

A Linux Backup Strategy – Version 1.0

Introduction

There are many applications available to do varying degrees of backups in a Linux environment. The following details the steps I take to insure an adequate backup of data on my production PC. It is hoped that the following content will help the reader avoid a situation where they are not able to adequately recover from either a major hardware and/or software failure on their PC.

Background

The following is a germination stemming from a couple of real data recovery disasters that I experienced over the last couple of years. Following a very grueling ordeal of trying to rebuild some critical data from scratch, I started to evolve the following strategy. Feel free to use whatever parts you feel are appropriate to your situation.

Disclaimer

The following details the backup strategy that is used for my current hardware and software configuration. I assume no liability for any damage that may occur on your system from using the attached.

Intended Audience

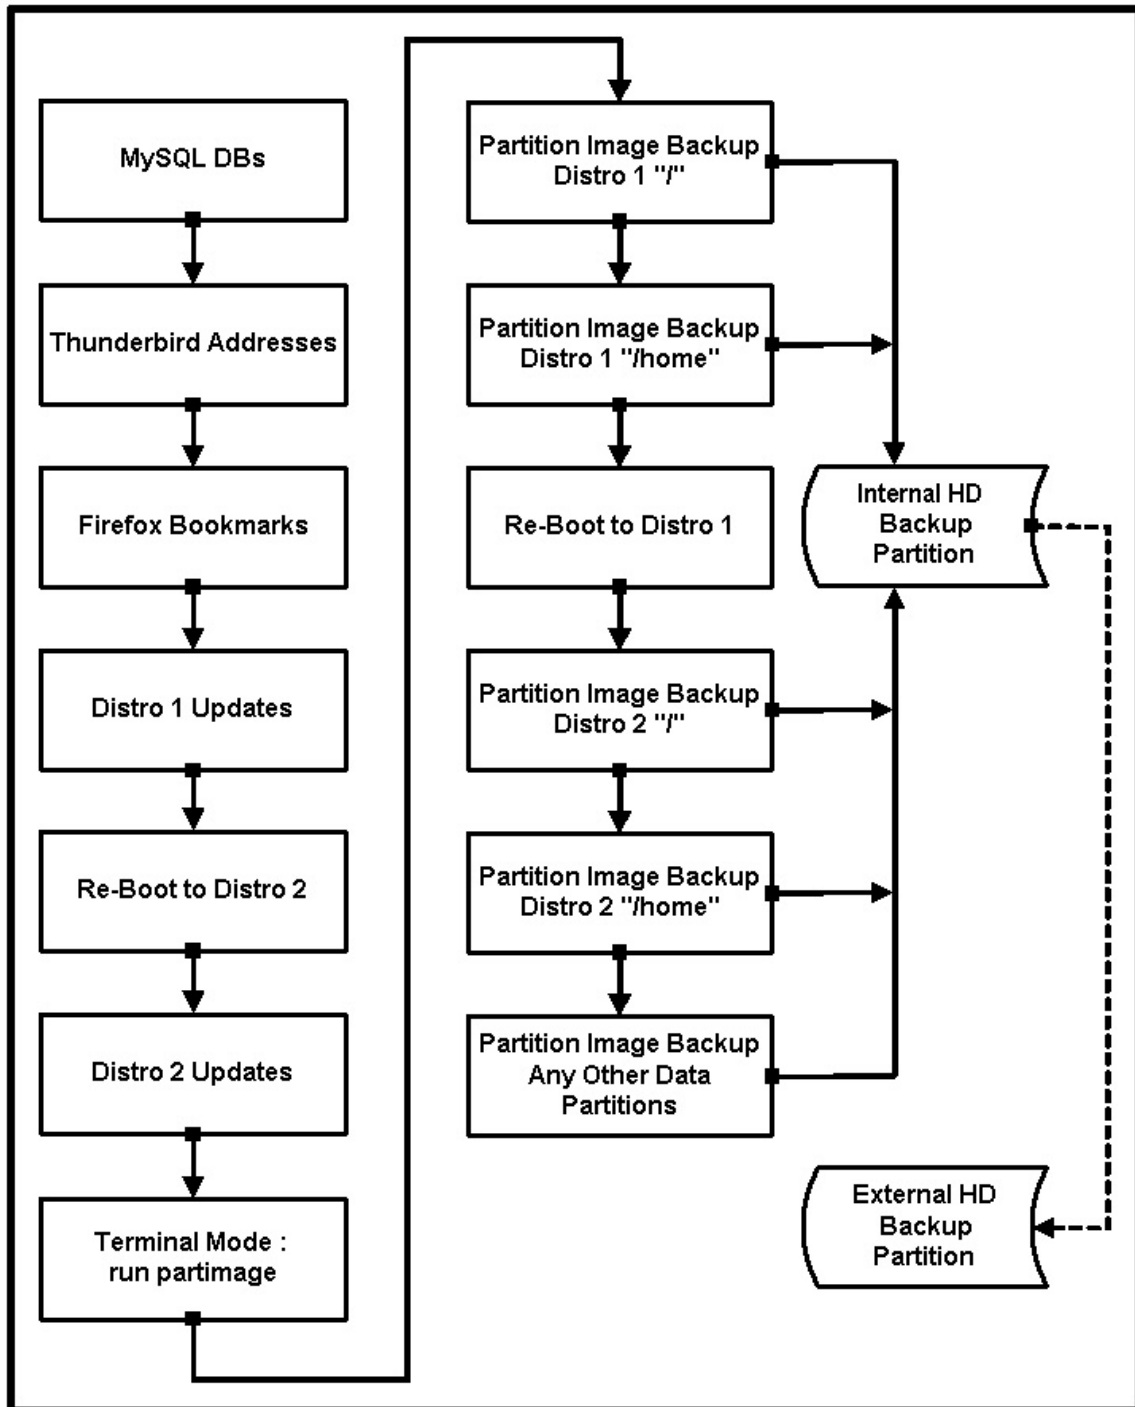
This paper is structured to users who are using one of more Linux operating systems. The Windows environment has an ample supply of commercial products available (for a cost) that will, for the most part, do what can be done in the Linux environment for no cost. The reader should be familiar with the basics of the Linux operating system and be comfortable doing a few things in a terminal mode. Parts of the following may have to be modified to accommodate the specific Linux distro you may be using as the frame of reference for a lot of the material is the debian derivative found in the Ubuntu-Kubuntu family of operating systems.

What Is Backed Up

On a weekly basis, the following are backed up:

- All active operating system partitions (3 Linux).
- The “/home” partition associated with each of the Linux operating systems.
- Multi-media partitions (if significant changes to the content occurred during the last week).
- The MySQL server databases that are in use on the local server.
- The address book from Thunderbird mail client.
- A Backup of the bookmarks in Firefox web browser.

The Backup Process Flow :



Timing

The full compliment of backups commences weekly on Saturday morning. Generally speaking, it takes about 1- 1.5 hours to complete the process.

This includes running the basic backups to a dedicated partition on the second internal hard drive as well as backing things up to an additional external hard drive. This provides a double redundancy for recovery purposes.

The operating system backups are run immediately after doing the weekly upgrades that are streamed from the repositories. This is the ONLY time that updates are run from the repositories. The rationale behind this approach is that the previous week's snapshot of the operating system partition, provides a recovery point if the current week's update doesn't fare well. Additionally, holding off on the updates until later in the week allows the developers to correct any glaring problems that may surface as the update is released.

Applications Required

1. Partimage installed across all Linux operating systems.
2. phpMyAdmin running to administer the MySQL databases on my production operating system.

Data Retention

The data stored on the internal hard drive partition backs up to either 1 week or the last event of backup. Data stored on the removeable hard drive is backed up to either two weeks or two events.

Backup Partition Structure

/media/LxBackup

 /addresses (used for the Thunderbird address book backup)

 /bookmarks (used for the Firefox bookmark backup)

 /ops-iso (stores the various iso images that may be burned to a boot cd etc.)

 /HD-BU (contains the partimage partition image files)

 /ServerBU (contains the MySQL and webserver backups)

 /Setup (various setup snippets etc.)

 /Documents (various categories of document folders and contents)

Backup Sequencing

Database Backup:

Each week a new folder is created. The folder is named to the date of the backup in the format of YYYY-MM-DD . For this use, the path is:
/media/LxBackup/Server Backups/MySQL/<YYYY-MM-DD> folder

A separate folder is created for each database. Within each of these folders, the database table structures is exported as a separate file. The contents of each table is exported as a separate file. Individual table files are used because of the extensive use of the InnoDB engine and foreign keys. In this environment, it is critical to load the tables in the correct order of referential hierarchy.

A backup copy of the phpMyAdmin, MySQL and the database specific to the LAMP server is also backed up. These backups are inclusive inasmuch as the structure and content are combined into one backup file for each database.

The database content are regarded as the most volatile of all of the data being backed up, That is the reason that they are sequenced first in this activity.

Address Book Backup:

An export of the address book contained in Thunderbird is made each week. The export is accessed via the following within Thunderbird:
Tools -> Address Book ->Tools -> Export

Bookmark Backup:

With version 3.x of Firefox, the user can now export a backup of the bookmarks vs an html extraction as was found in the earlier versions. The export is accessed (version 3.x) via:
Bookmarks ->Manage Bookmarks -> Export

Linux Operating System Update:

At the conclusion of backing up the bookmarks, the next activity is to run the updates for the operating system and applications for the week. The updates are run only a weekly basis. This approach provides some level of control and recovery if an update goes bad. Waiting until the end of the week also allows the updates to “age” a bit. If there are any problems with the initial launch of an update, the developers are generally quick to respond and re-issue the updates. Waiting until the later part of the week adds a little bit of additional confidence that the update is a “good-good” coming out of the box.

A Linux Backup Strategy – Version 1.0

Miscellaneous:

If any significant changes occur with respect to Documents, application iso files or setup snippets, these changes are copied over to the internal hard drive backup partition.

Partimage

The partimage application serves as the partition backup engine used in this backup strategy. It is easy to use and provides quick recovery to the state of the backup image used to restore. The application can be installed via the repository (either adept or synaptic or the equivalent in non-debian Linux distributions).

Preparation :

As discussed briefly in the above, it is strongly recommended that a partition dedicated to backups be used for storing the partition image files (and the other backup materials). Definitely DO NOT store the partition backup material on your operating system partitions.

Limitations:

You will NOT be able to run partimage against a partition that you are currently booted to. The same probably also applies to the currently booted “/home” partition.

Running partimage:

Access the terminal mode.

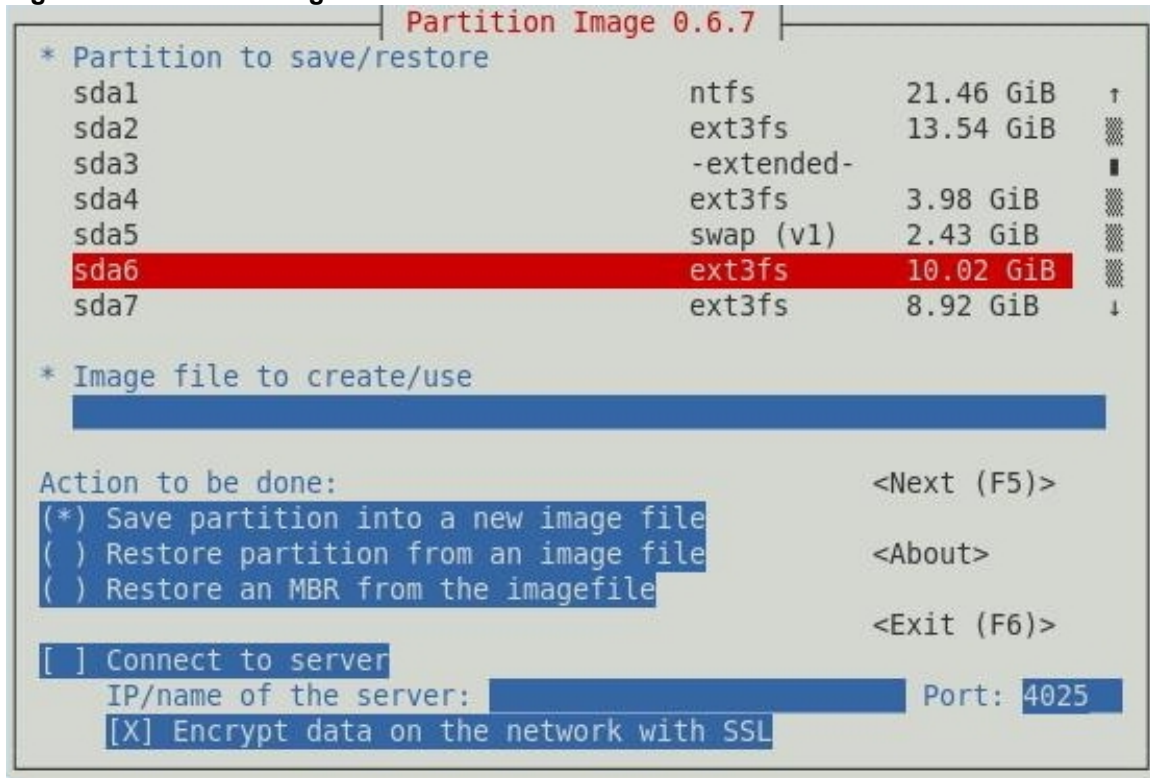
At the command prompt:

```
$sudo partimage
```

Enter your passord when prompted

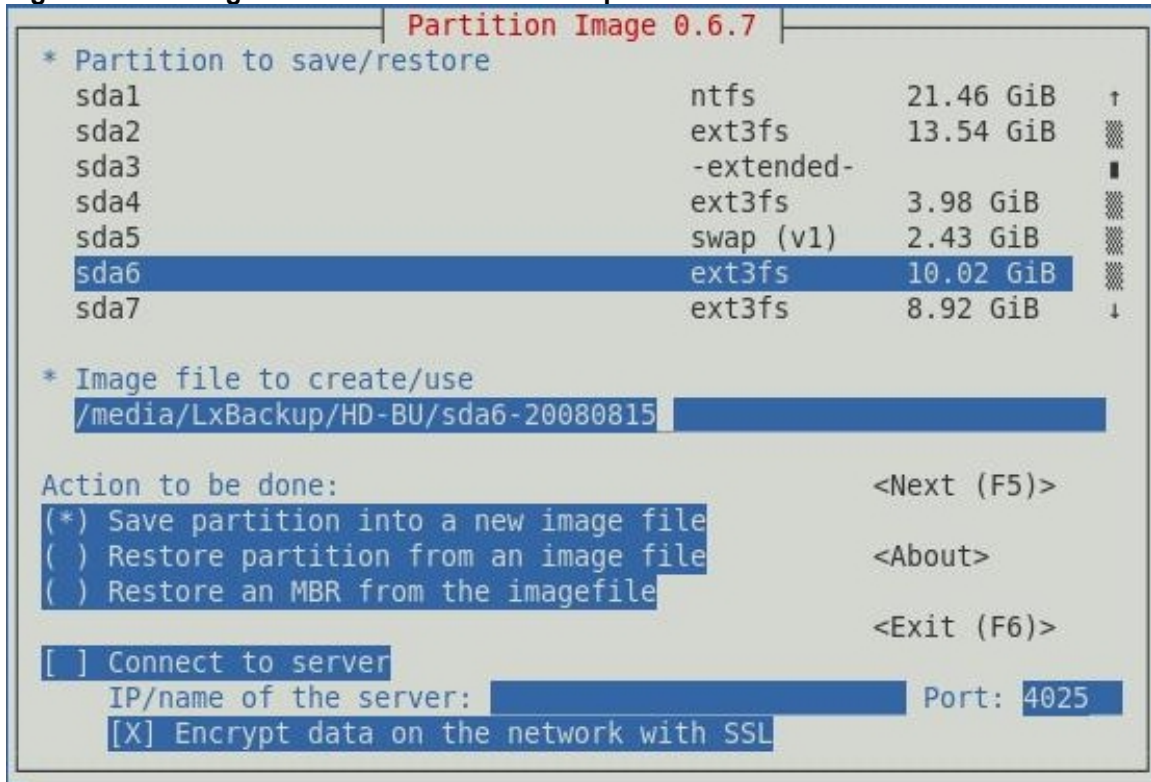
A Linux Backup Strategy – Version 1.0

Figure 1 : Initial Partimage Screen:



Using the up-down cursor keys, select the partition that is to be backed up. Click <enter> and click <tab> until the "Image file to create/use" is active.

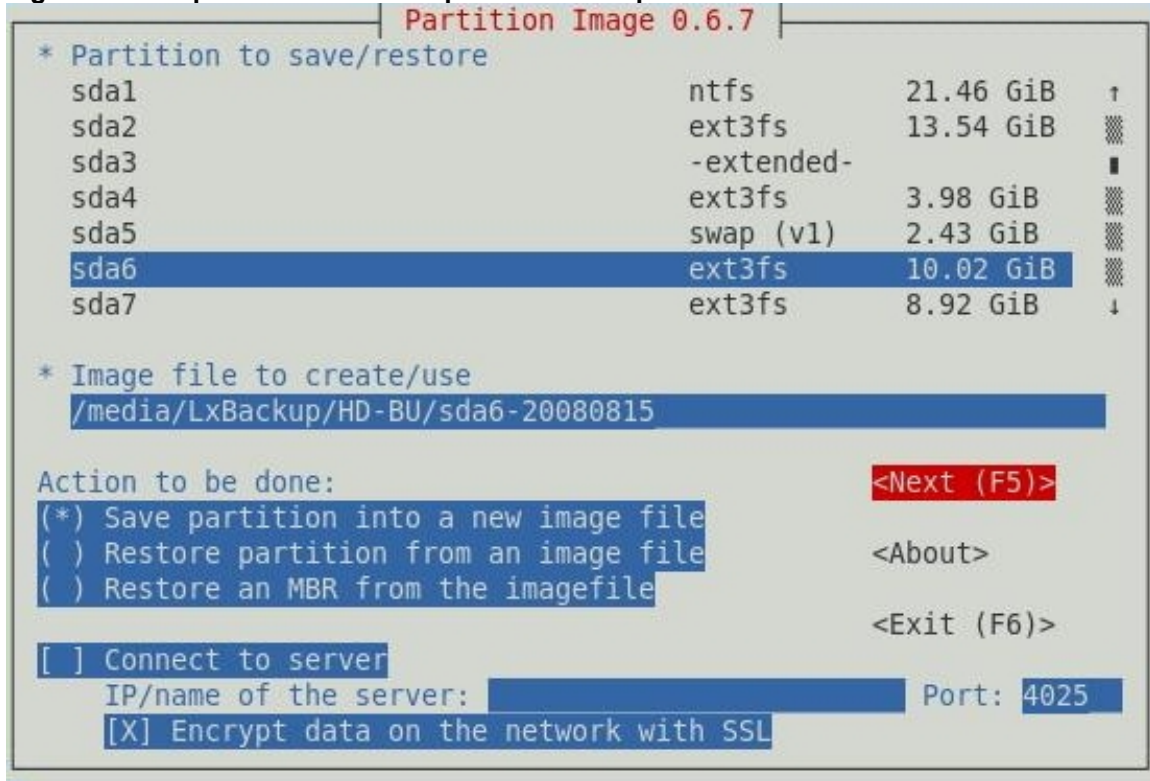
Figure 2 : Entering the destination of the backup file



Type in the complete path to the backup folder and the filename. In the example shown above, the path is “/media/LxBackup/HD-BU” and the filename is “sda6-20080815”.

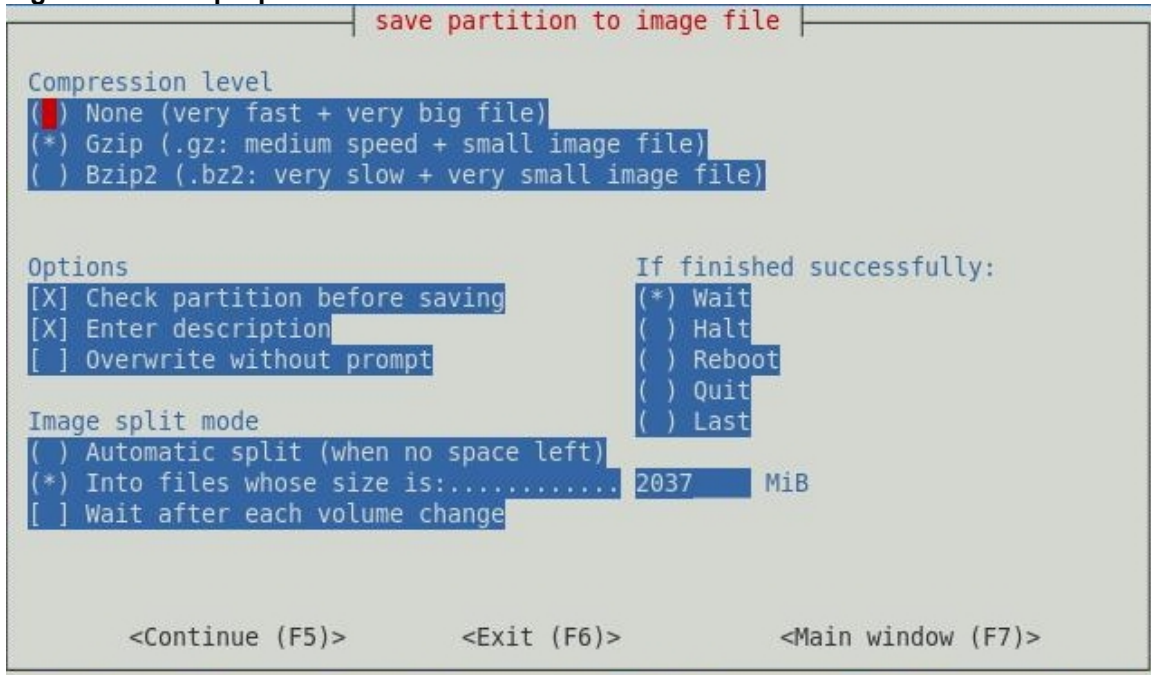
It is recommended that no special characters or spaces be used in the primary file name. Additionally, keep the filename short and meaningful.

Figure 3 : Completion of the initial partition backup details



After the path and filename have been entered, using the <tab> key tab until the “<Next (F5)>” is active. Click <Enter>.

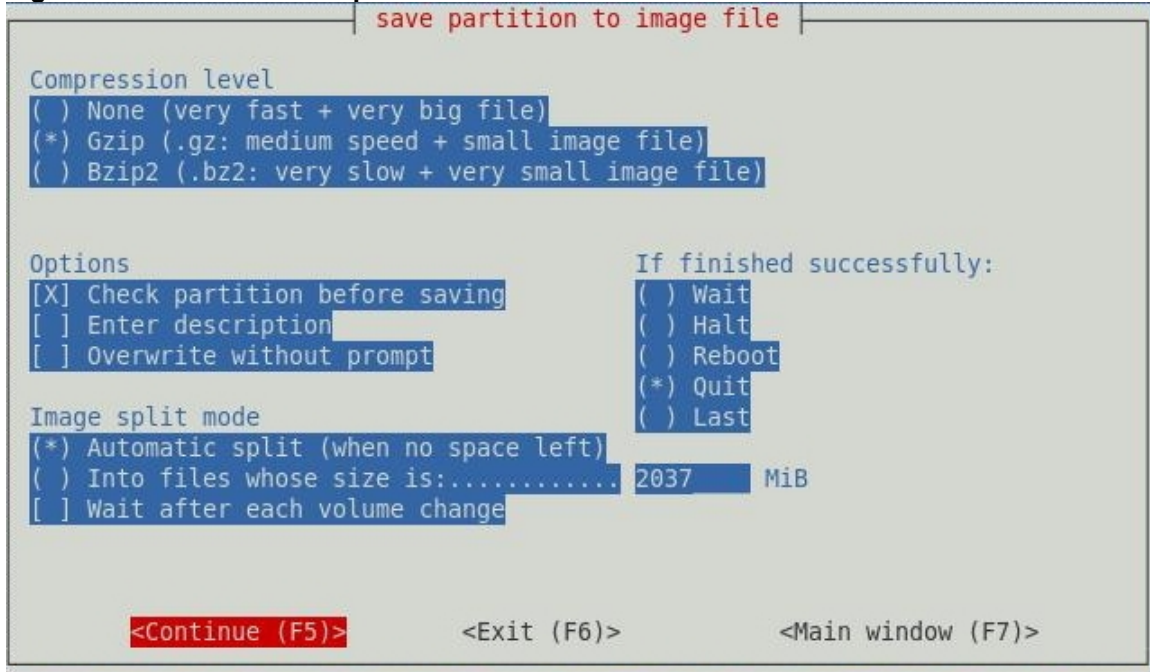
Figure 4 : Backup Options



This page presents the various backup options for the partimage session. Use the <tab> key to move between options. Use the <space bar> to toggle an option on or off.

When the options have been configured correctly, <tab> to the “<Continue (F5)>” option and click <Enter>.

Figure 5 : Recommended Options



A couple of comments on the options selected:

The Gzip option will offer a moderate amount of compression of the backup file. The tradeoff between Gzip and Bzip2 is time. Bzip2 will take a considerably larger amount of time to accomplish writing out the backup file.

It is recommended that the partition check always be run before backing up. You do not want to have a restore point made from a partition that contains problems with the file system or indexing.

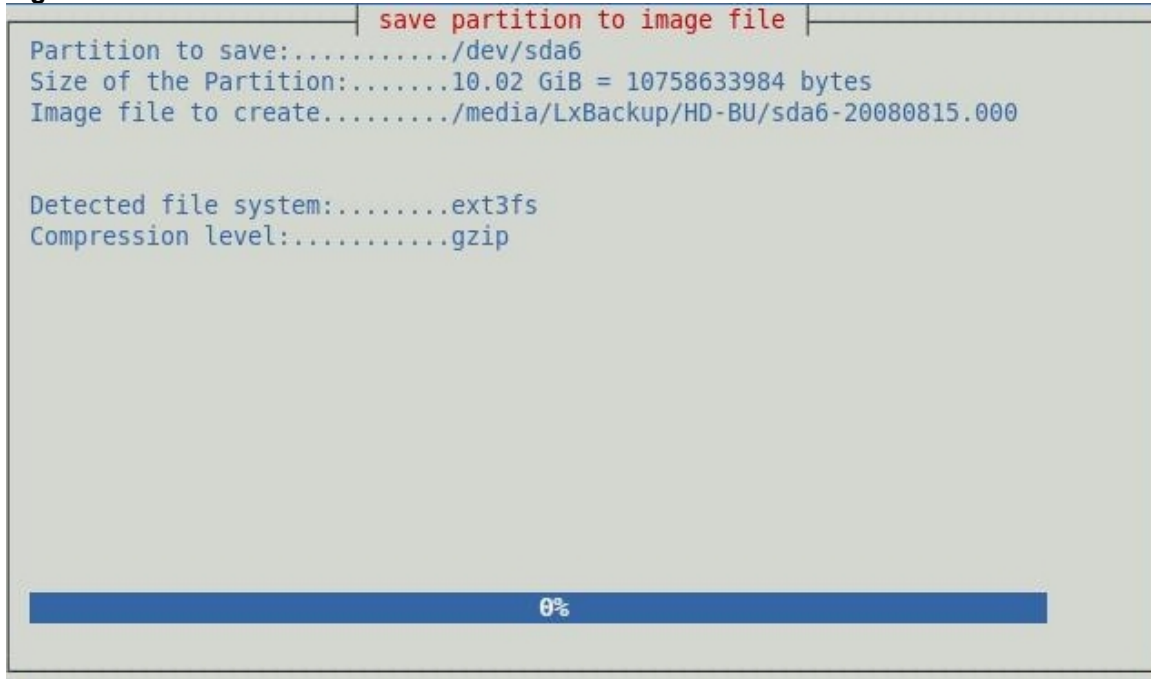
Since the backup file naming convention used is meaningful, little value may be gained by entering a description.

The “Quit” option in the “If finished successfully” section brings the desktop focus back to the terminal mode. Using the cursor <up> arrow in the terminal mode, the partimage application can be re-launched if needed for the next partition backup.

It is strongly recommended that the “Automatic split” option be used. Space permitting in the destination folder, this will allow the image file to be one file irregardless of size. This is of significance on the recovery end. If a recovery is necessary, using multiple files, the user will have to do multiple passes through the partimage recovery to accommodate each file for the partition.

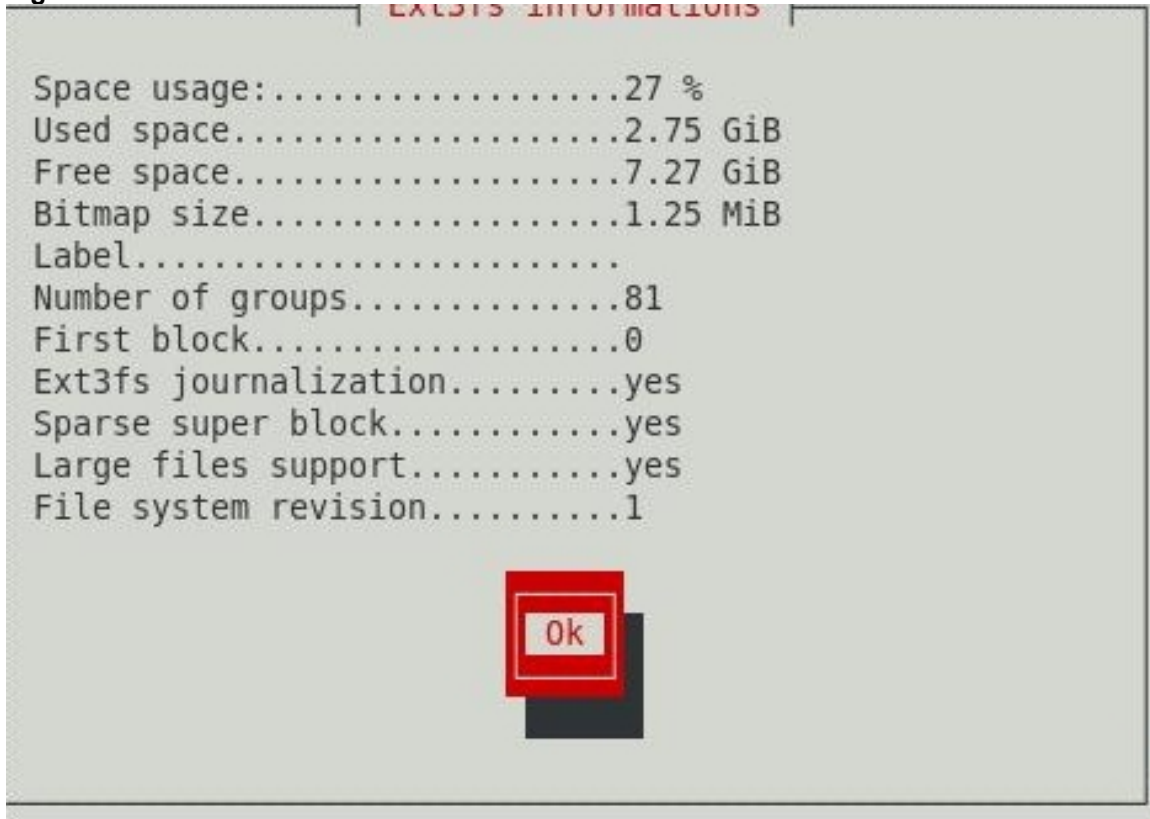
A Linux Backup Strategy – Version 1.0

Figure 6 : Filecheck Window



This is the window that will be present during the partimage filecheck.

Figure 7 : Filecheck verification:



At the completion of the file system check, click the <Ok> button to proceed with writing out the partition image file.

A Linux Backup Strategy – Version 1.0

Figure 8 : Image writing window:

```
save partition to image file
Partition to save:...../dev/sda6
Size of the Partition:.....10.02 GiB = 10758633984 bytes
Current image file:...../media/LxBackup/HD-BU/sda6-20080815.000
Image file size:.....4.79 MiBiB
Available space for image:...29.40 GiB = 31569608704 bytes
Detected file system:.....ext3fs
Compression level:.....gzip

Time elapsed:..... 5sec
Estimated time remaining:.... 5m:22sec
Speed:.....516.00 MiB/min
Data copied:.....43.00 MiB / 2.75 GiBB

1% 1%
```

Conclusion

It is hoped that the backup strategy presented here will benefit the reader in considering their own backup and restoration requirements. As mentioned in the introduction, there are a great many tools and applications to accomplish varying degrees of system-wide backups. The restoration of a partition as detailed above is relatively fast. On my last restoration, it took approximately 3-4 minutes to fully restore an operating system partition with a total size of approximately 6 Gig.

The combination of limiting the updates to immediately before running the partition backups provides a recovery point should an update go bad.

This document will be upgraded as new information is collected relative to system wide backup.

Tim Bonesho