

Drobo How-To Guide

Disk Pack Migration on Drobo 8-bay iSCSI SAN Storage



BeyondRAID™ technology in Drobo® enables users to very quickly migrate the disks in one Drobo (called a disk pack) to a Drobo of the same model or, in certain cases, to a larger or newer Drobo. Disk pack migration may be required when you need to replace a system or when you upgrade from DroboElite to Drobo model B800i, as described in this document. Following the steps in this guide will aid you through the process of going from:

- DroboElite to Drobo model B800i
- DroboElite to DroboElite
- Drobo model B800i to Drobo model B800i

Note that in most cases migrating drives is a one-way operation for supported migrations from one Drobo model to another. Once a disk pack is migrated, you will no longer be able to access the data it contained with the older Drobo.

Topics

- Preparing for migration and moving the disk pack
- Mounting datastores in VMware ESX/ESXi Server

What You Will Need

- Drobo 8-bay SAN storage systems
- The most recent firmware installed on both Drobos
- Drobo Dashboard management software version 2.0 or higher



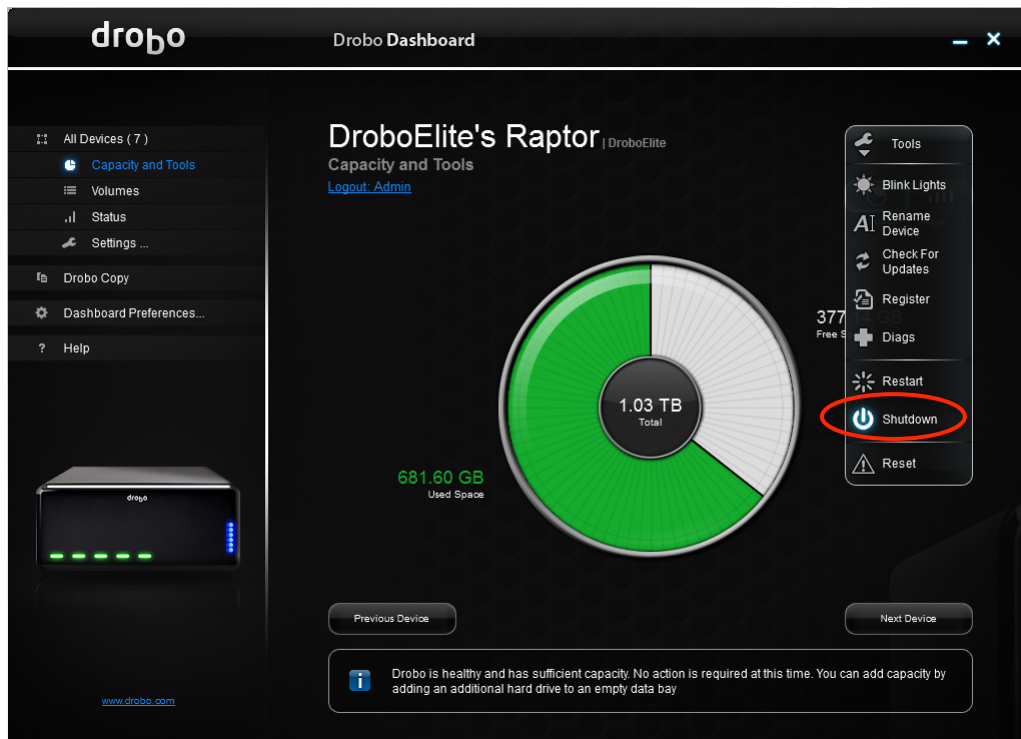


Preparing for Migration and Moving the Disk Pack

STEP 1

Shut down the existing Drobo 8-bay SAN system. In Drobo Dashboard 2.0 select **Shutdown** from the **Tools** dropdown menu (circled in red). (In Drobo Dashboard version 1.x.x, go to the Advanced > Tools tab.)

Ensure that all lights on the Drobo are off before performing next step



STEP 2

Now, simply remove all of the drives from the source Drobo. Since Drobo is slot independent, you don't need to keep track of which drives came out of which slots when you re-insert them into the destination Drobo.



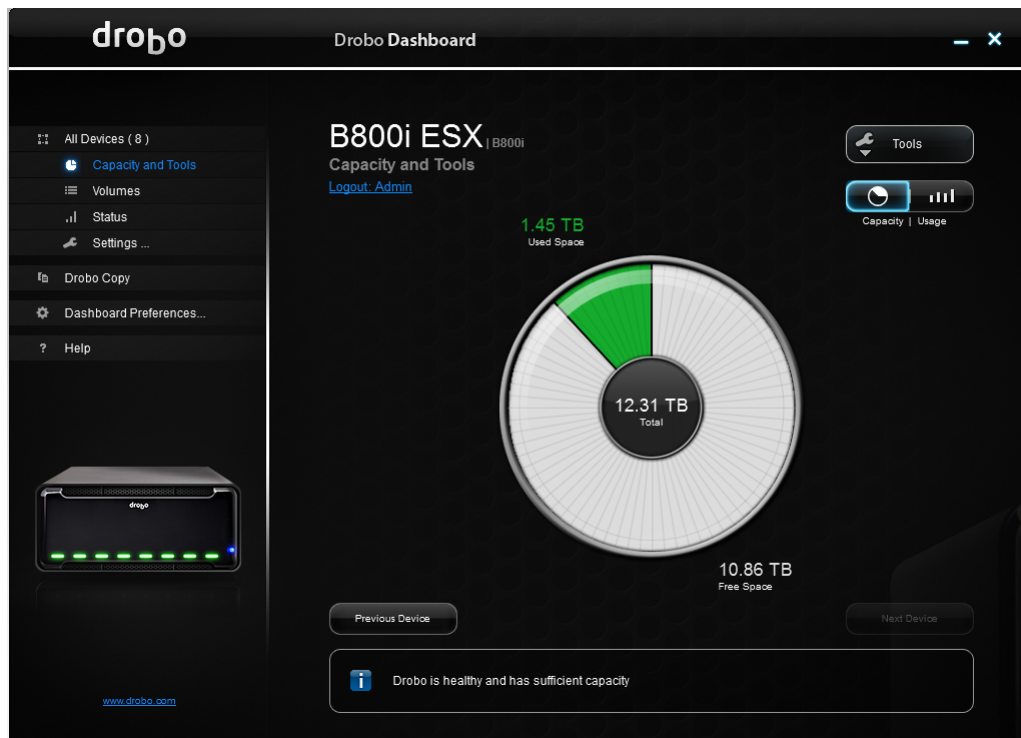


STEP 3

Insert the disks from the DroboElite into the Drobo B800i; the drive order doesn't matter Plug all cables into the destination Drobo and power up.

Drobo Dashboard automatically discovers the system.

NOTE: The IP address and other device settings such as the name are retained with the disk pack, so the destination system will appear with the same name and IP address as it had in the source Drobo.



Mounting Datastores in VMware ESX/ESXi Server

When a LUN is presented to an ESX/ESXi host, it gets a signature in the form of UUID. Depending on the state of the VMFS datastore residing on a LUN that has been detected, it can be mounted using one of the following options:

- Resignature the VMFS datastore and mount it on the selected host. Subsequent VMFS refresh is required on other hosts to see this newly resigned datastore.
- Keep the existing signature of the VMFS datastore and mount it on the selected host. This is also known as "force mounting" a VMFS datastore.

The steps in this section describe the second option, keeping signature(s) of the VMFS datastore(s).

Drobo How-To Guide

Disk Pack Migration on Drobo 8-bay iSCSI SAN Storage



STEP 1

Open a terminal session to the ESX/ESXi Server host and log in as root. If for any reason you are unable to connect to an ESX/ESXi Server host via an SSH session, refer to the VMware KB article [Unable to connect to an ESX host using Secure Shell \(SSH\) \(1003807\)](#).

```
root@R410:~  
login as: root  
root@172.32.16.80's password:  
Last login: Wed May 4 08:35:58 2011 from 172.32.16.130  
[root@R410 ~]# vmware -v  
VMware ESX 4.1.0 build-260247  
[root@R410 ~]#
```

STEP 2

Run: `esxcfg-mpath -b |grep -i dead`

This will return any “dead path” that may have resulted from the disk pack migration due to a change in Universal Unique ID (UUID). The ESX host uses a unique identifier to reference the volume. The name you specify when you create a VMFS volume is a user-defined device name that is a symbolic link to the UUID of the VMFS volume. This is done to solve the problem of changing the device name. When you change the volume name you are changing only the user-defined device name and not the volume UUID.

```
[root@R200 log]# esxcfg-mpath -b|grep -i dead  
vmhba33:C1:T10:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T10:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T14:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T14:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T7:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T7:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T28:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T28:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T13:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T13:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T9:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T9:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T23:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T23:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T12:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T12:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T8:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T8:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T11:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T11:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C1:T15:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable  
vmhba33:C0:T15:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable
```

STEP 3

To perform ESX/ESXi host rescan, run: `esxcfg-rescan vmhba33`, where xx is the unique ID, in this example, 33. If any syntax is returned, it means that there is a dead path and you must reboot the ESX/ESXi host. If reboot is not required proceed to step 6

```
[root@R200 log]# esxcfg-rescan vmhba33  
Dead path vmhba33:C1:T12:L0 for device naa.6001a620000742303935303830333732 not removed.  
Device is in use by worlds:  
World # of Handles Name  
4096 1 console
```



STEP 4

To verify the all paths-down (APD) state due to a dead path, run: `grep -ir "APD" vmkernel`

The APD state is a condition in which no working path exists to a storage device (LUN). In this case, it is a result of migrating the disk pack from DroboElite to Drobo B800i. Proceed to rebooting the ESX/ESXi host; ensure that no VM is active before rebooting.

```
[root@R200 log]# grep -ir "APD" vmkernel
May  4 13:04:50 R200 vmkernel: 40:21:59:00.223 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:05:55 R200 vmkernel: 40:22:00:05.210 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:07:00 R200 vmkernel: 40:22:01:10.210 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:08:05 R200 vmkernel: 40:22:02:15.220 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:09:10 R200 vmkernel: 40:22:03:20.230 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:10:15 R200 vmkernel: 40:22:04:25.238 cpu1:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
[root@R200 log]# esxcfg-mpath -b|grep -i dead
vmhba33:C1:T12:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable
```

STEP 5

To verify that you can keep the existing signature of the VMFS datastore and mount it on the selected host, run: `esxcfg-volume -l`.

This is also known as force mounting a VMFS datastore.

The Can mount value must be Yes before you can proceed to the next step. Record the UUID portion of the output. In this example, one of the UUIDs is 4cbf6c94-57b4c7f3-42bf-002219d4fc7f.

```
[root@R200 log]# grep -ir "APD" vmkernel
May  4 13:04:50 R200 vmkernel: 40:21:59:00.223 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:05:55 R200 vmkernel: 40:22:00:05.210 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:07:00 R200 vmkernel: 40:22:01:10.210 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:08:05 R200 vmkernel: 40:22:02:15.220 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:09:10 R200 vmkernel: 40:22:03:20.230 cpu0:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
May  4 13:10:15 R200 vmkernel: 40:22:04:25.238 cpu1:4210)WARNING: NMP: nmpDeviceAttemptFailover: Retry world failover device "naa
.6001a620000742303935303830333732" - failed to issue command due to Not found (APD), try again...
[root@R200 log]# esxcfg-mpath -b|grep -i dead
vmhba33:C1:T12:L0 LUN:0 state:dead iscsi Adapter: Unavailable Target: Unavailable
```

STEP 6

To persist the volume mount run: `esxcfg-volume -M <UUID>`

Where the <UUID> is the value recorded in Step 6.

NOTE: If you do not wish the volume mount to persist after reboot, use the `-m` switch instead.

Repeat this command syntax as many times as required.

```
[root@R410 ~]# esxcfg-volume -M 4cbf6c94-57b4c7f3-42bf-002219d4fc7f
Persistently mounting volume 4cbf6c94-57b4c7f3-42bf-002219d4fc7f
```

Drobo • 1705 Wyatt Drive, Santa Clara, CA 95054 • www.drobo.com • 1.866.97.DROBO

Copyright 2011 Data Robotics, Inc. Data Robotics, Drobo, DroboElite, DroboPro, BeyondRAID, and Smart Volumes are trademarks of Data Robotics, Inc., which may be registered in some jurisdictions. All other trademarks used are owned by their respective owners. All rights reserved. Specifications subject to change without notice. • HT-0028-00 • May 2011