

Re: Maximum Exchange Users

Source:

<http://www.tech-archive.net/Archive/Exchange/microsoft.public.exchange.admin/2006-01/msg02942.html>

- *From:* "John Fullbright" <fullbrij@xxxxxxxxxxxxx>
 - *Date:* Tue, 24 Jan 2006 19:56:44 -0800
-

AD/Exchange 2003 can easily handle 300K users.

A good place to start is:

<http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/default.mspx>

The planning and architecture section.

Some things you should start thinking about are:

- Level of consolidation
- Feature set
- Service Level Agreements
- Business Continuance & Disaster Recovery

I guess at this point I need to disclose that I do in fact work for a storage vendor, and specifically do large scale architecture and design for Exchange on that vendor's platform. Using experience drawn from previous deployments, I'd like to point out a few technologies you may want to consider.

1. iSCSI. At this point the iSCSI protocol is mature enough to have a proven track record of performance and reliability. For 1G ethernet, the maximum throughput using the software initiator (MS initiator 1.06 or 2.0) is roughly 115 megabytes per second. The overhead on the wire is very low compared to FCP, and actual observer throughput for Exchange 4K IOs is in the area of 25,000 IOPS. Virtual memory limits in 32 bit Exchange will kick in long before any given virtual server gets anywhere close to this IO level. I have actually personally observed this in multiple customer deployments. The vendor I work for has over 3000 customers deployed on iSCSI, and is the current iSCSI market leader with a 43% market share. Compare this throughput to 1G FCP (~30MB/sec with a 4K IO size) and 2G FCP (~ 58MB/sec with a 4K IO size). Yes, 4G and 10G trunked FCP is coming, but so is 10G ethernet. Certainly this is on an isolated swith infrastructure, however the cost of copper and copper switches is far lower than Fibre and Fibre Channel switches.

2. Snapshots. If you consolidate to any degree, there aren't enough hours

Re: Maximum Exchange Users

in the day to backup to tape. Snapshots only take a few seconds
irregardless of the size of the LUN you back up.

3. Disk to Disk backup and change delta replication. After the initial
baseline replication, replication change deltas to secondary storage is
clearly the way to go.

4. Remote replication. With iSCSI and WAM acceleration technologies, not
only is this possible, but a standard component of the business continuation
plan for many companies.

5. Geographically dispersed clusters or hot standbys. Ensure continued
operation even in the event of the failure of an entire datacenter.

One last point, carefully evaluate your storage performance requirements as
well as your storage space requirements. For Exchange, the number of
spindles is primarally driven by performance, not space. Don't fall into
the trap of buying fewer larger spindles; you'll regret it down the road.
The same goes for the choice of RAID type on your storage; whatever you
choose, remember to factor in the write penalty for your chosen RAID type.
When you evaluate storage vendors, don't get conned into RAID 5. You'll pay
for that mistake though performance pain or more costly purchases on that
back end of the deal to add enough spindles to make the pain go away.
Unless the vendor has a mechanism to limit the impact of comingling, don't
mix LUNs withing a storage group. Comingling leads to costly and hard to
troubleshoot IO problems because IO against a LUN on a given server can
negatively impact IO on another unrelated server if the LUNs share the same
physical spindles. On the last two points, write penalty and comingling,
there is only one storage vendor that has no write penalty and has a
mechanism to limit the impact of comingling. That vendor is the one I work
for, that also happens to be the iSCSI market leader. That vendor is
Network Appliance. Do the math. I can meet or beat the performance of any
other vendor in 25% or more less spindles.

John

"CR" <CR@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:43DD5D21-AB7A-4A45-A1BD-50D91E8AED21@xxxxxxxxxxxxxxxxxxxx
> I'm looking for a good Exchange architecture sizing whitepaper, document,
> anything. We're looking to implement a private email system for about
> 300,000 of our users in the US utilizing Exchange. Each user would be
> limited to 10MB of disk space so I know we'd be looking at a few TB total.
>
> Obviously we'd be looking at numerous quad-clusters connected to a SAN but
> I
> was wondering if AD/Exchange could even handle that many users in a
> multiple

Re: Maximum Exchange Users

- > front-end/back-end configuration.
- >
- > Anyone have any ideas where to start?

- *Follow-Ups:*

- ◆ *Re: Maximum Exchange Users*

- ◇ *From: lseamans*

- Prev by Date: *Re: Deleted Folder...Recovery?*
- Next by Date: *exchange webbased page problem!*
- Previous by thread: *Re: Maximum Exchange Users*
- Next by thread: *Re: Maximum Exchange Users*
- Index(es):
 - ◆ *Date*
 - ◆ *Thread*