

## Re: Clustering Newbie – SAN Advice

---

*Source:*

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.clustering/2007-11/msg00063.html>

---

- *From:* "Geoff N. Hiten" <[SQLCraftsman@xxxxxxxxxx](mailto:SQLCraftsman@xxxxxxxxxx)>
  - *Date:* Tue, 20 Nov 2007 22:30:18 -0500
- 

I have not used that particular unit, but it does meet my basic criteria of having the controllers and cache on-board the array. I would see if Dell can give you access to one to do some I/O testing.

You have hit upon why I am so frustrated with the "Active/Active/Passive/whatever" nomenclature. Simply put, you cannot have a single database file accessed and managed by two host computers. You can have two different databases on different disk sets that normally live on different nodes, but can failover to a single node if necessary.

Of course, any single node must be able to handle the entire load under the worst conditions. Because that is when the system will fail. :)

—  
Geoff N. Hiten  
Senior SQL Infrastructure Consultant  
Microsoft SQL Server MVP

"Aaron" <[Aaron@xxxxxxxxxxxxxxxxxxxxxxxxxxxx](mailto:Aaron@xxxxxxxxxxxxxxxxxxxxxxxxxxxx)> wrote in message  
[news:5AB08E84-4804-4308-8EFC-959811EDC752@xxxxxxxxxxxxxxxxxxxx](mailto:news:5AB08E84-4804-4308-8EFC-959811EDC752@xxxxxxxxxxxxxxxxxxxx)

I have a similar situation to Craig. I was looking at using the Dell MD 3000 for storage. I know that it is SCSI, but cost is an issue at this point. Any thoughts on that box. Additionally, (this may be a stupid question so forgive me), but can we have one server with an active instance of one database and a passive instance of another. This way my two node cluster would be using both servers actively and spread the load. Does this make sense.

Aaron

"Geoff N. Hiten" wrote:

What you will need is one or more fibre channel adapters (called Host Bus Adapters or HBAs) in each host computer. These can connect directly to the SAN/Smart array or through a fibre channel switch. The configuration is different for each. Fibre channel is really just a wrapper protocol around

Re: Clustering Newbie – SAN Advice

the standard SCSI block access protocol, along with a fiber and an electrical (copper) spec.

A SAN generally has gigabtes of cache and uses large internal block sizes and aggressive caching techniques to optimize IO. A SAN can also support many host computers where a smart array is intended to support only a few. SANs are much more scalable. A Smart array does some abstraction of the internal disks, but not nearly to the degree a SAN will.

The SAN or Smart array will dictate what internal connection the disks have. Some have fibre channel (sometimes over copper wire) all the way to the disks. Some have SCSI or SAS attachments to the actual disks.

When dealing with a SAN, the topic of LUN alignment is critical. This is where the OS allocation blocks line up with the SAN internal stripe boundaries. If they are not lined up, the SAN must execute excessive internal IO operations for each physical disk request. This can lose you up to 40% of the IO capacity of the unit.

Whatever you buy, run some IO tests so you will have a realistic expectation and a baseline comparison when the unit is in production.

--  
Geoff N. Hiten  
Senior SQL Infrastructure Consultant  
Microsoft SQL Server MVP

"craig\_amtdatatechnologies@xxxxxxxxxxxxxxxxx"  
<craigamtdatatechnologiesdiscussionsmi@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>  
wrote in  
message  
news:C485D29A-13F2-4F88-BF41-25C3F4B4B8AD@xxxxxxxxxxxxxxxxx  
> Hi,  
>  
> Thanks for reply, certainly helpful :)  
>  
> I note the need for clarity in nomenclature, correction welcome.  
>  
> If your (or anyone else), were to put names (not brands) to the bits, > what  
> nomenclature should I be using to described the various components and  
> bits  
> connecting them together?  
>  
> Basically, my knowledge of storage systems is pretty weak (as you wil  
> lhave  
> guessed) ... so I need a connector out the back of my sql boxes (a bit > of  
> fibre ... 1 or 2 bits each?), something to plug the fibre cables into > (an

## Re: Clustering Newbie – SAN Advice

> array), some disks to stick in the array. If fibre connector is used, > do  
> these disks need to be anything special?  
>  
> Sorry for noddy style questions ... !!  
>  
>  
> Thanks again  
>  
>  
>  
> "Geoff N. Hiten" wrote:  
>  
>> Lets see.  
>>  
>> x64 is definitely the best bang for the buck in the range you are >>  
>> looking  
>> at. More memory will help. Put in as much as the box can hold.  
>>  
>> Obligatory naming rant: Active/Passive really doesn't describe how  
>> things  
>> work. You are building a two-node, single-instance cluster. That  
>> nomenclature is much more accurate and it scales. (Try describing a  
>> cluster  
>> with 6 nodes and 12 instances using the words "active" and "passive".)  
>>  
>> One very important consideration when building a cluster is to use >>  
>> some  
>> type  
>> of smart array or a true SAN. The key is that the disk cache must >>  
>> reside  
>> in  
>> the array, not on an onboard controller. Controller write cache is  
>> always  
>> disabled in a clustered setting to maintain data consistency.  
>>  
>> If performance is key, then RAID 1+0 is the configuration you want.  
>> Don't  
>> worry about swap files, a properly configured dedicated SQL server >>  
>> should  
>> not have any significant paging activity. Do split tempdb onto its own  
>> fast  
>> disk set.  
>>  
>> As for connections, I strongly prefer fibre channel where possible.  
>> First,  
>> it is a full-duplex data path so you can read and write at the same >>  
>> time.  
>> SCSI is a bus so a large checkpoint can hammer read performance. >>  
>> Second,  
>> you can only have two nodes in the cluster with SCSI. I am also very  
>> reluctant to unplug a SCSI cable on a downed node should I need to >>

Re: Clustering Newbie – SAN Advice

repair  
>> something. Not so much with a fibre channel connector.  
>>  
>> You can always bolt mirroring on top of clustering for some geographic  
>> separation. Choosing one does not necessarily eliminate the other.  
>>  
>> -- >> Geoff N. Hiten  
>> Senior SQL Infrastructure Consultant  
>> Microsoft SQL Server MVP  
>>  
>>  
>>  
>>  
>> "craig\_amtdatatechnologies@xxxxxxxxxxxxxxxxx"  
>>  
<craigamtdatatechnologiesdiscussionsmi@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>  
>> wrote  
>> in  
>> message  
news:130DBB73-7072-4E58-A36D-DFCD35CC6057@xxxxxxxxxxxxxxxxxxxx  
>> > Hi,  
>> >  
>> > We are looking into implementing an Active:Passive SQL 2005 Cluster  
>> > for  
>> > a  
>> > reasonably large and important system (< 1Tb)  
>> >  
>> > I'm not looking for advice re: particular brand etc., but on advice >> >  
>> on  
>> > type  
>> > / type of connection to SAN that may be appropriate. The SQL 2005  
>> > instance(s)  
>> > will be x64.  
>> >  
>> > Please, any wannabe MVPs do not provide sarcastic/useless answers >>  
>> just  
>> > to  
>> > up  
>> > your profile. To pre-empt some of these:-  
>> >  
>> > Yes I have indexes on current database  
>> > Yes the queries have been tuned  
>> > Yes, I have detailed stats on usage, performance, but you do not  
>> > need  
>> > these to answer the general thrust of the question  
>> > etc.  
>> >  
>> > Now, I went through the Dell SQL configurator to see what they came  
>> > up  
>> > with  
>> > .... they suggested £50k would be a reasonable price ... maybe it >> > is,

## Re: Clustering Newbie – SAN Advice

>>> maybe  
>>> it isn't. This had all manner fancy fibre connection things, disks >>> at  
>>> £1000  
>>> each, other buzzers and bells. I wonder if I really need all this?  
>>>  
>>> Basic setup is 2 \* servers (Windows 2003 R2, standard, x64), spec'd  
>>> (each):-  
>>>  
>>> 1 \* Quad-core CPU  
>>> 2 \* 15,000rpm disks (RAID 1) for O/S and SQL program files etc. >>>  
(and  
>>> swap files)  
>>> 8Gb RAM  
>>> Network cards, etc.  
>>>  
>>> I would like, an external disk array to hold database drives (pretty  
>>> essential really for a cluster!), of roughly the following format >>>  
(each  
>>> physical disk probably 146Gb capacity):-  
>>>  
>>> 2 \* 15,000 rpm RAID 1 (tempdb)  
>>> 2 \* 15,000 rpm RAID 1 (transaction logs)  
>>> At least 6 (probably more) \* 15,000 rpm Stripped and Mirrored for >>>  
> main  
>>> database files  
>>>  
>>> What sort of storage kit/spec. (fibre? not fibre? etc. ???) should I >>>  
be  
>>> looking for to support the cluster? ie. what will do, not what do  
>>> HP/Dell  
>>> etc. wish to sell me.  
>>>  
>>> Yes the system wil grow over time.  
>>>  
>>> One last thing ...  
>>>  
>>> There may be alternative solutions to an Active:Passive cluster, but  
>>> some  
>>> big name players (financial/multinational etc.) are using the >>>  
system,  
>>> and  
>>> we  
>>> need to do some serious bomb-proofing.  
>>>  
>>> I would prefer to have a single high-performance cluster, to reduce  
>>> number  
>>> of required SQL licences (sorry Bill!), and also avoid having to  
>>> maintain  
>>> too  
>>> many live copies of the database (yes Mirroring is a cool >>>  
technology,

