Re: IOPS and megacycles calculations

Source:
http://www.tech−archive.net/Archive/Exchange/microsoft.public.exchange.design/2007−03/msg00051.html

- *From*: "Neil Johnson" <neiljohn@xxxxxxxxxxxxx>
- *Date*: Wed, 21 Mar 2007 14:09:20 −0000

Hi Tony,

The important thing is to understand "what" your peak load and not especially "when" it occurs. So, if you look over your data for the month, pick the period of peak load during working hours. Use the average for a 2hr period across the peak to calculate your IOPS/mailbox figure.

Once you have an IOPS/mailbox figure you can use it to scale as high as you want from a storage perspective. Dont forget that we dont support more than 4000 mailboxes per Exchange 2003 server though due to the 32bit memory constraints.

Neil.

"TonyP" <TonyP@xxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:46B30508−7444−41F4−BAEB−E9961C83ACB5@xxxxxxxxxxxxxxx

> Hello
>
> I am trying to capacity plan for a migration of 30,000 users to Exchange 2003. Currently the users are based on a Sun iPlanet environment using only POP and IMAP connections.
>
> In the exchange environment initially they will use POP/IMAP and MAPI then overtime we forecast about 20,000 MAXIMUM users to be configured for MAPI. Actually we cannot even see more than 15,000 ever being moved to MAPI.
>
> Currently we also have a four node exchange 2003 cluster, 3 Active 1 Passive hosting 20,000 users in MAPI with a backend SAN environment.
>
> Our clustered environment is hosted at our Data Centre and all our offices over the city connect back to us via Fibre. All users connect via MAPI and there a few using RPC over HTTPS. There is also a front end NLB consisting of 3 servers for OWA but this is mainly used for OWA and OMA connections.
>
> I am using the Microsoft white paper Optimizing Storage for Microsoft Exchange Server 2003 to work out IOPS and Megacycles per mailbox.
>
> I have monitored data for 2 weeks period on the current cluster and I seem

to have no 2 hour busy period which is consistent (Monday mornings), it seems to change from day to day and hour to hour .. my users are in Spain and working habits here are not the same!!

I intend to collect data for about one month then analysis this data and use statistical averages etc to come up with a figure. I will come up with an IOPS and Megacycles for a 2 hour busy period every day. Then analysing this set of information to come up with an average IOPS and Megacycles based on data collected every day.

My intention is come up with an IOPS and Megacycle for the current 20,000 users. Then use this calculation to size hardware for the 30,000 users I intend to migrate.

I wanted to check on two things:

1) If my approach to use the IOPS and Megacycles from the 20,000 users can be used for planning for the 30,000 users?

2) MORE importantly looking at tonnes of documentation on the web I am seeing some people suggesting to monitor logical disk transfers and others suggesting Physical disk transfer in Event viewer. Which do I use? Four node cluster on HP hardware backend is Storagetek SAN.

Thanks

T

.