

Slow file system performance on ESXi ...



patparks1 23 posts since

Jun 16, 2009

I've got 3 servers in my environment running ESXi version 3.5.0 (163429). The hardware is the same on all 3 machines. In this environment we have NO san storage space, all local disk

Servers: HP DL 360/Gen 5's, Dual Quad Core CPU's, 20GB of RAM

Datastore 1: 2 x 300GB 10K rpm 2.5 SAS drives in a RAID 1 mirror,

Datastore 2: 4 x 300GB 10K rpm 2.5 SAS drives in a RAID 5 mirror,

As a test, I created a virtual machine running CentOS linux5.3 and gave it a 2GB hard drive. I installed CentOS on the drive. This virtual machine lives in datastore1.

Now, I go into the VMWare Infrastructure client, go into the datastore browser and I copy the file into a subfolder on datastore1 called backups. The file is 2,097,152.00KB and it took 3 minutes and 42 seconds to copy. That equates to about 9MB/sec. Which to me, seems terribly slow.

As a side test, using VMWare Server 2.0 which is installed on my Dell laptop, I created a 2GB file for a VM and copied it under Windows Vista using Windows explorer in 1 minute and 43 seconds. So, that is 2x as fast on a slow laptop drive.

As far as I can tell, the actual performance of the virtual machines themselves is just fine. Nobody is complaining, they load fast and seem to be working great. It's just painful to make an image and then copy it from one folder to another to setup another machine.

Any suggestions on how to improve these speeds, or is this just the way it is?



patparks1 23 posts since

Jun 16, 2009 **1. Re: Slow file system performance on ESXi 3.5 using HP DL360G5 all local disk** Jun 19, 2009 3:32 PM

It's possible after some more reading that my problem could be the LACK of a BBWC (Battery Backup Write-Back Cache) module on the

SmartArray controller which is causing this performance issue:

I'm going to have to physically get to the data center and check it out to see if it has the battery or not. It's the SmartArray P400i controller. It seems to be MFG#383280-B21

If anybody else has ideas, please throw 'em my way.



J1mbo 565 posts since

May 20, 2009 **2. Re: Slow file system performance on ESXi 3.5 using HP DL360G5 all local disk** Jun 20, 2009 12:12 AM

in response to: patparks1

You should be able to check the BBWC status from Insight Manager or possibly the VMware console itself, if you're using the HP embedded version that is.

Slow file system performance on ESXi ...

Also on the "slow management network" thread it is noted that 10MB/s seems to be a programatic limit perhaps to preserve the performance of running machines? What is the guest performance?

Finally, unless there is specific reason for the seperate mirror, most likely the overall IOPS rate would be improved by moving all the disks into the RAID 5 volume, since contention would be set against more spindles.



[patparks1](#) 23 posts since

Jun 16, 2009 3. Re: **Slow file system performance on ESXi 3.5 using HP DL360G5 all local disk** Jun 22, 2009 10:06 AM

👤 in response to: [J1mbo](#) We aren't using the embedded version, instead 3 servers were purchased and we installed the ESXi 3.5 software onto the machine once in house. I physically looked at the machine this morning, and it does not have the BBWC module installed. I can see where it plugs in...but it's not there...so I plan to purchase 3 of these and get them installed.

As far as the mirror goes, the original reasoning was that I don't usually install an OS on a RAID5 array due to parity calcs on swap space and such. Our decision originally was to get as much fault tolerance as we could for the base OS and the base virtual machines.....and then have as much physical space as possible for secondary drives that would be as fast as possible. So, we were going to stripe 4 of the drives and accept the fact that a loss of any would result in total data loss since speed was very imporant for the machines in question. In the end, however, other decisions were made to have fault tolerance on the secondary datastore, but we didn't want to completely start over, so I went with a RAID5 array for the second datastore.

I can understand the limiting of throughput on the management interface to protect the VM's....but that wouldn't impact the speed of file copy from 1 folder within datastore1 to another folder within datastore1 would it (either by using the datastore browser, or by doing cp on the VMDK files themselves at the unsupported console)? Our actual guest performance is just fine...at least I have noticed any issues and nobody has complained. However, when I try to copy a 20GB image file, to bring up another machine, it takes like 40 minutes and that just doesn't seem right to me.