

Friday, January 18. 2008

SAN without a storage system

I had a strange idea some days ago. The new generation of our X- or T-class chassis have bays for 8 disk right now. There must be a good usage for this mass of bays. Think about it. This means you have 320 Disks in your rack. A few seconds later i remembered about a question, somebody asked at the Niagara Presentation Event some years ago in Berlin (must have been 2005): "72 GB per disk, okay ... your OS uses 4 Gig at maximum. How can i use the rest? It would be nice, if other systems can use the storage". I had no answer for it ... okay ... the standard ones ... SMB/NFS.

Well ... there is a better idea. We have an iSCSI Target in Solaris right now and COMSTAR enables Opensolaris to present FC targets. Okay, now take all this free storage, format it ZFS and create 2G sized emulated Volumes out of it and map each of it with the respective target emulation to the LAN or SAN. Whenever you need some storage, you can pick some of the several volumes you have created from all nodes with the exception of the node you work on.

By such a configuration, a rack of servers can easily be it's own SAN without a real storage subsystem. And you can use all the unused space on all the disks, or use own disks for such a construction with 8 or more bays you have more than enough room. With 128 SPARC Threads on 2 HU or 16 x86 threads on 2 HU there should be enough spare cycles on the CPU as well.

Just an idea, but i will think about it the next days.

Posted by Joerg Moellenkamp in English, Oracle at 12:03

Herzlichen Glückwunsch, Sie haben das Google Filesystem erfunden. Nur komplizierter.
Anonymous on Jan 18 2008, 16:04

Hmmm ... wie in der Bosch-Werbung ...

Aber das GoogleFS ist doch shared reader/shared writer? Das kann die obige Lösung nicht.
Anonymous on Jan 18 2008, 17:21

Ob es jeder Admin toll findet, wenn "sein" Rechner als Storage mitbenutzt wird? Darüber hinaus muss man eine zusätzliche FC-Infrastruktur aufbauen... aber iSCSI-LDOMs koennten mir gefallen...
Anonymous on Jan 18 2008, 20:34

Ja, okay ... das mit dem FC stimmt. Aber mit dem iSCSI geht schon heute. Und wenn er mosert, einfach den Admin in einem "Container stecken" ... dann sieht er nichts davon, das die Maschine noch was anderes macht ... ansonsten ist es ja eh seltener, das in einer groesseren Farm jeder rechner von einem Admin gewartet wird.
Anonymous on Jan 19 2008, 10:16

Trotzdem widerspricht die Lösung dem Konsolidierungsgedanken. Stoarges gehören in das Zentrum einer IT-Struktur. Es ist egal wer oder was die Daten verarbeitet - wichtig ist wo sie liegen. Trotzdem erscheint mir die Lösung für bestimmte Anwendungszwecke sicherlich interessant.

MfG
Patrick
Anonymous on Jan 19 2008, 15:50

Hey, storage and servers are consolidated in this example
Anonymous on Jan 19 2008, 18:40