

Re: Exchange Hardware Question

Source:

<http://www.tech-archive.net/Archive/Exchange/microsoft.public.exchange.design/2007-01/msg00030.html>

- *From:* "John Fullbright [MVP]" <fjohn@donotspamnetappdotcom>
 - *Date:* Fri, 12 Jan 2007 17:45:38 -0800
-

Store = RAID 1 or RAID 10 depending on the usage.

1. The dell, with three drives in a RAID 5 configuration, implies the OS was built using the dell config utility. This is a worst case for performance. When an OS is built using such a utility, the file system begins life as FAT and is converted to NTFS. This results in an allocation unit size of 512 bytes. This means a write to the page file (assumes page file on the os drive which is typical in small deployments). This means every 4K write to the page file will result in a split IO and 8 IOs to the drive. That goes for anything else on the OS drive. You can see this by looking at the perfmon physical disk counter split IOs per second. I have one customer who did a study comparing this default configuration to a mirror and manually building the OS with a 4K (the default) allocation unit size. The measured performance increase was 400%.

2. You don't mention the type of drive, but with a perc it will be SCSI. For a 20ms response time you can use a P value of 85 for 10K drives and 130 for 15K drives. For RAID 5 read perf = $P * (n-1)$ where n is the number of spindles and P is the IOPS/spindle of the drive type you are using. Write perf = $P * (n-1) / 4$. For exchange you have to apply the read:write ratio to determine the mixed value. assuming 10K drives and a 3 spindle array, read perf = 170 IOPS and write perf = 42.5 IOPS. With a 1:1 read:write ratio, the mixed perf is 106 IOPS.

For a 2 drive RAID 1 mirror, read = $P * N$ and write = $P * N / 2$. With 10K spindles write perf = 170, read perf = 85 and mixed perf = 127.5 IOPS. That's a 20% performance increase over RAID 5 with a third less spindles.

I suppose this is why the MS paper "Optimizing Storage for Exchange Server 2003" states "In general, Raid-5 does not provide the best trade-off between reliability/availability and performance". Considering the math, I'd call that an understatement. I would not recommend RAID 5 except in extremely small Exchange 2003 deployments (under 50 IOPS load).

<http://technet.microsoft.com/en-us/library/bb124875.aspx>

You don't state the number of users or your IOPS requirement, however the configuration Bharat talks about provides:

Re: Exchange Hardware Question

85 IOPS for the OS
85 IOPS for the logs
127.5 IOPS for the databases

If you move to 6 spindles and replace the mirror with a RAID 10, then

85 IOPS for the OS
85 IOPS for the logs
255 IOPS for the databases

This configuration would require 6 drives and cover many small organizations. Moving to 15K drives in RAID 10 would give you

130 IOPS for the OS
130 IOPS for the logs
540 IOPS for the databases

In RAID 1 you would get

130 IOPS for the OS
130 IOPS for the logs
270 IOPS for the databases

You could even go three mirrors to provide fault tolerance for all the volumes and, in 6 15K spindles, have

270 IOPS for the OS
270 IOPS for the logs
270 IOPS for the databases.

John

"Bharat Suneja [MVP]" <bharatsuneja@xxxxxxxxxxxx> wrote in message
news:efYhVZpNHHA.1252@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

Before you start – how many users? What's the usage? Disk I/O required?

Generally speaking, no... but if you want to optimize (without knowing more details..), the rule of thumb **given the number of drives you have =6**:

OS: RAID1 mirror
Transaction Logs: RAID1 mirror
Store (in this situation): RAID1 mirror

--

Bharat Suneja
MVP – Exchange

Re: Exchange Hardware Question

www.zenprise.com
NEW blog location:
www.exchangepedia.com/blog

"-=gu=-" <gu@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:D84C18AC-F2A9-405A-8549-765C9D0765AB@xxxxxxxxxxxxxxxxxxxx

Hello,

We are a single site running Exchange 2003 on 2003 server.

I am not too far from running out of disk and would appreciate your thoughts on migrating to a different server. Currently, my priv1.edb and stm are totaling close to 40 GB and I have only 15 GB Free. The existing server is ancient.

The company is going to grow (almost double) in 2007.

I have a decent server waiting in the wings doing nothing which I can use. It's a Dell PowerEdge 2850 that currently has a PERC raid controller and (3) 147 GB drives. There are an additional 3 empty bays, which I could also fill with 147 GB drives. Under a RAID5 configuration, I should get 700+ GB of disk.

I know that the optimal configuration of Exchange would be to use separate redundant (not partitioned) drives for: system \ pagefile \ logs and exchange.

So my question is...
If I buy the 3 additional drives and put them all on a single (or even two RAID5s), am I asking for trouble in terms of a disk I/O bottleneck?

Thank you all very much for your responses.
-=gu=-

Re: Exchange Hardware Question