

## Re: Multi-Channel Raid VS SAN Storage

---

*Source:*

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.setup/2008-07/msg00000.html>

---

- *From:* Eli Silverman <[EliSilverman@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxx](mailto:EliSilverman@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxx)>
  - *Date:* Mon, 30 Jun 2008 16:14:01 -0700
- 

Thanks for the quick feedback.

The idea of keeping the database in memory is appealing, but can that be implemented in a clustered environment?

Also how would one implement that any way.

We actually are running RAID 1+0 and our application is definitely more read than write intensive.

We get tolerable performance at best. but the second someone tries to run a report or copy a large datafile performance slows to a crawl. A page that typically loads in 1-2 seconds can take 20-30. and the report that takes 10-15 seconds to run when nobody is on can take 2-3 minutes.

Our report SPs are vry tempDB and memory intensive.

A typical report may require 15-20 tables

We have found that in order for them to run in a timely fashion we have sometimes needed to process the data in several separate stems creating 3 or 4 temp table to hold the data and then join all of the temp tables together at the end.

Maybe we should be using @ memory tables for processing the reports instead of # temp tables but for the most part I seemd to get better performance out of the # temp tables.

Probably because our SQL server only has 3 GB of RAM.

There are a number of inherited inefficiencies from the original design that we are working on eliminating so I think the IO is much more intensive than it needs to be.

Additionally there is a lot of IO from the fileshare.

Also we have 45 people that access the database continually in house and roughly 3000 people that have access to one of the 5 client web sites that access this particular database. We also have 3 other databases that reside on the same server but they are smaller and see much less activity. Our sites use pooled connections but we typically have about 60 connections to the SQL server at any given time.

I think my best choice at the moment is to try and convince them that SQL needs it's own SAN device, or at least a dedicated IO channel on the SAN. And to mak ecertain that we can get as much memory as possible.

Re: Multi-Channel Raid VS SAN Storage

"Andrew J. Kelly" wrote:

Eli,

Explaining SAN storage can take up a whole week but some general comments from what you have said so far. First off a 1.5GB database is extremely small for SQL Server these days. You should be able to fit the entire database in memory so most of the IO should be logical and not physical anyway. The exception would be the transaction log file. But again with a db so small you probably don't have much I/O to begin with. If you are currently using a single Raid 1 array and share that with Exchange and a file share and get OK performance now you should only see improvement with a SAN. Even though you will most likely still share an array on the SAN with other apps you will most likely have many more spindles in the array than you have now with a Raid 1. The SAN also most likely has more cache to buffer the writes which will also help. The key in your case is to ensure you have enough memory to keep the db in cache and you should be pretty set.

--

Andrew J. Kelly SQL MVP  
Solid Quality Mentors

"Eli Silverman" <EliSilverman@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message [news:4F01EC36-1F3A-4B73-8004-313E186878B3@xxxxxxxxxxxxxxxxxxxx](mailto:news:4F01EC36-1F3A-4B73-8004-313E186878B3@xxxxxxxxxxxxxxxxxxxx)

I am hoping someone can give me some advice about the performance difference between Multi-Channel Raid vs SAN storage.

My basic question has to do with Performance and optimization. Assuming I have the same spec drives (size, speed, throughput, etc...) is SQL better off with a Multi-channel raid controller or a SAN storage array?  
Are there such things as multi-channel SAN storage Arrays?  
Can anybody direct me to some good white papers about SAN storage and SQL?

We currently run a two node Clustered SQL 2000 environment. Unfortunately the nodes are attached to a common 1 channel raid device so everything runs off the 1 Raid array.

We are getting ready to upgrade our hardware, however our parent company is insisting that we implement a SAN for all of our data storage and I am wondering if that is the best solution for our SQL server. We will also be upgrading to SQL 2005 at the same time.

## Re: Multi-Channel Raid VS SAN Storage

I know we could benefit from moving the logs, tempdb and possibly some of our indexes over to separate IO channels, but I am uncertain as to how a SAN environment affects these concepts.

I don't think the database is all that big, roughly 1.5GB but there are a couple of sizeable tables. one audit table has over 5 million records and our expense tables have about 300K and will probably grow by at least 25-30K records per year.

The RAID array has been divided into multiple volumes. One dedicated to SQL, another to Exchange and a third to general fileshare. But they are still on a single SCSI channel.

My concern is that when we migrate to our new hardware they are going to try and implement the same architecture and drive throughput is going to continue to be stolen from SQL by the other servers and applications. I want to make certain tha I ask all the right questions and make all the right points to avoid having SQL crippled by our new infrastructure.

Thanks in advance for any advice.

Eli Silverman