

Re: IO Bottleneck

Source:

<http://www.tech-archive.net/Archive/Exchange/microsoft.public.exchange.design/2006-08/msg00115.html>

- *From:* "Al Mulnick" <amulnick_No_SPAM@xxxxxxxxxxxxx>
 - *Date:* Sat, 26 Aug 2006 10:49:39 -0400
-

Just to point out: it's not the space as much as the separate spindles that you're after. That separation is one of the best performance related investments you can make. As was explained earlier and as you've noticed, Exchange is very disk dependent. Very. Several things can be done to help reduce the impact, but at the end of it all, the io has to be handled at some point.

Understand that there are several different things going on for the life span of a message submission/delivery/retrieval event (all of these can be going on at the same time.) A lot of it is going to utilize the disk. There will be several different io types i.e. sequential write, random read/write, and sequential read are primarily what you'll have to account for. Putting them all on the same physical spindles can have a serious impact on performance due to disk contention. Of those, the one io type most noticeable to users is the sequential write of the log files. This is an artifact of the two-phase commit database technology employed by the jet engine. You want that to occur, trust me. But if that write to the log file is slowed, then store gets blocked and your users may notice it.

So what can you do about it? For starters, you can optimize the disk subsystem so that these events don't impact one another nearly as much. Separating io types is the way to achieve that and by far putting the log files on their own separate set of write-optimized physical disks is a low cost, and easy way to do that. Using raid5 for log files is not recommended for that reason. raid 10 or 1 depending on the sizing information is often recommended and used with great success.

Another thing you can do is use cached mode for your users. This has a bigger impact when first fired up, but the user is less impacted later.

You may also want to check out your global catalog servers. If the majority of your users are OWA users, then the impact on the server should be much different than what you describe and disk should be less of a concern in many cases. Your case may be different, but generally speaking. You should also take the advice posted and find out what the message patterns are: are there a lot of large blobs being submitted? Are there other applications affecting you? etc.

Re: IO Bottleneck

My \$0.04 anyway.

<jamestechman@xxxxxxxx> wrote in message
news:1156557888.578680.31840@xx
Yes log files grow in increments of 5MB. 73GB is overkill for logfiles.
36GB is the smallest HD you can get for the HP servers, which will
suffice.

James Chong

Ariel wrote:

Thanks, I will try to collect some data with exmon.

One thing I've noticed is that my log files never grow above 5MB each...
35MB total... is that normal? Should I really use two 73GB disks to store
just 35MB?

"Simon Walsh" escribió:

Run an EXMON on the server.
Look for unusual user activity in real time.
Start with the worst offenders. Turn their Outlook client off. Any
change?
What version of Outlook are they using? Have you installed all the
latest
fixes for those clients?
Are their bunches of users that are sending/receiving the exact same
amount
of bytes/second?
Just a few tips

/Simon.

"Ariel" <Ariel@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:0B83F4E4-B6EB-45AD-8AFF-DD46FEF3C47F@xxxxxxxxxxxxxxxxxxxxxxxx

My databases are 36GB in size.
I am not running blackberry.
Most of my clients use OWA.
I am not sure how to determine if the users have desktop
search, I saw
a
couple of ones, but it's probably 5 users in total.
I am not running indexing or message tracking.

PTA says that user activity is unusually high.

Re: IO Bottleneck

Thanks for all your help

"jamestechman@xxxxxxxxx" escribió:

You're absolutely correct, this server should be able to support substantially more than 300 users. Let's try to get some more info as to the source of your I\O bottleneck in terms of user IOPS. How big are your user's MB? Are you running Blackberry? Are your clients running any desktop search engines, ie. google desktop search, yahoo desktop search... Are you doing any Exchange indexing? What other applications are running on the server?

James Chong
MCSE M+, S+, MCTS, Security+
msexchangetips.blogspot.com

Ariel wrote:

Bummer, I really don't know what I'm gonna do now... probably change the delay for the RPC dialog so I don't get so many complaints, until I am able to buy a new server. What really bothers me is that this server, properly configured, should be able to service much more than 300 users...

"John Fullbright [MVP]"
wrote:

Re: IO Bottleneck

Write cache
could
remorve 1
IO
operation
out of 4 for
RAID 5.
Instead of
defining
write
performance
as $P*N/4$,
you could
descibe it as
 $P*N/3$.

For your 4
drive RAID
5 array
using 10K
drives, your
current
write
performance
would be
something
like $85*3/4$
or 61 IOPS.
Adding
the
write
cache could,
depending
on
controller
architecture,
change that
to
 $85*3/3$ or
85 IOPS. I
doubt that,
in practical
terms, it
would make
much
difference.

"Ariel"
<Ariel@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx>

Re: IO Bottleneck

wrote in
message

news:C812D79F-A273-4DA4-908A-79514F8E5096@xxxxxxxxxxx

John:

Thanks
for
your
response.

The
way
you
describe
is
exactly
how

I
would
design

a
new
server.

However,
this
is
our
only
exchange
production
server,
and

I
need
to
improve
performance
as
much
as
possible
without
reinstalling
it.

So,
would
adding
the
battery
backed

Re: IO Bottleneck

write
cache
improve
performance
in
a
significant
way?
And
would
I
be
better
of
moving
something
else
to
the
logs
disk,
say,
the
smtp
queue
or
temp
dirs?

"John
Fullbright
[MVP]"
escribió:

At
a
minimum,
you
should
have
a
mirror
for
the
OS,
a
mirror
for
the
logs,

Re: IO Bottleneck

and
a
mirror
(or
RAID
10)
for
the
databases.
I'd
hazard
a
guess
that
the
write
times
on
the
the
RAID
5
array
far
exceed
20ms.

"Ariel"
<Ariel@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>
wrote
in
message
news:F0839239-546B-4AA9-8B55-5B0E7A4DC9

I
have
recently
become
the
administrator
of
an
Exchange
2003
Server.
These
are
the
server
specs:

Re: IO Bottleneck

HP
Proliant
DL380
G3
Dual
Xeon
2.8GHz
2GB
RAM
Smart
Array
5i
Controller
with
64MB
and
read-only
cache
4
x
300GB
Ultra320
10K
Hard
Drives
in
a
RAID
5
array

I
have
been
assigned
the
task
to
investigate
a
possible
bottleneck,
since
users
were
seeing
the
RPC
dialog
frequently.

Re: IO Bottleneck

Re: IO Bottleneck

I
ran
the
Performance
Troubleshooting
Analyzer,
and
of
course,
the
results
indicated
a
disk
bottleneck.

The
server
has
room
for
2
more
hard
drives.
I
tried
adding
a
73GB
disk,
and
using
it
for
the
database
logs,
then
ran
the
PTA
again,
and
now
the
report
says
that
I
have

Re: IO Bottleneck

a
disk
bottleneck
in
all
the
roles
that
remained
on
the
primary
array.

I
am
going
to
add
another
73GB
disk
to
make
a
RAID
1
array,
and
move
the
logs
anyways,
but
I
was
wondering
if
I
should
transfer
another
role
to
that
new
volume.
I
am
also
considering

Re: IO Bottleneck

Re: IO Bottleneck

buying
an
upgrade
for
the
controller
to
add
128MB
of
battery-backed
write
cache,
if
it
would
help.

I
would
appreciate
any
advice
you
have
to
offer.

Regards,

Ariel.