

Re: Add new cluster and use existing LUN's?

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.clustering/2005-10/msg00120.html>

- *From:* "skyline" <skyline@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Fri, 21 Oct 2005 11:53:04 -0700
-

Anthony,

This is exactly what I was looking for. We have dual HBA's using the Storport driver and Powerpath on all of our SAN host servers so we are trying to make the best of our fibre connections.

In the end I think that I may play it cautious and create a new RAID group, move our SQL data over temporarily, and destroy/recreate the 4 disk RAID 10 into 2 RAID 10 groups with 2 drives a piece and then add 2 drives to our remaining open slots and create a RAID 10 group out of them for the second virtual servers t-logs.

My final decision will probably come based upon our drive usage analysis, but as you said I am going to be careful :)

"Anthony Thomas" wrote:

- > I am speaking from experience here: spindle contention can big the BIGGEST
- > bottleneck to a SQL installation. Recommendation, be careful. I think you
- > are on the right track though. A few points to help guide you.
- >
- > 1. A physical spindle (real disk) has limited IOPS and bandwidth capacity,
- > anywhere between 75 and 150 IOPS (I/Os per second) theoretical maximum
- > sequential reads. Random reads and any type of writes decrease from there.
- > To overcome this limitation, Meta-volumes are constructed out of the RAID
- > sets/Hyper-volumes.
- >
- > 2. Whether or not these RAID sets from which you construct the Hyper and
- > Meta-volumes are dedicated to a single host, single activity (OS, swap file,
- > TLogs, Data Files, or TempDB) or shared among many hosts with similar or
- > varied activity (DBMSes, Messaging Server, File Server, Web Servers,
- > Application Servers, or Middleware components), the physical spindles retain
- > the same theoretical IOPS and bandwidth capacities. Make sure what you
- > construct can handle the anticipated loads: both in IOPS and total
- > throughput.
- >
- > 3. The various activities of all hosts as well as those of a particular host
- > need to be isolated: O/S, page file, Quorum, MS DTC, TLogs, Data Files,

Re: Add new cluster and use existing LUN's?

- > TempDB, Exchange Server jet files, Web Server pages, etc. You could
- > however, create a large RAID set that could be shared across multiple hosts
- > from which you construct several Meta-volumes that are dedicated to, say,
- > page file space. The point is ISOLATE the TYPE of activity, then share with
- > as many hosts as the IOPS and bandwidth will support.
- >
- > 4. Some of the physical spindle limitations can be addressed through the SAN
- > cache, which can be configured fractionally as dedicated to Read and or
- > Write operations. This piece requires a detailed understanding of the SAN
- > configuration and the anticipated host activity, whether or not that
- > activity has been made homogeneous throughout the SAN or whether or not the
- > various activities have been segmented to dedicate RAID sets.
- >
- > 5. Consider the number of SAN fibre controllers and available ports, each
- > one has a theoretical limit that can be overloaded and the number limits the
- > total activity that can be assigned to a given SAN cabinet installation.
- >
- > 6. Network switches, each port and switch have theoretical limitations and
- > which will be serving up all connected hosts. The total usage must be
- > aggregated to properly configure.
- >
- > 7. HBAs on each host have limitations. Even if you go through all of the
- > trouble of isolating activity on the SAN, carving up and presented isolated
- > LUNs to each host dedicated to the separate activities, if you try to run
- > all of that through one or two HBAs, you may still cause the same sort of
- > contention. Seriously consider running Windows 2003 and using STORport
- > drivers over SCSIport minidrivers. You will also need to look into a MPIO
- > solution like Veritas Volume Manager or EMC Power Path. This will allow you
- > to leverage the bandwidth across all HBAs and allow you to get by with fewer
- > overall, not to mention this will provide you another
- > Single-Point-of-Failure isolation.
- >
- > 8. Finally, I've already said it, but it so important that I'm going to say
- > it again. Many engineers only aggregate and examine total bandwidth real
- > and/or anticipated consumption, in the components, networks, or storage
- > solutions. That is, size capacity and throughput. How much data do I need
- > to store and how much do I need to transfer per second. However, for a
- > heavily used OLTP system, size is less important. Number of concurrent
- > transactions, and thus IOPS, will be the limiting factor. I don't know how
- > many times I have asked my sub-system engineers to look into either the
- > network or storage components just to have them come back and tell me I am
- > only using 10% of the bandwidth or total storage capacity. To which I
- > respond, I am running over 1,000 TPS and up to 1,000,000 IOPS. The typical
- > response is "Huh?". Which is never a good response to hear in a crisis
- > situation.
- >
- > Hope this helps.
- >
- > Sincerely,
- >
- >

Re: Add new cluster and use existing LUN's?

Re: Add new cluster and use existing LUN's?

> Anthony Thomas
>
>
> --
>
> "skyline" <skyline@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
> <news:9993F283-E654-4F30-8316-E491670CCF35@xxxxxxxxxxxxxxxxxxxx>
>> I don't have access to the tool you mentioned but will research it
> (Dell/EMC
>> is correct by the way). I still have the fear of affecting performance on
>> both virtual servers by forcing them to share the same 4 disks because
> RAID
>> 10 is not even bulletproof enough to gurantee a good speed if we are high
> on
>> the disk I/O side. I didn't want to semi-blindly make a decision to
> purchase
>> another cabinet and additional drives...hence why I came out here looking
> to
>> see if a few MVP's reccomended against it. I figured real world
> experience
>> may help me in the decision.
>>
>> Like I told Rodney, I know testing is in store for us—a smoking gun for
> the
>> decision. All I have now are my own opinions and reservations and that
> isn't
>> enough to guide my decision to spend or not spend money.
>>
>> Thanks for the advice.
>>
>> "Geoff N. Hiten" wrote:
>>
>>> Even a SAN has a maximum I/O capacity. Whether you are I/O bound on the
>>> disks in question will determine whether using the remaining space will
>>> impact your current system performance. Your SAN vendor (EMC from the
> hints
>>> you dropped) should have tools (likely for an additional fee) that can
>>> measure IO at the LUN/RAID set/disk level. You will definitely need
> that
>>> capability in order to measure the impact of your proposed change. A
> quick
>>> test would be to run a simulated SQL I/O load on the newly created LUNs
>>> while monitoring I/O performance at the SQL, OS, and SAN levels. Do the
>>> same without the load and see how much you can get away with.
>>>
>>> Personally, I prefer to use LUNS that map one-to-one to RAID Sets,
> mostly to
>>> avoid this problem. Since that isn't a real option now, you may want to
>>> consider purchasing an additional drive tray. It isn't the cheapest
>>> solution, but it may be the most effective.
>>>

Re: Add new cluster and use existing LUN's?

Re: Add new cluster and use existing LUN's?

>>> BTW, you can share a LUN with more than one server. That is how a
> cluster
>>> is built. Access Logix whines about it but you can force it to do what
> you
>>> need. .
>>>
>>> ---
>>> Geoff N. Hiten
>>> Senior Database Administrator
>>> Microsoft SQL Server MVP
>>>
>>>
>>> "skyline" <skyline@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
>>> news:DA48AC82-6A9F-4738-BDB6-08D10D16EAB4@xxxxxxxxxxxxxxxxxxxx
>>>> The SAN is indeed capable. Creating a LUN out of the 225 GB of
>>>> unformatted
>>>> data and assigning it to the new SQL server storage group is quick and
>>>> easy.
>>>> The first 5 drives of our SAN share the system binaries, Snapview
> reserved
>>>> LUN's, and 3 additional LUN's all set as RAID 5 (making a total of 509
> GB
>>>> divided between several different LUN's.) It is possible to share the
>>>> same
>>>> drives with more than one server, however you can't share one LUN with
>>>> more
>>>> than 1 server.
>>>>
>>>> However we would have 1 virtual server accessing the original LUN of
> 175
>>>> GB
>>>> and the other virtual server accessing the newly created LUN of 225
> GB—my
>>>> only hang up being the fact that the 2 LUN's exist on the same set of
> 4
>>>> drives and if you have 2 completely independent SQL servers accessing
> the
>>>> same drives is it going to put a hurt on performance. Again, t-logs
> stay
>>>> seperate (IE the 2 virtual servers would not be sharing the t-log
> drives,
>>>> they would each have their seperate set of RAID 10 drives) so that is
> my
>>>> only
>>>> saving grace thinking this implementation may be possible.
>>>>
>>>> I am here because I need a second opinion (and 3rd or 4ths if
> possible)
>>>> from
>>>> some higher SQL guru's. I personally am the Net Admin so my SQL
>>>> experience

Re: Add new cluster and use existing LUN's?

Re: Add new cluster and use existing LUN's?

>>>> is not broad however I understand the concept of the SAN well and know
>>>> what I
>>>> am looking to do with the SQL servers.
>>>>
>>>> Thanks!
>>>>
>>>> "Rodney R. Fournier [MVP]" wrote:
>>>>
>>>>> It depends on the SAN. You may not be allowed to share the drives
> with
>>>>> different servers.
>>>>>
>>>>> Cheers,
>>>>>
>>>>> Rod
>>>>>
>>>>> MVP – Windows Server – Clustering
>>>>> <http://www.nw-america.com> – Clustering Website
>>>>> <http://msmvps.com/clustering> – Blog
>>>>> <http://www.clusterhelp.com> – Cluster Training
>>>>>
>>>>> "skyline" <skyline@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
>>>>> <news:A4415101-D95B-4871-A666-4D4A2497275F@xxxxxxxxxxxxxxxxxxxx>
>>>>> >I realize that testing is the ideal way to come to the complete
>>>>> >conclusion
>>>>> >on
>>>>> > our pathway but my first and foremost question really is has anyone
>>>>> > successfully used a set of drives and have 2 luns for 2 different
>>>>> > virtual
>>>>> > servers on that same set of drives?
>>>>> >
>>>>> > I'd like to know if I am doing something "average" or if the
> general
>>>>> > consensus is against such an idea. That helps me to determine how
> much
>>>>> > of
>>>>> > a
>>>>> > possibility this really becomes and how much money I will be
> spending.
>>>>> >
>>>>> > Thanks for your input thusfar.
>>>>> >
>>>>> > "Rodney R. Fournier [MVP]" wrote:
>>>>> >
>>>>>> > You need to stress test, look at the results, and then test some
> more.
>>>>>> > Testing will tell you what the correct answer is for your
> environment.
>>>>>> >
>>>>>> > Cheers,
>>>>>> >

Re: Add new cluster and use existing LUN's?

Re: Add new cluster and use existing LUN's?

>>>>> Rod
>>>>>
>>>>> MVP – Windows Server – Clustering
>>>>> <http://www.nw-america.com> – Clustering Website
>>>>> <http://msmvps.com/clustering> – Blog
>>>>> <http://www.clusterhelp.com> – Cluster Training
>>>>>
>>>>> "skyline" <skyline@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
>>>>> news:18DB25CC-37FC-4FEF-AF56-64696BE9D0B8@xxxxxxxxxxxxxxxxxxxx
>>>>>> Hi All,
>>>>>>
>>>>>> We have a Win2003 Enterprise/SQL 2000 "active/passive" cluster,
> and
>>>>>>> we
>>>>>>> want
>>>>>>> to split a few larger DB's off to a seperate and new
>>>>>>> "active/passive"
>>>>>>> cluster. The reason for the split is to allow this new cluster
> to
>>>>>>> control
>>>>>>> the one application and the other cluster to control the
> existing
>>>>>>> app.
>>>>>>> Currently we have:
>>>>>>>
>>>>>>> SQL LUN occupying 4 Raid 10 disks with 175 GB assigned to this
> LUN.
>>>>>>> Doing
>>>>>>> the math leaves us with appx. 225 GB left on these 4 discs that
> is
>>>>>>> unassigned and unformatted.
>>>>>>>
>>>>>>> T-log/Quorum LUN occupying 2 Raid 10 disks with all of the space
>>>>>>> assigned
>>>>>>> to
>>>>>>> this LUN (appx 134 GB).
>>>>>>>
>>>>>>> We are most likely not going to utilize the entire 175 GB for
> the
>>>>>>> data
>>>>>>> collection side let alone the additional 225 GB. At what impact
> to
>>>>>>> the
>>>>>>> disk
>>>>>>> I/O are we going to suffer if we create the new a/p cluster
> using a
>>>>>>> newly
>>>>>>> assigned LUN that occupies the previously unformatted 225 GB on
> the
>>>>>>> SQL
>>>>>>> data

Re: Add new cluster and use existing LUN's?

Re: Add new cluster and use existing LUN's?

>>>>>>> disks? So essentially we would have 2 separate virtual servers
>>>>>>> operating
>>>>>>> off
>>>>>>> of the same spindles, however 2 different LUN's. My gut feeling
>>>>>>> says
>>>>>>> that
>>>>>>> we
>>>>>>> would take some type of I/O performance hit since 2 separate
> high
>>>>>>> I/O
>>>>>>> applications are accessing the same spindles. We would however
> have
>>>>>>> 2
>>>>>>> new
>>>>>>> RAID 10 disks for the t-logs and quorum of the new cluster and
> since
>>>>>>> they
>>>>>>> are
>>>>>>> the highest I/O portion of SQL and would not be sharing this is
> the
>>>>>>> only
>>>>>>> reason I think the idea may be feasible.
>>>>>>>
>>>>>>> The reasons for trying to use the unformatted space on those 4
.

• **References:**

- ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: Rodney R. Fournier [MVP]
 - ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: skyline
 - ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: Rodney R. Fournier [MVP]
 - ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: skyline
 - ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: Geoff N. Hiten
 - ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: skyline
 - ◆ **Re: Add new cluster and use existing LUN's?**
 ◇ From: Anthony Thomas
- Prev by Date: **Re: Add new cluster and use existing LUN's?**
 - Next by Date: **SQL 2000 test cluster: SP4 not applying**
 - Previous by thread: **Re: Add new cluster and use existing LUN's?**
 - Next by thread: **Re: Account for SQL Server services**
 - Index(es):
 - ◆ **Date**
 - ◆ **Thread**