

Monday, August 10, 2009

PSARC/2009/423 - ZFS logbias property

The PSARC/2009/423 - ZFS logbias property is an interesting addition to the separated ZIL functionality of ZFS. Not every synchronous write need the shortest possible latency, there are many write that are in need of throughput. And then there are synchronous writes that need so short latencies, that you want to remove all requests from a separated log, that doesn't need the help it.

With the logbias property you can control this behaviour. To stick with the example in the PSARC case: Redo Logs need short latencies, so you would use `zfs set logbias=latency bigoracledataset/redologs`. The data files of oracle need throughput, so you would configure `zfs set logbias=throughput bigoracledataset/datafile`. You don't have configure distinct pools (one with sZIL, one without it) to get the optimal behaviour). Furthermore the throughput mode makes another important change: Also data will be written immediately to spread the write load thus making for quicker subsequent transaction group commits to the pool.

Posted by Joerg Moellenkamp in English, Solaris at 13:54

Sounds great, but I cannot find any hint on the web with an estimation when we can expect to be able to use. Do You have any?

Thanks
Kalle

Anonymous on Aug 11 2009, 07:44