

Re: StorEdge 3310, disk sets, 2 hosts

Source: <http://unix.derkeiler.com/Newsgroups/comp.unix.solaris/2004-09/1995.html>

From: noone (noone_at_noone.org)

Date: 09/25/04

Date: Sat, 25 Sep 2004 18:04:51 +1000

Elias wrote:

> noone wrote:

>>

>> *Forget about the mirroring of the arrays via disksuite for the mean time.*

>>

>> *Assume for the moment that there is just 1 array, the question still*

>> *stands:*

>>

>>

>> *1) One or two controllers on the 3310*

>>

>> *Do I need 1 or 2 controllers (one controller on 3310 per host) ?*

>>

>>

>> *2) Any issues of taking the controller from another 3310 and placing*

>> *it on my existing 3310.*

>>

>>

>> *3) Finally, has any done this before and is this setup (diskset, 2*

>> *hosts, NO clustering and NO HA) supported by Sun ?*

>>

>>

>

>

> *The second controller does not add functionality, it only adds*

> *redundancy. So, there is no need to get a second controller for either*

> *array.*

>

That's what I thought, but they were suggesting that the second controller IS required.

> *As for earlier comments about array based vs host based mirroring, in*

> *your current case host based mirroring is a better way to go. Any time*

> *you have things in one physical enclosure you have risk. By spreading*

> *the mirror across multiple arrays you reduce risk. There are plenty of*

> *instances where a single array (this is vendor neutral) that didn't have*

> *a single point of failure goes down. This is usually caused by bugs in*

- > *the firmware.*
- >
- > *The type of DB failover you want to do is probably the most risky in*
- > *terms of having a successful failover. One of the value add features of*
- > *cluster software is the cluster file system. If you have locks or*
- > *reservations that don't get properly released the second server may*
- > *refuse to import the volume.*

But solstice disksuite allows you to reserve the diskset forcibly in the event that host A, which previously had the reservation, have a hardware failure and can no longer boot for whatever reason ... so that host B can reverse the diskset forcibly:

<http://docs.sun.com/db/doc/806-3205/6jccc20rb?a=view>

Doesn't that solve the problem ?

- >
- > *There are other options though. Most of the options that have more*
- > *reliable failovers involve replicating the data from server to server*
- > *instead of through a mirror. This would mean that you put one 3310 on*
- > *your 280 and one on the 220, with no array doing any multi-hosting. A*
- > *list of things that are known to work for this are:*
- > *– Sun Network Data Replicator*
- > *– Oracle DataGuard*
- > *– Veritas Volume Replicator*
- > *– Scripted Log Shipping*
- >
- > *You should be seeing a theme here. Replicate the updates over the*
- > *network to the standby server while it's in recovery mode and have it*
- > *read in all the logs, or replicate the volume over the network. By*
- > *having each server use independent storage you don't have to worry about*
- > *importing and exporting volumes. You do have to make some decisions on*
- > *synchronous or asynchronous replication. i.e. does the commit happen*
- > *when both servers write the data or only when the primary server writes*
- > *it. Your risk is transactions that haven't been sent and the ones that*
- > *are in flight. It's all a matter of how much risk you want to take.*
- >

Yes I know .. we investigated the and evaluated Sybase Replication Server (reading transactions logs from one database to applying these transaction logs to another database) to implement a warm-standby. It works well ... but the licensing cost is very prohibitive for a small company.

- >
- > *Elias*