

## Re: Database Size, Database Corruption, Message Store & Storage Group

**Source:**

<http://www.tech-archive.net/Archive/Exchange/microsoft.public.exchange.admin/2005-01/2848.html>

---

**From:** Ben Winzenz [Exchange MVP] (*ben\_winzenz\_at\_NOSPAMdotmessageonedotcom*)

**Date:** 01/20/05

Date: Thu, 20 Jan 2005 15:44:14 -0600

To be honest? More careful monitoring of the server. Checking daily that the backup job completes successfully probably would have also helped. Daily checking of the event logs for any warnings or errors. Quarterly (or thereabouts) testing restores to make sure that they are good. I'm not trying to come across as being hard on you, but there really is no good reason for discovering all the things that you mentioned "at time of disaster". They are all things that should have been discovered as soon as they happened. What you can do differently in the future is make a list of things that you should be checking on a daily basis so that problems can be remediated as soon as they occur. Again – take this as constructive feedback rather than bashing on you. :-) Most of us have been there at one point or another. The best thing you can do is learn from what went wrong.

Another thing I want to mention is that you may be better served filling up an entire storage group before creating a new one. The reason I say this is because each extra database requires additional memory when Exchange starts. Each extra storage group also requires more memory at startup, but it is quite a bit more than a mailbox store requires. Microsoft generally indicates that the best practice for server utilization is to fill up the entire storage group first for this very reason. If you have 8 mailbox stores, you may consider consolidating down to 2 storage groups. If you are using too much physical memory, you will start to have excessive disk paging, which can have an impact on performance.

Also it is important to know if there was actually database corruption. If there is, then there will be corresponding events logged to the server's application log. Events such as 1018, 1019 or 1022 errors (these refer to corruption within the edb file – there is another set corresponding to the stm file). If there was database corruption, then more than likely it was the result of faulty hardware. In this case, it could have possibly been flaky hard drives, but the raid controller is also something you should look at. The figure that gets tossed around is that 99.9% of all database corruption in Exchange is caused by faulty hardware.

In terms of backups, you should look at "how" you are performing backups.

Are you simply doing a normal online backup of the Information Store? Are you also doing Brick-level (Mailbox-level) backups via a 3rd party backup agent? Do you have file-level antivirus installed on the Exchange server?

As far as individual database size, that is really up to you. Normally, I would say that database size should be dictated by your SLA's. If your restore window is only 4 hours, then you should size your databases so that they can be restored within that time frame. You should also factor in maintenance. Normally, Exchange doesn't require daily maintenance – it is pretty good at taking care of itself. However, if you should have to run eseutil against the database, it can take a while.

--

Ben Winzenz

Exchange MVP

"bc" <bc@discussions.microsoft.com> wrote in message

news:BA9B0773-8B6F-4062-9231-8E91EEC31521@microsoft.com...

> We're just recovering from an Exchange disaster. We discovered that we  
> could

> not get database backups off of two of our four storage groups. We then  
> discovered we were down one drive in our RAID 5 array. Then we discovered  
> that we could not rebuild onto a new drive, probably due to issues on a  
> second drive. We then attempted to save as much data as we could by  
> exmerging data out of the two damaged stores and moving mailboxes to a  
> second  
> server. In the midst of this, the server crashed and would not reboot.

>  
> We've since rebuild and recovered what we could, but we want to know if we  
> could do things differently in the future that might help us recover more  
> quickly. Especially since we found that one of the moved stores wouldn't  
> back up.

>  
> We have 4 storage groups on the server in question. Two databases in each  
> storage group with a total of about 100 GB between the 8 of them. We  
> found  
> exmerging and/or moving the data to try to save it to be extremely time  
> consuming. We'd like to hear from anyone who might have suggestions or  
> similar experiences. Should we decrease the size of our databases by  
> splitting them up? Should we try doing an off-line consistency check?  
> Are  
> there online alternatives to repairing a database that won't back up?