

# Ben Armstrong's Virtualization Blog

Hyper-V Program Manager

## Why does it take so long to create a fixed size virtual hard disk?

[Benjamin Armstrong](#)

11 Dec 2008 12:22 AM

15

If you have ever created a fixed-size virtual hard disk that was larger than, oh - 2GB, you probably noticed that it takes quite a while to create. The reason why this takes so long is that when we create a new fixed-size virtual hard disk we take the time to explicitly zero-out all of the disk space that is being assigned to the new file.

Now - we could do this practically instantaneously by not zeroing out the data - but this has an interesting potential security problem.

Imagine the following situation:

- You have a virtual machine with a bunch of confidential data running on a central server (e.g. your company payroll).
- This virtual machine gets moved to a new physical server in response to increased work load.
- You create a new virtual machine which is given to someone on from the in-house dev team - but the virtual hard disk data was not zeroed out.
- Developer then runs data recovery tools on his new, blank virtual machine and is able to recover data from the old payroll server (yikes!)

You see - data is never actually deleted from a disk when a file is moved or deleted (it is just dereferenced) so to avoid the above scenario - we must take the time to "do the right thing" and zero out the VHD contents.

Cheers,  
Ben

**Update: We have provided a tool to create quick, but not secure, fixed virtual hard disks. Details [here](#).**

Tweet

2

Like

0

Share

Save this on Delicious

### Comments



11 Dec 2008 1:28 AM

Kieran Walsh

It would be nice if there was a switch to bypass this for the more usual times when you are creating the fixed sized disks on a brand new server.



11 Dec 2008 2:22 AM

Jon

Doesn't NTFS already guarantee that sectors read from a new file will be zeroed? I think you're just duplicating work the filesystem is doing for you.



11 Dec 2008 3:08 AM

Jevgenij

Having an option to skip file zeroing out would be a great thing.



11 Dec 2008 3:09 AM

**[Benjamin Armstrong](#)**

Kieran Walsh -

We have discussed this, but the problem is that we would be providing a "do this in an insecure fashion if you know what you are doing checkbox" which would need a heck of a lot of text to try and explain to people why you do not want to do it - and then most people would not read the text anyway :)

Jon -

Actually - a couple of folk at Microsoft have just been emailing me on this too. NTFS will zero out a blank file for you - but it zeroes from the beginning of the file up to the point where you tried to write to - which would cause unexpected performance problems. Alternatively it is possible to disable this behavior in NTFS (which is what I was referring to as part of "creating it quickly") which would cause the problem I highlighted above.

Cheers,

Ben

---

11 Dec 2008 8:59 AM



**mityak**

It's good idea not to trust developers ;-). But this approach hurt IT pros when they build new VMs on brand new HDD.

Dmitri

---

11 Dec 2008 10:46 AM



**Ryan**

If not an option in the graphical wizard, how about a powershell switch so we can quickly create them when the need arises. It would go unused unless someone knew enough about what they were doing to do it in the shell.

---

11 Dec 2008 1:17 PM



**Kieran Walsh**

Thanks for the reply Ben.

All these replies show that it's certainly a pain point out in the field.

---

11 Dec 2008 2:23 PM



**Robert**

Is there a script for converting a dynamic disk to fixed? Thanks.

---

11 Dec 2008 10:26 PM



**jsheehan**

NTFS will never return data from a previous file on disk. That would violate the government security standards that it adheres to. Also, NTFS on Win2K8 only allocates blocks as you use them (it has had this capability since Win2K). So, you will not see the performance problem you're mentioning.

---

12 Dec 2008 12:02 AM



**Benjamin Armstrong**

Jsheehan -

When you write to a new file in a location beyond the current valid data length (VDL) NTFS will zero fill the file up to that point and extend the VDL. This isn't a problem on small files or files with sequential data access - but it is problematic on large files that are written to non-sequentially (like a VHD). You can disable this behavior in NTFS by using SetFileValidData to set the VDL to the logical end of the file - see: [http://msdn.microsoft.com/en-us/library/aa365544\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa365544(VS.85).aspx)

Of course this will cause the problem I mentioned above.

Cheers,

Ben

---

16 Dec 2008 12:56 PM



**Dean Steadman**

Is there any official documentation or KBs on this behavior? This doesn't mix well with SAN based thin or on-demand provisioning and I'd like to include this as a best practice doc that I'm working on.

Thanks for the great post and discussion in this thread!

---

16 Dec 2008 6:50 PM



**marcelo sauaf**

"would need a heck of a lot of text to try and explain to people why you do not want to do it"

so MS would never deliver VPN in RAS or ISA, as user machines would be dangerous to server LANs (virus etc - that time didn't have NAP). MS would never deliver format.exe or del in cmd.

Etc etc etc.

Things always must have options to the techies. If it takes too time for creating VHDs, enshort it with an optional, non-default calling parameter, at least.

---

24 Dec 2008 11:39 AM



**IT guy**

The right thing to do is to ask. If you think people should go a particular way set a default.

Instead you waste every one else in the world's valuable time zeroing out what isn't payroll data 99.99% of the time.

This "MS knows best" is what professionals hate about MS products.

---

26 Dec 2008 8:10 PM



**Troy Martin**

I know this is off-topic from this blog-entry, but I was searching for "destroy delete vm slow") and found this blog entry. I manage 8 vm's in my Hyper-V environment and it takes absolutely TOO LONG to delete a VM (and snapshots for that matter too). I've deleted VM's in all sizes and number of snapshots. The more snapshots, the longer. However, even with just one or two snapshots, it can easily take up to 30 minutes to delete the VM. I have seen the same behavior across multiple hosts.

As I type, I'm deleting a VM that has 15 snapshots (I'm building a distributed lab environment and like to get snapshot increments of as I get things working properly), I started deleting it about 45 minutes ago and it is only 11% complete!! That is just not good. My prediction is that it is going to take approx. 8-10 hours to delete ONE VM!! That is crazy...

Again...I'm deleting a VM...not a snapshot.

I am sure the Hyper-V product/feature team knows about this. This is a GLARING performance issue...

---

13 Nov 2013 8:08 AM



**Dave**

Is there any way to calculate the time it will take based on the fixed drive size created? Just a ball park figure?

---