Queue Depth | | 27690 | 44296 | 68821 |
| AS 256 Queue Depth | | 27910 | 45131 | 67965 |
| NOOP 256 Queue Depth | | 27453 | 45030 | 68548 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46711 | 76261 | 97558 |
| CFQ 256 Queue Depth | | 46682 | 76272 | 97658 |
| AS 256 Queue Depth | | 45888 | 69941 | 87235 |
| NOOP 256 Queue Depth | | 43108 | 76717 | 97685 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46541 | 76268 | 97827 |
| CFQ 256 Queue Depth | | 46646 | 76079 | 97796 |
| AS 256 Queue Depth | | 45886 | 70075 | 86901 |
| NOOP 256 Queue Depth | | 45887 | 76700 | 97924 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46642 | 76778 | 98086 |
| CFQ 256 Queue Depth | | 46503 | 76814 | 98125 |
| AS 256 Queue Depth | | 35825 | 57822 | 80052 |
| NOOP 256 Queue Depth | | 46350 | 77173 | 98134 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 19565 | 33836 | 61943 |
| CFQ 256 Queue Depth | | 19469 | 34482 | 63085 |
| AS 256 Queue Depth | | 18759 | 35093 | 59232 |
| NOOP 256 Queue Depth | | 19190 | 38024 | 65667 |

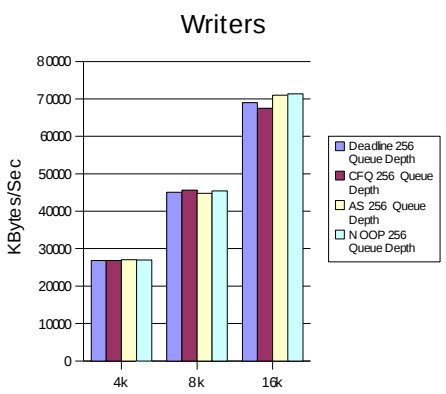


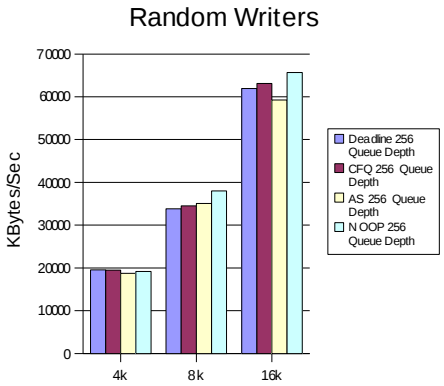


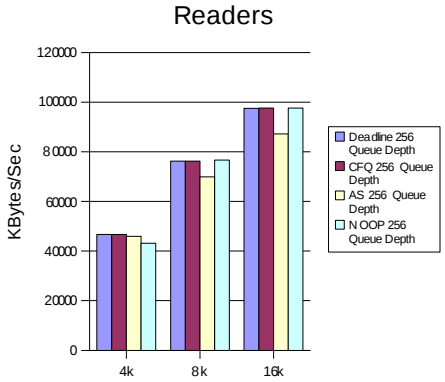


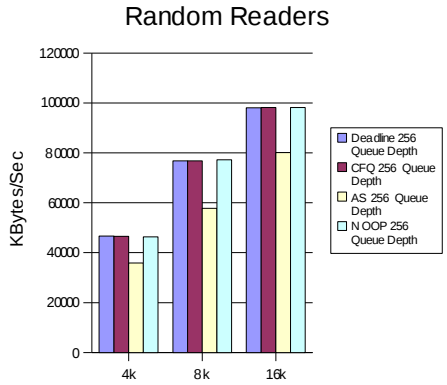


Using Direct I/O no caching.
LVM on 8 Hypers stripped using 128K stripe size.

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 27543 | 46032 | 71688 |
| CFQ 256 Queue Depth | | 27465 | 46655 | 71622 |
| AS 256 Queue Depth | | 27492 | 46415 | 71825 |
| NOOP 256 Queue Depth | | 27430 | 46681 | 71442 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 28132 | 47388 | 70477 |
| CFQ 256 Queue Depth | | 28132 | 47664 | 70619 |
| AS 256 Queue Depth | | 28034 | 47688 | 70160 |
| NOOP 256 Queue Depth | | 28039 | 47652 | 70625 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46180 | 75963 | 97829 |
| CFQ 256 Queue Depth | | 46468 | 75856 | 97818 |
| AS 256 Queue Depth | | 47719 | 76558 | 96238 |
| NOOP 256 Queue Depth | | 46115 | 75939 | 97576 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46481 | 76448 | 98000 |
| CFQ 256 Queue Depth | | 46385 | 76264 | 98139 |
| AS 256 Queue Depth | | 48040 | 76650 | 96209 |
| NOOP 256 Queue Depth | | 46178 | 76122 | 98095 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 47019 | 77356 | 98123 |
| CFQ 256 Queue Depth | | 47027 | 77108 | 98231 |
| AS 256 Queue Depth | | 34327 | 55639 | 79124 |
| NOOP 256 Queue Depth | | 46994 | 77129 | 98198 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 25694 | 44867 | 71917 |
| CFQ 256 Queue Depth | | 25762 | 44837 | 71731 |
| AS 256 Queue Depth | | 25657 | 45028 | 71276 |
| NOOP 256 Queue Depth | | 25686 | 44818 | 71856 |

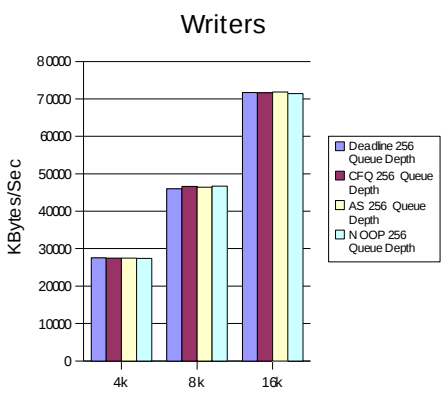


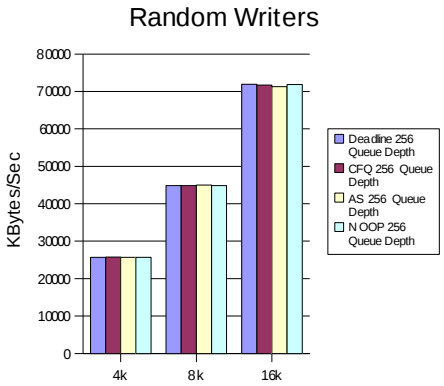


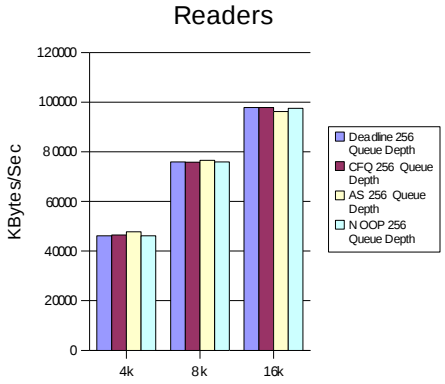


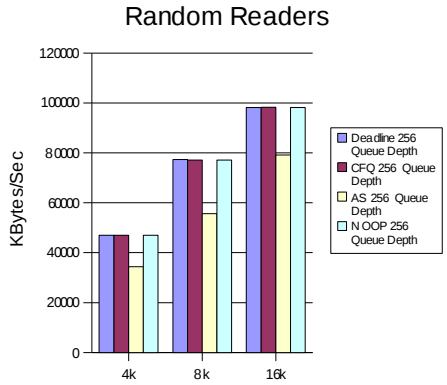


Using Direct I/O no caching.
LVM on 12 Hypers stripped using 128K stripe size.

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 27573 | 46575 | 71779 |
| CFQ 256 Queue Depth | | 27707 | 46746 | 71133 |
| AS 256 Queue Depth | | 27658 | 46885 | 71217 |
| NOOP 256 Queue Depth | | 27560 | 46516 | 71238 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 28193 | 47382 | 70583 |
| CFQ 256 Queue Depth | | 28165 | 47423 | 70660 |
| AS 256 Queue Depth | | 28101 | 47888 | 70442 |
| NOOP 256 Queue Depth | | 28220 | 47828 | 70551 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46515 | 76534 | 97890 |
| CFQ 256 Queue Depth | | 46417 | 76277 | 97819 |
| AS 256 Queue Depth | | 47787 | 76892 | 96221 |
| NOOP 256 Queue Depth | | 46372 | 76758 | 97869 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 46467 | 76796 | 98108 |
| CFQ 256 Queue Depth | | 46327 | 76480 | 98075 |
| AS 256 Queue Depth | | 47912 | 77164 | 96449 |
| NOOP 256 Queue Depth | | 46398 | 76848 | 97981 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 47037 | 77493 | 98246 |
| CFQ 256 Queue Depth | | 47042 | 77504 | 97850 |
| AS 256 Queue Depth | | 34371 | 55810 | 79106 |
| NOOP 256 Queue Depth | | 47059 | 77455 | 98217 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 25733 | 44911 | 71928 |
| CFQ 256 Queue Depth | | 25781 | 44944 | 72243 |
| AS 256 Queue Depth | | 25756 | 45106 | 71469 |
| NOOP 256 Queue Depth | | 25738 | 45110 | 71659 |

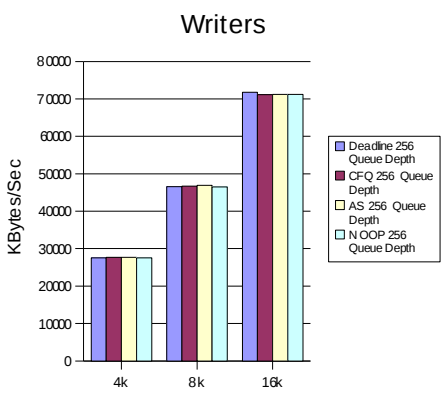


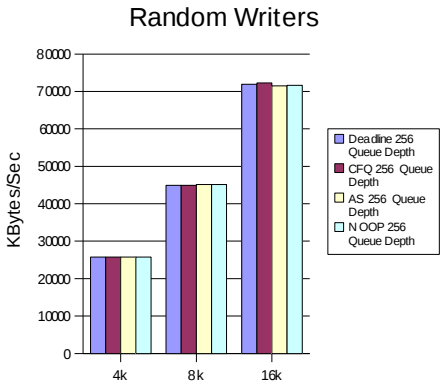


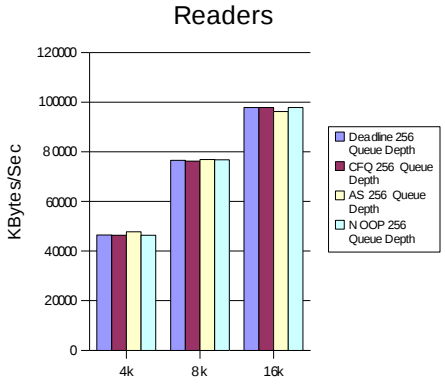


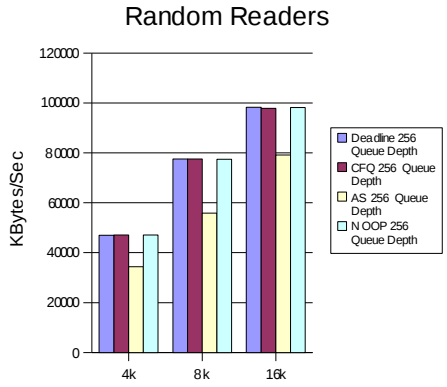


Using Direct I/O no caching.
LVM on 12 Hypers stripped using 128K stripe size.

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 36231 | 44310 | 62864 |
| CFQ 256 Queue Depth | | 34540 | 48864 | 62560 |
| AS 256 Queue Depth | | 35838 | 46497 | 58944 |
| NOOP 256 Queue Depth | | 34570 | 46015 | 55713 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 38819 | 54004 | 70268 |
| CFQ 256 Queue Depth | | 36794 | 55340 | 69995 |
| AS 256 Queue Depth | | 39472 | 53732 | 65152 |
| NOOP 256 Queue Depth | | 39546 | 56320 | 63638 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 73223 | 90232 | 103264 |
| CFQ 256 Queue Depth | | 67804 | 95504 | 107153 |
| AS 256 Queue Depth | | 73897 | 96250 | 99450 |
| NOOP 256 Queue Depth | | 72506 | 98914 | 92977 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 78178 | 105384 | 117272 |
| CFQ 256 Queue Depth | | 74833 | 105550 | 102064 |
| AS 256 Queue Depth | | 77432 | 103824 | 102747 |
| NOOP 256 Queue Depth | | 76592 | 112327 | 120950 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 3203 | 6440 | 13910 |
| CFQ 256 Queue Depth | | 3466 | 6894 | 13704 |
| AS 256 Queue Depth | | 3601 | 7045 | 13520 |
| NOOP 256 Queue Depth | | 3672 | 6376 | 12479 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** | |
|---|---|---|---|---|
| Deadline 256 Queue Depth | | 34540 | 45831 | 54786 |
| CFQ 256 Queue Depth | | 35113 | 45933 | 53900 |
| AS 256 Queue Depth | | 35734 | 46532 | 60201 |
| NOOP 256 Queue Depth | | 33995 | 44802 | 53074 |

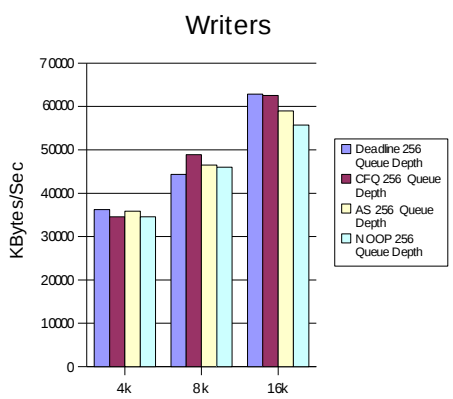


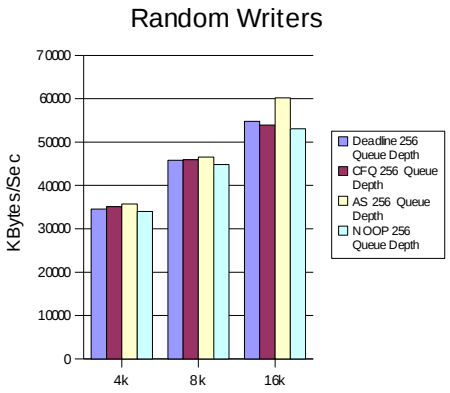


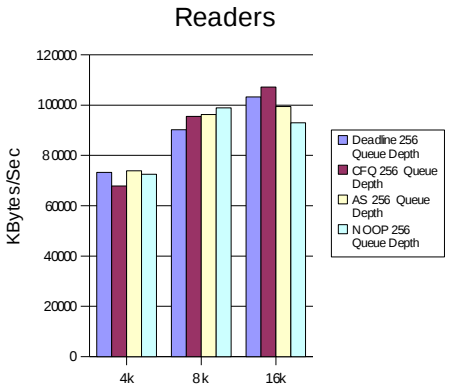


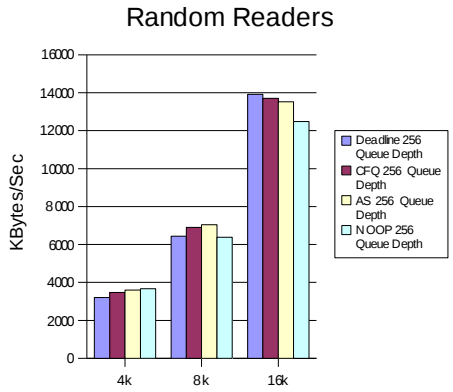


Using Direct I/O no caching.

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

| **Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 256 Queue Depth | 36418 | 40015 | 63248 |
| CFQ 256 Queue Depth | 34696 | 49236 | 60110 |
| AS 256 Queue Depth | 36289 | 40063 | 52079 |
| NOOP 256 Queue Depth | 34822 | 48326 | 63093 |

| **ReWriters 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 256 Queue Depth | 37521 | 55501 | 70520 |
| CFQ 256 Queue Depth | 39436 | 55615 | 64659 |
| AS 256 Queue Depth | 39360 | 43310 | 59220 |
| NOOP 256 Queue Depth | 39925 | 55779 | 71126 |

| **Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 256 Queue Depth | 73688 | 99272 | 108722 |
| CFQ 256 Queue Depth | 74118 | 96838 | 102163 |
| AS 256 Queue Depth | 19013 | 18447 | 23859 |
| NOOP 256 Queue Depth | 74533 | 98386 | 107382 |

| **ReReaders 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 256 Queue Depth | 78080 | 112486 | 118406 |
| CFQ 256 Queue Depth | 78079 | 104157 | 107112 |
| AS 256 Queue Depth | 16349 | 21857 | 28416 |
| NOOP 256 Queue Depth | 77430 | 109681 | 126173 |

| **Random Readers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 256 Queue Depth | 3590 | 6972 | 13429 |
| CFQ 256 Queue Depth | 3802 | 7120 | 13115 |
| AS 256 Queue Depth | 2129 | 4122 | 7898 |
| NOOP 256 Queue Depth | 3693 | 6805 | 13445 |

| **Random Writers 8 Threads** | **4k** | **8k** | **16k** |
|---|---|---|---|
| Deadline 256 Queue Depth | 35063 | 46839 | 55473 |
| CFQ 256 Queue Depth | 35332 | 46238 | 58521 |
| AS 256 Queue Depth | 26434 | 38453 | 53635 |
| NOOP 256 Queue Depth | 33893 | 45876 | 53605 |

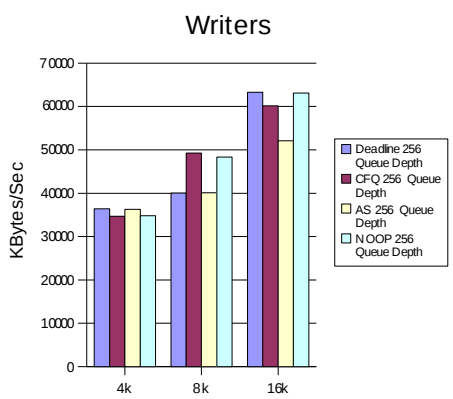


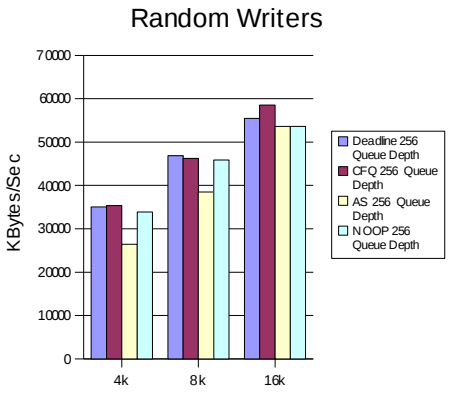


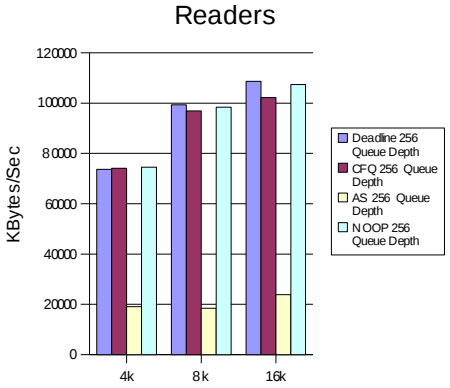


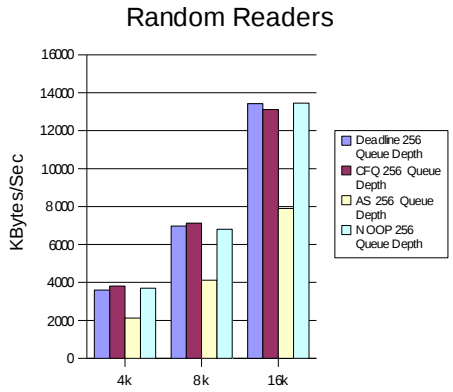


Using Direct I/O no caching.

LVM – refers to EXT3 laid on an LVM volume made up of either 8 or 12 hypers from the EMC

30 and 256 queue depths refer to that setting for the HBA 'emulex driver.

Filer – refers to the 12 disk aggregate from the 920 filer 'where the lun was created then delivered over the SAN

256T is the 12 disk filer SAN volume with OS NAS tuning in place to see if it caused any appreciable impact.

William R. Welty Infrastructure Architect, Managing Partner RedHawk.org November, 2007

| | | |
|---|---|---|
| System | Sun 4200 | |
| OS | RHEL 4 U5 | |
| Memory | 16GB | |
| CPU | 2cpu 2.8GHZ | |
| Test | IOZONE | iozone -+u -I -i0 -i1 -i2 -m -t 8 -O -r 4k -s 512m -Rb /home/bwelty/direcio/ld101-san4ko-cfq30x8.wks |
| | | Directio |
| | | Cpu Stats |
| | | MB/s and IO/s |
| | | 4,8,16K record sizes |
| | | 512MB file size |
| | | 4,6,8,10,12 Threads |
| | Scheduler | CFQ, Deadline, AS, NOOP |
| | Queue Depth | 30, 256 |
| NAS | 12 Spindles | |
| | Netapp 920 Filer | |
| | Jumbo Frames. | |
| SAN EMC – DMX800 | 8 HYPER LVM | LVM on 12 Hypers stripped using 128K stripe size. |
| Emulex 11000 cards. | 8 HYPER Meta | |

General Notes

During tests with NAS and SAN the iowait hit in the 42% range, Load 10-14
LVM performed better than Meta volume at 8way.
Both LVM and Meta 8 way peaked spindle usage to 100%, LVM fully 100%, vs meta just under.

Queue Depth change worked on 12 vs 8 way where we saw no difference.

/etc/sysctl.conf tuning.

```
net.ipv4.ipfrag_low_thresh = 262144
net.ipv4.ipfrag_high_thresh = 393216
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.core.rmem_default = 262144
net.core.wmem_default = 262144
net.ipv4.ip_local_port_range = 1024 65000
fs.file-max=327679
fs.aio-max-nr=3145728
sunrpc.tcp_slot_table_entries = 128
vm.page-cluster = 40
vm.lower_zone_protection = 1024
vm.min_free_kbytes = 200000
vm.swappiness = 0
vm.dirty_background_ratio = 10
vm.dirty_expire_centisecs = 1000
vm.dirty_ratio = 20
vm.dirty_writeback_centisecs = 100
vm.vfs_cache_pressure = 40
```