

Wednesday, June 4, 2008

Solid State Disks at Sun

We've announced yesterday, that we will deliver solid state disks in our servers in the next fiscal year. Many journalists wrote about the complete substitution of rotating rust disks by solid state disks or drives (strange wording for a technology without disks or any moving parts), but they forgot the most obvious usage of this. One example is: Sun stretches for Flash dance

Rotating rust as it's advantage and solid state has other advantage. So why don't use both. Solaris and ZFS can do exactly this: Separated ZIL and L2ARC were designed to use the advantages of SSD where rotating disks can use some help. You don't have to move all your data to solid state just to get the speed advantages of solid state. Yet another point, where you have to see the the complete package to understand a move made by Sun.

Posted by Joerg Moellenkamp in English, Oracle at 13:50

Indeed it seemed obvious to me when I heard about this SSD news that they would be combined with regular disks in ZFS.

SLOG and L2ARC are two very nice applications. However, I have the impression that ZFS is missing a way to use the SSD as a write cache.

Anonymous on Jun 4 2008, 15:04

Why do you need a write-cache, when you have a seperated ZIL ?

Anonymous on Jun 4 2008, 15:10

In order to write faster?

The seperated ZIL means you can write fast, but the sync can still take a lot of time (you have to actually write to those slow disks). For some applications, before continuing, you want to make sure what you have written has been committed to stable storage. The SSD could be used here, and the data later moved to the slower disks. This is what I mean by a write cache (might not be the best name). It might be better suited to SAMFS rather than an extension to ZFS though...

Anonymous on Jun 5 2008, 13:21

With such an construction you would double the load on the storage connects and systems, as the data has to be moved twice. And by the way: Your description sounds very similar to the concept of the seperated ZIL.

Anonymous on Jun 5 2008, 18:53

I guess I was wondering how ZFS could best optimize a raid0 setup that includes disks of very different speeds.

Anonymous on Jun 6 2008, 15:23