

Friday, March 21. 2008

### **iSCSI for I/O-intensive tasks**

I had an idea to vastly reduce the network load for iSCSI harddisks while waiting for the water to boil for my first tea. There are two new features in development for ZFS: L2ARC and Separated ZIL. L2ARC enables ZFS to use a disc as a second level cache (thus speeding up reads). The Separated ZIL enables ZFS to use a disc as to use a dedicated disc for the ZFS Intent Log (thus speeding up writes). Both features were designed with solid state disks in mind to augment "rotating rust" hard disks. But this features are not limited to this high end task.

Think about the following usecase: A server has in most cases some local boot disks even when you use iSCSI or FC. From a 146 GB disc you need only roundabout 20 GB (2x 7.5 for OS incl. Live Upgrade Partition and 5 GB for swap) or so. Thus you have 126 GB left. Okay, let's assume, that you have two X4500 exporting zvols with iSCSI.

You have another machine in your network. You create an mirrored ZFS pool out of the remote iSCSI volume. But now the both new features come into the game. We augment this iSCSI volume with a local separated ZIL and a local L2ARC. With the first one, you keep the log writes from the network and with the local L2ARC you get an really huge local cache without sacrificing precious memory resources. Many read/write requests stay local to the system and you can vastly reduce the load to the network. Thus you are capable to put even more I/O intensive tasks on this system.

Posted by Joerg Moellenkamp in English, Solaris at 06:34