

## Re: Determine Disk Block Size ?

**Source:**

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.server/2005-02/0681.html>

---

**From:** Anthony Thomas (ALThomas\_at\_kc.rr.com)

**Date:** 02/01/05

Date: Mon, 31 Jan 2005 23:03:26 -0600

Well, that doc confirmed at least part of what I read previously:

"SQL Server limits a single read-ahead request depth to 128 pages on most editions. However, Microsoft SQL Server Enterprise Edition raises the limit to 1,024 pages."

Sincerely,

Anthony Thomas

--

"Anthony Thomas" <ALThomas@kc.rr.com> wrote in message  
news:OIoyclBCFHA.3588@TK2MSFTNGP11.phx.gbl...

As always, thanks Andrew. Yes, I am well aware of having to "filter" read information and make judgement calls on the sources reliability. The article I read seemed respectable, though. If I come across it, I'll post a "What's up" response.

Sincerely,

Anthony Thomas

--

"Andrew J. Kelly" <sqlmvpnoospam@shadhawk.com> wrote in message  
news:ewn3Iz7BFHA.2260@TK2MSFTNGP14.phx.gbl...

Not sure which document that came from either. The read ahead is actually dynamic in the size of each read ahead request based somewhat on how fragmented the data is. It can be 8K or 128K, it just depends on if it is possible to do so or not. Most of the documentation for 2000 is wrong in that it may state you need multiple files to spawn multiple threads. That is definitely not the case and this is supposed to be fixed in the 2005 docs. Yes it can spawn a thread of each file but it is not limited to a single thread per file. Even with a single file it is possible to spawn multiple threads to read the same file. Here is what looks like a very interesting article on IO in 2000. I haven't had a chance to read it yet (hope to in the next few days) but it looks very promising.

<http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/sqlIObasics.msp>

--

Andrew J. Kelly SQL MVP

"Anthony Thomas" <ALThomas@kc.rr.com> wrote in message  
news:u8xc256BFHA.1564@TK2MSFTNGP09.phx.gbl...

>I just read that SS2K SE spawns Read Ahead requests in 4xExtents, 128 KB  
> chunks, and spawns a thread per data file in each Filegroup. It also  
> mentioned that EE can exceed the number of concurrent extents per request  
> and managed the thread spawning more efficiently. But, damn if I can't  
> remember where I came across the information.

## microsoft.public.sqlserver.server: Re: Determine Disk Block Size ?

```
>
> Anthony Thomas
>
>
> --
>
> "Andrew J. Kelly" <sqlmvpnoospam@shadhawk.com> wrote in message
> news:uNjLEV5BFHA.2180@TK2MSFTNGP10.phx.gbl...
> An I agree with your comments except for the last one:
>>>>
> Furthermore, more files per Filegroup will spawn
> additional disk threads per request.
> <<
>
> That was true with 7.0 but is not necessarily true with SQL2000. In 2000
> Sql Server can spawn multiple threads to read a single file as well. So
> just having multiple files does not mean it will be more effecient or
> spawn
> more threads than a single file.
>
>
> --
> Andrew J. Kelly SQL MVP
>
>
> "Anthony Thomas" <ALThomas@kc.rr.com> wrote in message
> news:ep16rwlBFHA.3592@TK2MSFTNGP09.phx.gbl...
>> Not that I disagree with the RAID-10 argument, I don't; however, I do
>> disagree with your comments about NTFS Cluster Sizes (a.k.a., block
>> size).
>>
>> The MS SQL Server Operations Manual makes a recommendation of using 64-KB
>> cluster sizes. We have done so with a dramatic total disk I/O
>> throughput.
>> It makes sense. SQL Server reads and writes data in 8KB data pages, but
>> issues Scatter-Gather I/O requests in 4 or more Extents (8 x 8KB pages)
>> simultaneously. If your block sizes are on these 8KB boundaries, SQL
>> Server
>> Disk I/O requests can be more efficient.
>>
>> Your biggest disk throughput bottleneck is total number of I/O Operations
>> a
>> single spindle can handle. You're throughput will be based on type of
>> access, sequential versus random, and the size of each request.
>> Optimizing
>> the request sizes coupled with a fixed maximum number of I/O Ops, should
>> make your system as efficient as possible.
>>
>> Nevertheless, and along the same lines as the RAID arguments already
>> presented, this is also suggested from the Operations Manual: more
>> smaller
>> disks are more efficient than fewer larger disks. This would come into
>> play
>> in the 0 part of the RAID-10 argument. Striping allows more physical
>> spindles to fulfill requests, each one capable of servicing a theoretical
>> maximum number of I/O Ops. Furthermore, more files per Filegroup will
>> spawn
>> additional disk threads per request. This coupled with number of disks
>> can
>> greatly improve throughput.
>>
>> Sincerely,
```

Re: Determine Disk Block Size ?

microsoft.public.sqlserver.server: Re: Determine Disk Block Size ?

```
>>
>>
>> Anthony Thomas
>>
>>
>> --
>>
>> "Mike Epprecht (SQL MVP)" <mike@epprecht.net> wrote in message
>> news:uY2FhGsBFHA.3840@tk2msftngp13.phx.gbl...
>> Hi
>>
>> Rather make sure that your drives are configured correctly.
>> Make sure your logs on RAID-10, and preferably your data on a separate
>> RAID-10 volume. RAID-5 hurts more than it helps.
>> As Geoff said, block size is not important, but getting maximum IO
>> throughput is.
>>
>> Regards
>> -----
>> Mike Epprecht, Microsoft SQL Server MVP
>> Zurich, Switzerland
>>
>> IM: mike@epprecht.net
>>
>> MVP Program: http://www.microsoft.com/mvp
>>
>> Blog: http://www.msmvps.com/epprecht/
>>
>> "Geoff N. Hiten" <SRDBA@Careerbuilder.com> wrote in message
>> news:uhfJBgXBFHA.2180@TK2MSFTNGP12.phx.gbl...
>>> Unless you have an overriding reason, go with the hardware
>>> manufacturer's
>>> recommendations. These are typically the default settings. Changing
>>> the
>>> NTFS block size usually has no measurable effect on system performance.
>>>
>>> --
>>> Geoff N. Hiten
>>> Microsoft SQL Server MVP
>>> Senior Database Administrator
>>> Careerbuilder.com
>>>
>>> I support the Professional Association for SQL Server
>>> www.sqlpass.org
>>>
>>> "Joe K." <Joe K.@discussions.microsoft.com> wrote in message
>>> news:B77A9710-86AA-4319-B539-6F65D91E5ECA@microsoft.com...
>>> >
>>> > I have a SQL Server 2000 Enterprise database that processes 300
>>> > transactions/sec, approximately 300 users, and 250 GB in size. How
>>> > can
>>> > I
>>> > determine the the appropriate disk block size that should be on my
>>> > system.
>>> >
>>> > Thanks,
>>> >
>>> >
>>>
>>>
>>>
>>>
>>>
```

microsoft.public.sqlserver.server: Re: Determine Disk Block Size ?

>  
>