

Re: Resilient VMS server config

Source: <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2004-08/0490.html>

From: Michael Austin (maustin_at_firstdbasource.com)

Date: 08/10/04

Date: Tue, 10 Aug 2004 16:35:37 GMT

Chris Hale wrote:

- > *Hi all,*
- >
- > *I could do with some advice regards a highly resilient server/cluster.*
- >
- > *We have an application which is not processor intensive, and it does*
- > *not require huge amounts of disk space. In fact we have had it running*
- > *on an Alpha1000 with 2.1G disks without any problems.*
- >
- > *However we need replace old hardware and to ensure that it is highly*
- > *available.*
- >
- > *At the moment its running on a cluster of 4 Alpha's with shared SCSI*
- > *disks between each pair of nodes, and host based volume shadowing.*
- >
- > *We are now proposing a pair of DS25's with a SAN providing the*
- > *storage, and possibly host based volume shadowing across logical*
- > *volumes provided by the SAN. The logical volumes themselves can then*
- > *be members of a raid array or mirror. The proposed cluster*
- > *interconnect would be fibre, with an Ethernet backup.*
- >
- > *As well as any general advice or comments, could anyone address the*
- > *following:*
- >
- > *1. Should we bother with host based volume shadowing, when the SAN*
- > *provides raid or mirrors?*

Each has it's pro's and con's If they are on the same Storage Controller, then the Mirror/RAID will remove the overhead of doing this on the CPU

- > *2. Does a SAN help with image backups whilst maintaining the*
- > *availability of the disk? For example can we take a 'snapshot' of a*
- > *volume with very little loss of availability of the volume?*

There is NO loss of availability** Do online backups of your database to a backup LUN and then back that up to tape – yes it is a 2 step process, but can be done with NO interruption of service depending on the database engine you are using.

comp.os.vms: Re: Resilient VMS server config

- > 3. *The application does not require a cluster, but is able to failover*
- > *to another node. Is there a good reason to go for three nodes instead*
- > *of two?*

If you have 2 DS25's for production and failover and let's say a DS15 for Maintenance stuff – like doing the actual backup to tape...

- > 4. *Can anyone comment on the reliability of SAN solutions?*

If built correctly— EXTREMELY reliable!!! see recommendations below

- >
- > *Any help of advice would be appreciated.*
- >
- > *Thanks.*
- >
- > *Chris.*

Questions:

What database engine/version?

Which storage controller will be employed on the SAN? EVA? HSG80? XP?

When considering your backups, you want to purchase enough storage that you can keep 2 full copies plus incrementals ONLINE. If you have to restore from tape in an environment that can only stand downtime 3 times a year, you don't want to rely on tape as your primary medium. With disk be "cheap" compared to the cost of downtime, this is a no-brainer. You need to do a DR test. Hand your sys admin (you??) a stack of tapes and go restore them somewhere. How long did it take? can you afford that kind of downtime time? I doubt it.

I worked in a shop that had these sorts of challenges. We used the technology to do online backups to disk then a separate process to move them from disk to tape as a last resort only type of recovery. (200+ Alpha servers in 2 node cluster configurations and a "maintenance cluster" for doing the actual backups. 36+ HSG80's (BA and EMA) 6 EVA's and now at least one or two XP1024's all fully loaded with (18, 36, 72, 146 drives – over 4500 spindles and more than 300TB of raw storage.)

If you are in need of a consultant to help implement this, let me know... It is cool stuff, but unless you do it right, it won't be very efficient.

Also, building your SAN right the first time will help keep you out of trouble in the future.

Recommendations:

each system has a minimum of 2 HBA's #1 goes to SAN-A and #2 goes to SAN-B
you should have 2 redundant SANS – not the "star" interconnect where every switch is ISL'ed to every other switch – this makes SAN upgrades nearly impossible without serious downtime and needlessly wastes your port count. The

comp.os.vms: Re: Resilient VMS server config

newer switches with hot code uploads make it better, but IMPO dual-redundant SANS are much more resilient.

Use Brocade – IMPO they are still the leader and management is much easier than the others.

Use EVA with enough spindles to have 2 copies of your backups online.

Use Snapshots (Business Copy) to move data from your backup drives to tape.

This was for free.

--

Michael Austin.

Consultant - Available.

Donations welcomed. <http://www.firstdbasource.com/donations.html>

:)