

Monday, March 12. 2007

Thoughtgame: Leveraging ZFS, iSCSI and the Thumper with SamFS

It's a thing i will try out in the next few days. Can i use ZFS, iSCSI and Thumper to improve the already excellent SamFS HSM filesystem? Can i solve the problem of migration in storage media life cycle management en passant?

For the people new to SamFS: It's a hierarchical storage management system(HSM). The HSM consists out of two storage systems: The cache to speed up access. And the hierarchy. Based on a ruleset, data is stored to certain connected storage devices. An example: Right after writing a file to disk, it's likely that you will access it soon. So you keep it in cache. At the same time you copy the file to different medias. When you don't used a file, and you don't have space for new files, it gets deleted from your cache discs. Whenever you need it, the HSM fetches it from the media you stored it. You can use tapes, magneto opticals, discs or something similar to build up hierarchies.

Obviously the access times differ by a large amount: Data on hard discs are available in a matter of fractions of a seconds, as they are connected to the system and spin already. Data on tapes need more time. You need some seconds to load the tape in a drive and you need some seconds to spin up the tape and you need seconds to find the data by winding to the position. So you would use multiple stages of hierarchy: The data least accessed goes to the slowest storage device. An HSM controls this lifecycle by a set of rules .. well, before i explain SamFS completely, you should read this document. It explains the concepts to some extend.

The interface to all this logic for the user or application: A posix compliant filesystem. You use it like any other filesystem. HSM systems are know for years. But perhaps we can use ZFS and Thumper to spice up the old HSM technology a little bit.

Okay: Storage on a X4500 is really cheap ...Get a 10-pack of thumpers. Round about 450.000\$ list price. Gives you 100 TB (20 TB per Thumper nets with parity. 20 TB with 10 systems gives you 200 TB. Divided by two, as it's sensible to create 2 copies of your data equals 100 TB). You define a large ZFS emulated volume with activated compression and checksumming on your each of your thumpers. Share it via the integrated iSCSI. Connect the Thumper and your SamFS system via 10 GBit/s Ethernet.

Next step is configuring the iSCSI initiator to use the iSCSI targets on your thumper. Now you use this devices as disk archiving devices in your SamFS configuration. As the iSCSI devices looks as ordinary scsi devices, this should be a piece of cake.

What you you get? You get an really large and f... mindblazingly fast hierarchical storage management with storage devices secured by checksums to ensure data validity and online data compression. And data encryption is already a planned feature for ZFS.

And now the real kicker of the combination: When you know, that your drives goes to the end of their projected lifetime (you remember Common wisdom in rotating rust), you buy a Thumper with new discs, define the Thumper emulated medias in SamFS as damaged and SamFS will automatically create new copies on your new thumper and migrates all the data on your SamFS filesystem onto different locations.

Posted by Joerg Moellenkamp in English, Oracle at 22:38

Hi,

Have you tried this yet? The article is from March 07 and says that you would be trying this in a few days.

If you have tried this I would be very interested in your results.

thanks

Anonymous on Jan 30 2008, 00:40