

# Re: DC RAID Configuration

---

*Source:*

[http://www.tech-archive.net/Archive/Win2000/microsoft.public.win2000.active\\_directory/2006-11/msg00284.html](http://www.tech-archive.net/Archive/Win2000/microsoft.public.win2000.active_directory/2006-11/msg00284.html)

---

- *From:* "Joe Richards [MVP]" <[humorexpress@xxxxxxxxxxxxx](mailto:humorexpress@xxxxxxxxxxxxx)>
  - *Date:* Sat, 25 Nov 2006 13:20:28 -0500
- 

Microsoft's current standard config for Domain Controllers is DL585's with 8-10 spindles. 6 spindles are for the DIT, they are in a hardware based 0+1 RAID configuration. If the machine has 6 spindles then OS/Logs go on a hardware mirror, if 10 then they go on a 4 disk hardware based 0+1 RAID.

The number of spindles isn't for amount of disk, it is for IOPS capability. A DC in a large environment with Exchange can take a lot of pounding and you want as much spare IOPS as possible to keep up with it. Their DIT is only in the teens in size so the disks themselves could be small, you just need many to get the high IOPS (read OPS really) that make AD really run smoothly. x64 does help out but that is primarily with LDAP only functions. Authentication is not helped out nearly so much and disk is still a bottle neck there, this is pretty clearly visible in the x64 whitepaper MSFT has put out.

The only official recommendation I have seen from MSFT on RAID configs is in the deployment guide but that is for smaller environment. That document talks about 3 RAID sets, usually 3 mirrors (OS/Logs/DIT), however they mention that the DIT set could be a 0+1/1+0. I think that is 2 RAID sets too many for any deployment unless the DC is something other than just a DC.

Personally, most folks with larger environments seem to deploy DCs that you can get only 6-8 disks into and for those occasions I tend to recommend a single RAID 0+1/10/5. The perf difference between 0+1/10 and 5 is usually 10% or less. If you absolutely want the fastest RAID, go with 0+1/10 over the 5. At least that is all of the IO metric testing I have seen recently. Go back 4-6 years and I actually recall seeing the Dell PERC card's RAID 5 beating out its RAID 10 config by about 5%. Probably just poorly written firmware for RAID 10 I would guess.

30 DCs with 2500 users total, unless the environment had a lot of very heavy use AD apps I would likely be ok for even a single mirror for those DCs. I would load that many users on a single DC in a heartbeat though I would never use just a single DC.

Yep, I also don't care to split out OS/Logs/DIT. On a normal busy DC that is simply a DC, the OS and Logs are but a rounding error in IOPS next to the DIT. That means you are wasting IOPS on the OS/Logs sets that the DIT could make good use of. The page file is also something I am not terribly worried about on a DC because most of the memory should be used for caching the DIT, that is never cached out, if memory pressure is too great, the cache will trim to allow other things into memory. Very little to nothing else really should be running on DCs so you don't have other worries about paging.

The one time I have seen a need for a separate Logs drive was for a test done in the Microsoft Enterprise Computing Center where one of my friends built a 2TB DIT for ADAM in a month. He was adding tons of new objects and of course the Logs were seeing tremendous amounts of IO for that. Very much a corner case with AD as AD is very much read oriented.

## Re: DC RAID Configuration

If you like running all sorts of stuff (crap) on a DC, then there is no generic formula, you get to sit in a lab and work out the perf on what is best for you and how your DCs are being used. The main thing to keep in mind is that everything you add to a DC is adding additional surface area and attack vectors on the machine that is the core of your Microsoft security infrastructure.

--

Joe Richards Microsoft MVP Windows Server Directory Services  
Author of O'Reilly Active Directory Third Edition  
[www.joeware.net](http://www.joeware.net)

---O'Reilly Active Directory Third Edition now available---

<http://www.joeware.net/win/ad3e.htm>

Mike Shepperd wrote:

I have not seen MSFT produce a definitive guide because every individual situation and environment differ. With many architectural decisions, there are many ways you could go. For my small business customers, we often use six disks on a built-in RAID controller (both SCSI and SATA). First two disks are mirrored and the last four are RAID5 with one hot swap. Then we also backup to tape or online backup provider so that we're not relying on either RAID array to be the only form of backup.

I have heard about how MSFT does their own internal DC's and I have never seen such redundant redundancy in a server spec before or since. I don't remember all the specifics, but they did something like a mirrored stripe set with four disks for the OS, then a mirrored stripe with parity for the second volume with six more disks on a separate controller. Something along those lines. I'd never seen so many disks used in a single server that wasn't hosting huge amounts of data.

If I'm guessing right about the load you'll have on these servers (if they're only DC's and not running other major server apps), with 30 DC's servicing 2500 users distributed across the system geographically, you're averaging between 50 and 100 users per DC. With a load like that, you have a lot of flexibility in your configuration. If you were serving 2500 users on a single pair of DC's, it would be much more important where you put the dit file and the pagefile.

The one thing I would require on a server I was building is the battery backup on the RAID card. It doesn't make a lot of sense to me to spend the kind of money you're talking about, with good redundancy and then end up dropping data off the RAID controller.

In this scenario, I would be tempted to create two stripe sets (with or without parity depending on the performance I needed) and mirror them. I have read a lot of arguments against using RAID5, but in my personal experience, 10 years spanning life as a sysadmin, support engineer at msft and a small business IT consultant, I've never seen two disks in one RAID5 fail at the same time (I hope I didn't just tempt fate).

## Re: DC RAID Configuration

When I worked at msft I occasionally reviewed AD architecture plans for large consulting organizations. I (and my teammates) would give approval on any setup that provided adequate redundancy and I never saw any specific configuration that was used more than another. It was always dependent on the overall scenario more than anything else.

If my assumptions about the load on these DC's are right, I would mirror two 3 disk RAID5 sets on each DC. That's certainly not the most high performance configuration for that hardware, but it also makes it very unlikely that you'll suffer a failure of one of those servers because of disk issues. Higher performance would be to mirror a stripe set (RAID0, no parity), you'd get the redundancy of the mirroring, but faster reads and writes.

If you want to provide more specifics about the environment and expected load on these DC's, that might make it easier to narrow down the scope of options you have.

Hope this feedback helps you decide which way is right for your environment.

Happy Thanksgiving!