



Is there a template to check the SnapDrive configuration and troubleshoot common disk enumeration issues?

https://kb-stage.netapp.com/Legacy/SnapManager_Suite/Is_there_a_template_to_check_the_SnapDri...

Updated: Wed, 24 Jun 2026 07:13:58 GMT

Applies to

Snapdrive

Answer

In order for SnapDrive to successfully enumerate the LUNs and properly communicate with the storage system for all other required tasks, perform the following steps on all the involved Windows servers and storage systems:

1. From each server, ping the storage system(s) by name (if ping is not working at any given time, resolve that first). When testing the name resolution, always ensure that reverse lookup also works (use the nslookup command for the name and then for the IP address).

2. Set up all the LUNs to be connected by the storage system hostname (that is when you connect or create the LUNs).
3. Set up the SnapDrive TCP protocol setting:
 1. Set the default protocol for any storage system that is not specifically configured.
 2. Optionally, set the storage system hostname by HOSTNAME (and not by the IP address) with a valid password for a storage system local user that is a member of the local storage system administrators group.

Run the following command to list the local users with Administrator privileges:

```
useradmin user list myfileruser1 -g administrators
```

If you are using RPC, there is no need to use the password. Ensure the SnapDrive service is running with the right credentials. To check that, click Start -> Run and then type services.msc. Under SnapDrive, click properties and then log on tab.

Ensure it is either a domain user or a local user as below:

```
mydomainuser1 or mylocalservermylocaluser1
```

If you use the pass-through authentication (local user on the server and storage system), ensure the storage system local user has the same name and password as the Windows server local user.

Note: Do not use a local user for the cluster service in a Windows cluster environment.

If you have a clustered SnapDrive configuration, use a domain account to run the cluster service, and all the nodes of the cluster must be in the same domain. However, the storage system can be in a different domain or workgroup.

If you use a domain user and want to check whether that user is also an admin on the storage system, run the following command:

```
useradmin domainuser list mydomainuser1 -g administrators
```

Note: When not including the SnapDrive related user to the local storage system's Administrators group, ensure to use the useradmin group list for the group name to which the user belongs. In the output, check the 'Allowed Capabilities' associated with the group; they should at least include: `api-*,login-http-admin`

3. When performing Step a and/or b above, ensure to use an account that is an admin on the storage system(s) and on the Windows server(s) as mentioned in the above note. Ensure that if you use RPC as the COMM protocol with the storage system, CIFS is started, set up, and working on the storage system. This is because CIFS opens the port that RPC requires. Check that by typing `cifs testdc` and/or `cifs domaininfo` on the storage system. In case of CIFS issues, type `cifs resetdc`. Make sure NetBIOS over TCP/IP option is either disabled on both DC and the storage system or enabled on both.
4. Check the SnapDrive service to see if it is running with a user which is a local server, remote server, and the storage system admin.
Perform these steps on all the involved servers.
5. Move to Step 6 if the Windows server is a member of the Windows cluster. Move to Step 10 if the server having the issue is not in a Windows cluster configuration.
6. Ping all the cluster nodes by name from each node.

7. Ping the storage system by name from each node. Perform an `nslookup` by the storage system IP address to make sure reverse lookup also works, as mentioned in Step 1.
8. Ensure the same version of SnapDrive is installed and working on all the cluster nodes and no firewall exists between the cluster nodes.
On the clusters, it is important that the cluster resource Network Name is (reverse) resolvable and that the resource is not in a failed state within Failover cluster management snapin. Within this snapin, you might encounter an error similar to the following:
`Cluster network name resource 'Cluster Name' cannot be brought online. The computer object associated with the resource could not be updated in domain 'escwin2k8dom.escalation.ams' for the following reason: Unable to obtain the Primary Cluster Name Identity token.`
This can occur, for example, if the Virtual Machine(s) (VM) are restored from a snapshot.
9. Run Steps 2, 3 and 4 from all cluster nodes. Then proceed to Step 10
10. Ensure that all services on which SnapDrive depends are running and the Workstation and WMI services are running.
11. If the storage system has multiple network interfaces but not accessible to the Windows server, set the SnapDrive `preferredfilerip` address option. This is for deciding the interface to choose for the traffic from the storage system back to the Windows server. It will not resolve any name resolution issues.
12. For enumeration issues with RDM LUNs/ VMDK disks, check the following:
 1. If the enumeration issues are with RDM LUNs, check that the user credentials that SnapDrive uses for contacting the Virtual Center (VC) has the minimum required permissions; this is explained in the section 'Minimum vCenter permissions required for SnapDrive operations' on page 49 of the SnapDrive 6.3 for Windows [Installation and Administration Guide](#). VC user should be in the format of `domain\user1`. If the user is directly giving the ESX credentials, the ESX local user should be an admin on the ESX box and should be in the format of `root` or `myadmin1`. The Data ONTAP WinDC file `SnapDrive_RegistryInfo.log` contains the user configuration.
 2. If the enumeration issues are with the VMDK disks, check the [Interoperability Support Matrix Tool](#) to check which version of VSC is installed on the VMware side (2.0.1 is the minimum supported version, but for example, SnapDrive 6.4 does not support VSC 4 and higher, SDW 6.4.2 is the first version to support ESX 5.1).
13. Disk enumeration is slow:
In this case, there might be known issues causing slowness or there could be configuration issues causing it.
 1. SnapDrive using HTTP with a domain user, in cluster environments is known to be slow, as per the [Installation and Administration Guide](#). Use a root account or other storage system's local account with local admin rights on the storage system instead.
 2. ZAPI calls through RPC might travel through multiple Windows client network interfaces. All interfaces need to be enabled for CIFS client purposes when using RPC. All interfaces need to be configured the same way as the storage system's interfaces and the Ethernet switches in between.
Examples: Both the storage system and the Windows client need to have the same jumbo frames setting; both enabled or both disabled. If jumbo frames are enabled on the storage system's NIC they intend to use, make sure the jumbo frames are configured on the client system too.

A quick test to check whether the Windows SnapDrive client is able to communicate with the storage

system's interface is to type the following command for the storage system's IP address (the IP address that DNS returns upon running `nslookup` of the storage system hostname):

```
ping -l8500 -f Filer_ip_address (Use -l with a number higher than the normal 4088 buffer size; use -f for 'dont fragment')
```

Finally, test the CIFS share connection with: `net use \fileripc$`

3. Some users report that they have resolved the issue by changing the NIC teaming settings on the network interface used for management and DNS (between the Windows server and the storage system, these interfaces were not used for iSCSI). Changing the policy from 'Auto / Load Balancing' to 'Fault Tolerance' makes the slow disk enumeration issue disappear. Basically packets sent from SnapDrive to the storage system take a longer time, which is also observed within a packet trace. Also, ensure that .Net Framework 3.5 SP1 is installed rather than version 4.
 4. One reason for slowness can be unsupported HBA Driver/Firmware installed on the Windows server(s) involved in the enumeration. Check [Interoperability Support Matrix Tool](#) for data taken from OneCollect to make sure the user environment is compliant with the [Interoperability Support Matrix Tool](#).
 5. SnapDrive 6.4 is supposed to be faster when it comes to enumerating LUNs in Data ONTAP 7 or 8.x in 7-Mode. SnapDrive 6.4 in 8.1 clustered ONTAP will be slower.
 6. If snapshot creation is slow when making a snapshot of a VMDK disk but not with a snapshot of a normal LUN, then you need to collect the VSC logs in verbose mode too. This is because the bottleneck is likely to be in the VMware/VSC area. So, from the SnapDrive end, check the VDISKAPI.logs if all calls from SnapDrive are sent to SMVI for checking snapshots of the VMDK disks or even enumerating VMDKs. On the other end (the SMVI end), check the SMVI logs (which in reality are called `server.log` logs.)
 7. NIC teaming on the client interface that is used on the Windows host to transfer iSCSI blocks is not supported by Microsoft and it can cause performance issues. In order to verify whether NIC teaming impacts performances, change the policy to failover only.
 8. [BUG 552607](#) shows excessive use of long running ZAPIs by OnCommand applications which might result in ZAPI timeouts and http connection limits to exceed. The issue is resolved starting with Data ONTAP 8.1.1P1. The symptom is that ZAPI calls do not reply until hours or days later. Normally, this causes snap mounts to fail, rather than disk enumeration slowness. One can easily spot this in the SnapManager and SnapDrive debug logs.
14. If the issue persists, log a case with NGS.

Additional Information

additionalInformation_text