



DM-Multipath Guide

Version 8.2



SBAAdmin and DM-Multipath Guide

The purpose of this guide is to provide the steps necessary to use SBAAdmin in an environment where SAN storage is used in conjunction with device-mapper multipath devices. When a system is using dm-multipath devices, there are several considerations that need to be addressed. The device naming, tools available, and device modules loaded become critical issues in discovery of devices and recreation of those devices upon restore. The information in this guide should provide the necessary information to create backups from systems using dm-multipath devices or for users who want to migrate a system backup to multipath devices. This guide is not intended to assist users in the initial setup of dm-multipath devices. For information on initially setting up dm-multipath devices, contact your Linux OS support vendors directly.

Requirements for DM-Multipath support

Software requirements

Support for dm-multipath devices has been available to Linux users for years, however, how the devices are created and supported by distributions has changed greatly over time as the technology matures. Below is the minimal level of the tools and file sets that SBAAdmin has tested and will support.

dm-multipath-tools version 0.4.5 (also known as device-mapper-multipath)
device-mapper version 1.02
udev version 039-10

Device detection requirements

Device naming

If you plan on creating backups from a system using dm-multipath devices, the SBAAdmin software must be able to recognize the devices by name prior to creating the backup. Multipath devices can be named several different ways such as by the World Wide Identifier (WWID), dm-[0-9], mpath[a-z], and mpath[0-9]. How the devices are named depends on three factors: The “**user_friendly_names**” setting in **/etc/multipath.conf**, aliases setup in **/etc/multipath.conf** and **UDEV rules**.

SBAAdmin requires that the “**user_friendly_names**” setting in the **/etc/multipath.conf** configuration file be set to “**yes**”. This setting ensures that the disk will NOT be named using the unique WWID, but will use a predictable name such as mpath[a-z] or mpath[0-9]. Some distributions supply a copy of the **/etc/multipath.conf** file already in place while others need to be created using a template supplied by the distribution. Refer to the documentation specific to your distribution for details.

If you need to create the **/etc/multipath.conf** file, at a minimum it should include the following entry:

```
defaults {  
    user_friendly_names yes  
}
```

The changes are not applied until you update the multipath maps, or until the multipathd daemon is restarted, such as at system reboot.

IMPORTANT NOTE: If the root filesystem is on a multipath device and you updated the “**user_friendly_names**” setting, it is recommended (and in some cases required) to recreate the initrd so that the updated multipath.conf file is used during the boot process. Please refer to your distributions documentation for details on recreating the initrd (initramfs).

1. Recreate the initrd with the command supplied by your distribution (here are two examples)
mkinitrd [see man page for additional options]
dracut [see man page for additional options]
2. Reboot the system
reboot

If the disks are still named with the WWID after completing these steps, the “**user_friendly_names**” directive is not properly being applied and you should contact support for your particular Linux distribution.

Device location

Another requirement is that a soft link to the device nodes should be created in the **/dev/mpath/** directory. Some Linux distributions already create these links. If your Linux distribution does not, SBAAdmin has a sample UDEV rule that you may apply. It is designed to create the proper links and allow SBAAdmin to discovery the devices. To apply the rule to your system:

1. Copy the rule into place
cp /opt/storix/config/dm_mpath_udev_rule /etc/udev/rules.d/99-storixmpath.rules
2. Trigger the UDEV rule with one of the following commands (depending on your version of UDEV)
udevadm trigger
udevtrigger

Device detection

At this point, your multipath devices should have the path **/dev/mpath/mpathX**. If your backups fail or appear to exclude the data on your dm-multipath devices, verify that the devices are listed in the **/dev/mpath/** directory.

Having the devices in the **/dev/mpath/** directory is required for device discovery. However, there are other factors that can prevent device detection. There are six (6) different device naming schemes that SBAAdmin has been tested against. Different Linux distributions and even different versions of the same distribution name these devices differently. The device names differ by the disk designation and the prefix used to note the partition.

For instance, the first detected SCSI or SATA disk in a Linux system will usually appear to the system as **/dev/sda**. There is not a partition prefix, so the first partition will be named **/dev/sda1**.

For example, the first partition on the first dm-multipath device it could be called one of the following:

```
/dev/mpath/mpath0p1  
/dev/mpath/mpath0-part1  
/dev/mpath/mpathap1  
/dev/mpath/mpathapart1  
/dev/mpath/mpathap1  
/dev/mpath/mpathapart1
```

The naming schemes are as follows:

```
/dev/mpath/mpath[0-9]p[1-15]      /dev/mpath/mpath[a-z]p[1-15]      /dev/mpath/mpath[a-z][1-15]  
/dev/mpath/mpath[0-9]-part[1-15]      /dev/mpath/mpath[a-z]-part[1-15]      /dev/mpath/mpath[a-z]_part[1-15]
```

If you are using an alias line in your **/etc/multipath.conf** file, you could also find that the alias is used instead of the word mpath. For example: **/dev/mpath/mydisk0-part1**. There are limitless numbers of combinations you could have for naming your devices. SBAAdmin has a configuration file that is used as a reference file for device naming schemes. This file is **/opt/storix/config/device_info**. Currently we have entries for the two of the most common naming conventions uncommented in this list. In most cases you will not need to update this file. However, you may need to edit this file to match the naming scheme in your environment. **Make sure to comment out any naming scheme that you are not using to prevent naming conflicts.**

So that changes to the device_info file are not lost after software updates, make a copy of the device_info file to the config directory inside your storix data directory. Once you have made a copy, edit the file located in the config directory. By default this would be **/storix/config on the client that is using dm-multipath devices.**

1. Make a copy of the device_info file before editing
cp /opt/storix/config/device_info /storix/config/device_info
2. Edit the device_info file and uncomment the mpath naming scheme that matches your environment. Be sure to comment out all others.

```
# vi /storix/config/device_info

#####
# Multipath (device-mapper)
#####

# The following defines up 16 devices, 15 partitions per device
# Because different linux distributions use varying udev rules
# uncomment the device names that match your environment
disk:mpath/mpath:0-15:p:1-15:Device-Mapper Multipath device
disk:mpath/mpath:a-p:_part:1-15:Device-Mapper Multipath device
#disk:mpath/mpath:a-p:-part:1-15:Device-Mapper Multipath device
#disk:mpath/mpath:a-p:p:1-15:Device-Mapper Multipath device
#disk:mpath/mpath:0-15:-part:1-15:Device-Mapper Multipath device
#disk:mpath/mpath:a-p::1-15:Device-Mapper Multipath device
```

If you are using an alias, you will need to create an entry for your own naming scheme. As in the example above for **/dev/mpath/mydisk0-part1**, the entry would be:

```
disk:mpath/mydisk:0-15:-part:1-15:Device-Mapper Multipath device
```

Please note that alias names **will not be** preserved during system recovery. All multipath devices will be renamed to the default mpath naming scheme used by your distribution.

Entries in /etc/fstab

Different Linux distributions handle dm-multipath devices in different ways. One of the areas that can cause problems with the SBAAdmin software is the mount entries in the **/etc/fstab** file. Some distributions use an underlying path device as the block special device used for mounting the filesystem. This causes problems

when the software attempts to query the device for size and geometry. The block special device listed in the **/etc/fstab** file should be the dm-multipath device instead of an underlying path device. Verify that if your **/etc/fstab** file is using the **/dev/disk/by-uuid/*** naming convention, that you change the entries to the dm-multipath device name or the **/dev/disk/by-name/*** naming convention. This will prevent problems when attempting to query the devices during system backups. **Note, after changing fstab entries you may need to recreate your initrd using the distributions mkinitrd command.**

Migrating to DM-Multipath Devices


Creating boot media

With most devices on Linux, there is a single module that needs to be loaded for the device to become visible to the operating system. In the case of dm-multipath devices, there could be two. One module for the Host Bus Adapter or HBA and one for the device handler module. The device handler module is specific to the type of SAN hardware you are using. If you know in advance that you are migrating to different hardware, you can preselect the module for the HBA and the device handler to be loaded upon booting from the boot media. For more information on pre-selecting modules, please refer to the **“Customizing Boot Media”** section in the *SBAAdmin Linux System Recovery Guide*.

Discovering dm-multipath devices from the recovery media

If you are migrating from a non-multipath system, by default, the necessary modules required for the operating system to discover the dm-multipath devices will not be available. Unless you have preselected the modules to be loaded at boot-time, you will need to load additional drivers to make the dm-multipath devices available to the operating system. In some cases, you may need to select both the module for the HBA as well as the device handler for your SAN.

From the **System Installation and Maintenance Main Menu**, select **System Recovery Utilities**, select **Load Additional SCSI/FC Adapter Modules** to display the following options:



After loading the module for the HBA, **only the underlying paths are listed and no dm-multipath devices were created**. At this point it is necessary to also load the device handler for the SAN.

```

Storix System Backup Administrator for Linux

Press ENTER to load or unload a module from the list below.
"=>" indicates a module currently loaded. Press F3/ESC when done.

Module      Description
pata_marvell  Marvell PATA driver
pata_netce
pata_rz100
pata_sil68
-----
lpfc
qla2xxx
qla1280
qllogicfc    QLogic ISP2100 SCSI-FCP host adapter
qla2xxx      QLogic ISP2xxx host adapter family support
qla4xxx      QLogic ISP4XXX host adapter family support
-----
Other RAID/Disk Controllers -----
cpqarray     Compaq PCI Hot Plug Controller
[More]
F1=Help
  
```

Now the list of devices created also shows the dm-multipath devices that were created and available for use during the restore.

```

Storix System Backup Administrator for Linux

Press ENTER to load or unload a module from the list below.
"=>" indicates a module currently loaded. Press F3/ESC when done.

Module      Description
qla4xxx      QLogic ISP4XXX host adapter family support
-----
cpqarray     EMC CX/AX/FC-family multipath
DAC960
i2o_config   disks: sda sdb sdc sdd sde sdf mpath/mpath0
i2o_block    cdroms: scd0
Press ENTER to continue ...
-----
dm-emc       EMC CX/AX/FC-family multipath
dm-hp-sw     HP StorageWorks and FSC FibreCat multipath
dm-rdac      DM Multipath LSI/Engenio RDAC support
-----
USB Storage -----
=> usb-storage  USB Mass Storage support
F1=Help
  
```

Troubleshooting

Because there are several device naming schemes possible when using dm-multipath devices (see the section titled “**Device naming**” above), it may be necessary to create a custom version of the **device_info** file that is used as a reference for device naming schemes. If the restore fails because your partitions are not discovered during install, reboot from the boot media and customize the **/opt/storix/config/device_info** file.

Here are the steps to customize the **/opt/storix/config/device_info** file from the boot media:

- 1) Power cycle the system.
- 2) From the **System Installation and Maintenance Main Menu**, select **System Recovery Utilities**, select **Start a Maintenance Shell** to access a terminal. The vi editor is available for making changes to the **/opt/storix/config/device_info** file.
- 3) Update the **device_info** file to match the naming scheme for your environment. Again, for information on dm-multipath device names, see the section titled “**Device naming**” previously referenced in this document.
- 4) Copy the updated **device_info** file to the local *SBAAdmin data directory*. The path is already stored in the variable `$STXPATH`. You can reference that path in the maintenance shell.

```
# mkdir $STXPATH/config
# cp /opt/storix/config/device_info $STXPATH/config/device_info
```

The custom version of `device_info` will be copied to the restored system for further use.

Restoring a system using dm-multipath devices

Restoring a system using dm-multipath devices is the same as restoring a normal Linux system. For further detail please refer to the ***SBAAdmin Linux System Recovery Guide***.