Saturday, September 19. 2009

## Aura

The Exadata V2 announcement contained a hint to a new product: The Aura card. The official name of this card is Sun Flash Accelerator F20 PCIe Card. Finally this product found it's way to the Sun Website. The specs are really nice. The 96 GB variant delivers 100.000 read IOPS and 87.000 write IOPS at 4k. The card contains a super-cap to ensure that there is power to write down the caches to the flashes. 1092 MB/s read throughput, 494 MB/s for writes. In addition to this it offers SAS connectivity for 8 drives, thus you loose nothing when you use this card instead of a normal SAS controller. You just have additional flash disks.

Posted by Joerg Moellenkamp in English, Oracle at 20:03

Looks like an Adaptec-Controller with 2 or 4 OpenFlash Modules and a larger battery. HighRes Photos would be nice...
Anonymous on Sep 20 2009, 17:09

It's something similar ... i have some more information about it ... i will post it, as soon as the docs of this card are available on docs.sun.com
Anonymous on Sep 20 2009, 20:56

Hi,

it`s good to see that Sun is finally offering something like the PCIe flash cards Fusion-Io is offering. I thought about trying out their products, but they didn't deliver the needed Solaris drivers so far. Too bad that it's not possible to use the fast flash device as L2ARC in regular Solaris 10 Update 7 and if you use it as a ZIL you cannot get rid of the dedicated log device yet without destroying the pool
Anonymous on Sep 20 2009, 22:16

I was wondering about that already (couldn't find it on docs.sun.com, but blamed it on my incompetence...).
Anonymous on Sep 20 2009, 23:34

Both complaints are not new, so just one comment to this: The solution is near

But: The problem of the non-removable ZIL is more a cosmetic problem. When the sZIL isn't available or failed, the system just used the pool ... you just have a failed zil in your configuration. ZFS ignores it.
Anonymous on Sep 21 2009, 06:49

I assume, the fact that Oracle mentioned it on the slide deck for Exadata V2 resulted in too many questions. Thus it was easier to put this page online.
Anonymous on Sep 21 2009, 10:31

doesnt this result in loss of data, that was written to sZIL and not committed to the pool?
Anonymous on Sep 21 2009, 10:53

All the data in the sZIL is in the ARC as well. That's the reason, why you just write to sZIL. You read only from it, when you have to recover from a crash without syncing. Thus as long you don't unplug the device and switch off the system before the transaction group hasn't commited to disk (a few seconds, depending on configuration) there is no loss off data.
Anonymous on Sep 21 2009, 10:57

bigger picture..
http://fixunix.com/solaris-rss/565496-welcome-aura-flashfire-hba.html
Anonymous on Oct 2 2009, 22:23