Re: Clustering for Performance

Source:
http://www.tech−archive.net/Archive/SQL−Server/microsoft.public.sqlserver.clustering/2006−01/msg00071.html

- *From*: "Andrew J. Kelly" <sqlmvpnooospam@xxxxxxxxxxxx>
- *Date*: Wed, 11 Jan 2006 13:11:19 −0500

SQL 2005 has lots of performance enhancements among other things. Just because it does not work like RAC does not mean it can not scale or handle large workloads. I just finished working on a system that was doing over 25K batch requests per second with upwards of 1000 concurrent (active and real connections) and it was hardly breaking a sweat. Please don't make any decisions based on appearance or misconceptions.

−−
Andrew J. Kelly SQL MVP

"Richard" <Richard@xxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:427BB31F−57CA−43A9−965C−E5F18196D696@xxxxxxxxxxxxxxx
> Thank you Geoff and Andrew for you swift responses.
>
> We are actually looking for a consultant to come in and help with this.
> However I am just doing some ground work beforehand.
>
> I was hoping that SQL Server 2005 would give us some performance
> improvements (such as true Load Balancing like Oracle Real Application
> Clusters), but the improvements over 2000 seem to be mostly for failover
> rather than scaling.
>
> Thanks again for your input
>
> Richard
>
> "Andrew J. Kelly" wrote:
>
>> Richard,
>>
>> Try not to take this the wrong way but if you are asking questions like
>> these and need to implement a system that large you are probably a bit
>> over
>> your head. I suggest you seriously consider hiring a consultant who has
>> been thru things like this before. Quite honestly there aren't many
>> people
>> who have dealt with systems that large. And if not done correctly it will

>> almost certainly turn belly up and die when you even get close to that
>> many
>> users. That said here are a few comments. One is I doubt you will really
>> have 40K concurrent users especially since SQL Server only allows
>> 32,767<g>.
>> Even with heavy use web based apps you rarely have as many concurrent
>> connections as you would think. And if you are talking anywhere near
>> this
>> amount you are talking some serious hardware to support it. As I
>> mentioned
>> in another post Clustering is not a load balancing option. Only one node
>> at
>> a time has control over a specific disk resource in the cluster. So even
>> with Active / Active (or more correctly Multi−Instance) you can't share a
>> database since it resides on only one disk resource. You can't design a
>> system like this in a newsgroup and if you try you will fail. It
>> requires a
>> very careful and well laid out plan to implement a large scale database
>> application.
>>
>> −−
>> Andrew J. Kelly SQL MVP
>>
>>
>> "Richard" <Richard@xxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
>> news:A115E5ED−8801−4142−B921−FD25665FE3C0@xxxxxxxxxxxxxxx
>> > Hello All. I need to set up quite a large SQL System 2005 system that
>> > needs
>> > to deal with 40,000 concurrent users, and one of the tables will
>> > contain
>> > blob
>> > data. 40,000 users throwing around Mb's of data at the same time
>> > worries
>> > me a
>> > little bit!
>> >
>> > My question is to do with Database Clustering and Mirroring. From what
>> > I
>> > can
>> > see, there is still no load−balancing with SQL Server 2005, so does
>> > this
>> > mean
>> > even in a clustered environment I am still basically only using a
>> > single
>> > database server? I have seen many posts that tell me that clustering is
>> > ONLY
>> > for failover and not for performance. I understand that with
>> > Active/Passive
>> > this is the case, but how about Active/Active? If I can set up
>> > Active/Active
>> > (2 nodes? 4 nodes? 8 nodes? how many are possible?) with a SAN and NOT

>> > have
>> > failover implemented (can I turn failover off?) then would I have a
>> > load
>> > balanced environment? I could then have all nodes running up to 100%
>> > (because
>> > I don't have to worry about the failover) and therefore give me a
>> > dramatic
>> > increase in performance compared with using a single server?
>> >
>> > If I can do this then I can set up 2 identical clusters and an extra
>> > server
>> > for the witness, and use database mirroring for failover? Of course I
>> > understand that mirroring will decrease performance on the clusters.
>> > But
>> > would this give me a super−fast database system that might cope with
>> > what
>> > I
>> > need?
>> >
>> > Also, I am thinking about taking the blob data of of the database and
>> > create
>> > a new database that just deals with the blobs. How would this affect my
>> > clustered/mirrored environment?
>> >
>> > Thanks
>> >
>> > Richard
>>
>>
>>

.

- ***Follow−Ups***:
  - ♦ ***Re: Clustering for Performance***
    - ◊ *From:* Geoff N. Hiten

- ***References***:
  - ♦ ***Re: Clustering for Performance***
    - ◊ *From:* Andrew J. Kelly
  - ♦ ***Re: Clustering for Performance***
    - ◊ *From:* Richard

- Prev by Date: ***SQLServer Agent Properties/Connection/Windows−SQL Auth***
- Next by Date: ***Re: Clustering for Performance***
- Previous by thread: ***Re: Clustering for Performance***
- Next by thread: ***Re: Clustering for Performance***
- Index(es):

- ***Date***
- ***Thread***