

Re: RAID – config Question

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.server/2005-02/2336.html>

From: Andrew J. Kelly (sqlmvpnoospam_at_shadhawk.com)

Date: 02/15/05

Date: Tue, 15 Feb 2005 00:25:40 -0500

SQL Server does not need read cache on the controller since it does it's own read ahead caching anyway and will usually do a better job at predicting what it will need. You are doing massive writes and your disks probably can't handle the load by them selves so going 100% write cache will help a ton. Since you are doing so much tempdb and log activity you want to make sure the tempdb is not on the same disk as the logs.

--

Andrew J. Kelly SQL MVP

"pdxJaxon" <GregoryAJackson@Hotmail.com> wrote in message
news:udwH1ftEFHA.960@TK2MSFTNGP09.phx.gbl...

> actually I misspoke.

>

> the entire job is moving a more than 10s of millions of records (~60GB).

>

> I have individual tables with 10s of millions of records.

>

> Not the biggest db I've ever played with, but it's non-trivial.

>

> By Volume I mean an array.

>

> I have 2 arrays with 3 disks each Both RAID 0.

>

> On array 1 I have C: (OS) and D: (Logs and TempDB)

>

> On Array 2, I have E: (Data and Indexes)

>

>

> You think making the Cache 100% writes will help ?

>

> I am reading from DB A and Writing to DB B

>

> If so, that is an easy change.

>

>

>

> GAJ

>

>

>

> "Andrew J. Kelly" <sqlmvpnoospam@shadhawk.com> wrote in message

> news:OuYeHJtEFHA.1348@TK2MSFTNGP14.phx.gbl...

>> By volume I hope you mean array and not simply a logical device. If

>> fault tolerance is not an issue then why not have one disk for the OS and

microsoft.public.sqlserver.server: Re: RAID – config Question

```
>> the Log files (both tempdb and the user dbs). Then either use the other 5
>> for the data files or take one or two for tempdb and the others for the
>> user data files. You absolutely need to separate the logs from the data
>> files. If you have that much tempdb you may want to split that data file
>> out as well but only testing will tell for sure. In either case change
>> the caching on the disk controller to be 100% write back and you should
>> see a big improvement as well. 8 hours to process only a few 10's of
>> millions of rows is pretty bad.
>>
>> --
>> Andrew J. Kelly SQL MVP
>>
>>
>> "pdxJaxon" <GregoryAJackson@Hotmail.com> wrote in message
>> news:emi80CseFHA.732@TK2MSFTNGP12.phx.gbl...
>>> hard to really answer but here is the question
>>>
>>> Your opinions please:
>>>
>>> I have 6 15k 72GB Drives in an ARRAY.
>>>
>>> This box is ONLY used for doing Data Transfer Activity (DTS Jobs) (NO
>>> OLTP, Etc)
>>>
>>> All I do on this box is copy mass amounts of data from DB A to DB B.
>>> This is dev platform only not a production environment.
>>>
>>> Generally speaking, I know that placing Logs and TempDB on one Volume
>>> and DATA\Indexes on another Volume is preferable.
>>>
>>> However, My thoughts are this.....
>>>
>>> If I create 2 seperate Volumes, Each volume would have 3 disks. NOT
>>> really that great for Striping (I'm using RAID 0 As I need no Fault
>>> Tolerance)
>>>
>>> If I use 1 Volume, I'm not getting the Seperation that we would
>>> typically want in a production OLTP Environment. However, with 1 single
>>> volume, I will have 6 disks to stripe across.....
>>>
>>> If data copy performance was your main objective, which would likely
>>> provide best performance ?"
>>>
>>> A) 2 volumes with 3 disks each
>>> B) 1 volume with all 6 disks
>>>
>>> This is an HP DL-380 with 6i Raid Controller and 128mb of Battery Backed
>>> Read/Write Cache (Configured 50/50)
>>>
>>> Thanks in advance
>>>
>>> Greg Jackson
>>> Portland, OR
>>> GAJ
>>>
>>
>
>
```