Description of Open-E Snapshots

A Step-by-Step Guide to How To Operate Snapshot with Open-E® DSS™

May 2009

Description of Open-E Snapshot

SNAPSHOT DEFINITION

The Active Snapshot function forces the file system to write new and modified files on unused space. During a regular “live” mount, the user has an extra point-in-time to mount the whole volume. This provides access to the data existing on the volume at the snapshot start time.

The original copy of the data continues to be available to the users without interruption, while the snapshot copy is used to perform other functions on the data for Backup and Data Replication applications or user access point-in-time data.

Description of Open-E Snapshot

Basic Explanation of Snapshot Functioning

The Open-E snapshot implementation supports several concurrent active snapshots.
Active snapshots require more space on the volume then one without snapshots. Because point-in-time data cannot be overwritten during an active snapshot, this will need the extra space. Deleted data is claimed as free space in a “live” volume mount, but in reality the deleted data is still available in the snapshot mount. The size of reserved space for snapshot depends on the amount of changed data while the snapshot is active. Daily scheduled snapshots will need less reserved space then weekly scheduled snapshot.
To calculate the size of the reserved space for snapshot, estimate the amount of expected data changes when the snapshot will be active. A good rule to use would be 2 or 3 times the size. For example, if a volume size is 1000GB and we are reserving space for daily snapshot and we expect about 10GB changes every day, the reserved space for snapshot will be 20 to 30 GB. If we decide for 30GB, it will be 3% of the whole 1000GB volume.

Description of Open-E Snapshot

OPEN-E SNAPSHOT TECHNOLOGY

- Snapshot is based on the Logical Volume Manager (LVM).
- Snapshot implements a “copy-on-write” on the entire block devices by copying changed blocks just before they are to be overwritten to the other storage, thus preserving a self-consistent past image of the block device. The file systems on this image can later be mounted as if it were on read-only media.
- The total number of snapshots is dependent on the LVM, but 255 is a safe number. There are, however, definite limitations as to the number of ACTIVE snapshots (32bit and 64bit system) running at the same time:
  - ✓ 10 per LV,
  - ✓ 20 per system.

Description of Open-E Snapshot

File Level Description

File system on the volume before snapshot is started. State at 7:59 A.M.

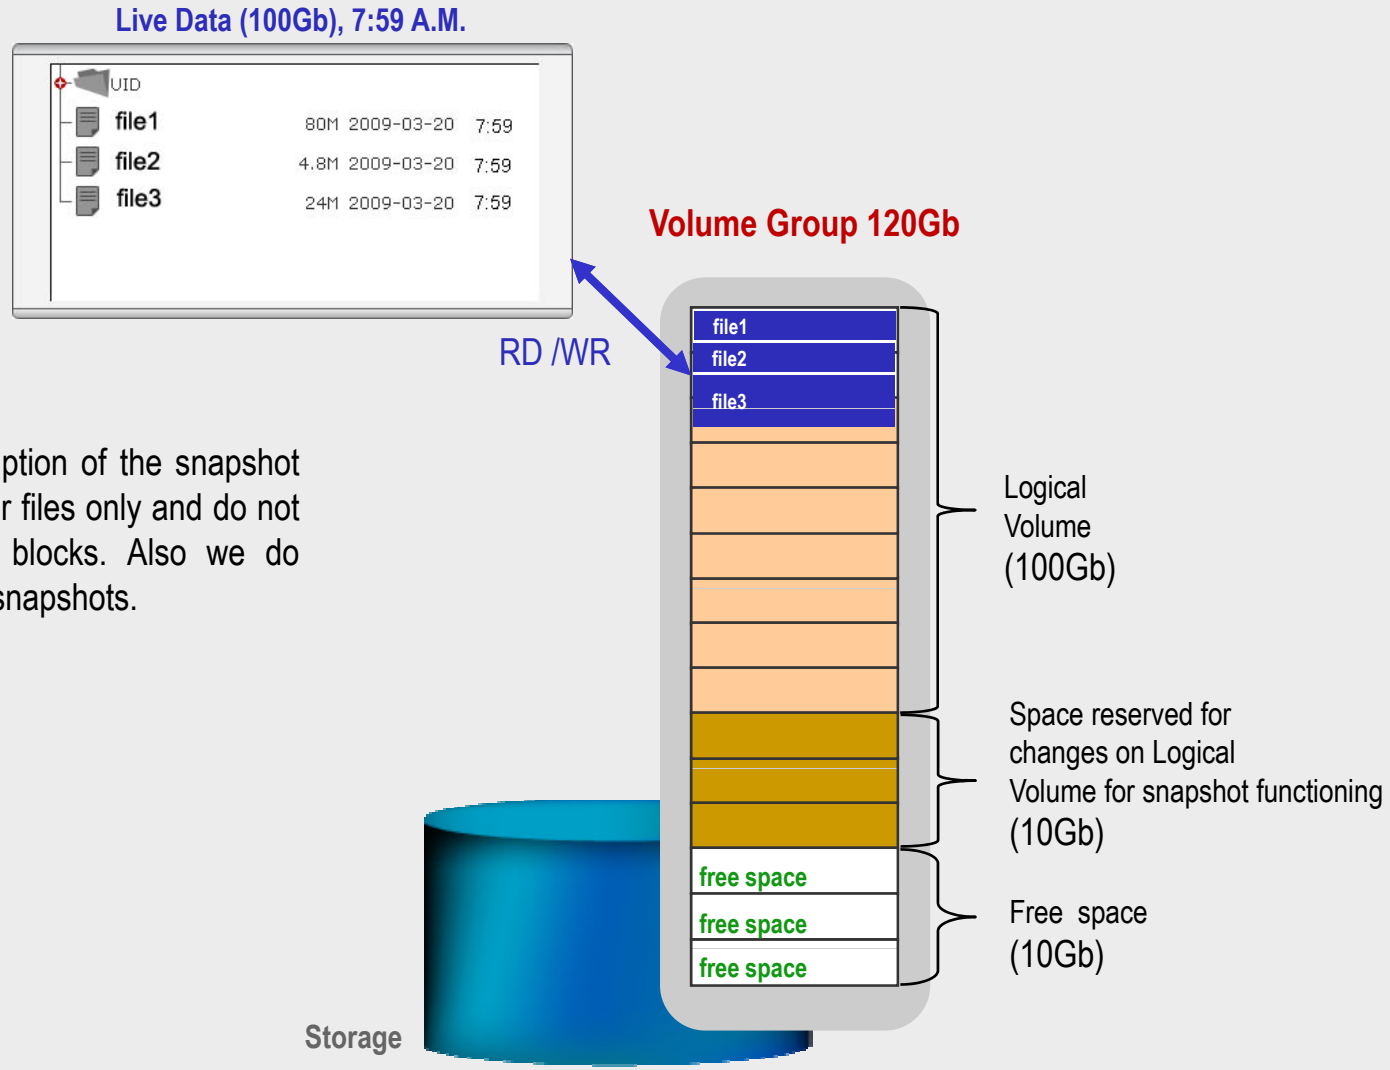

For a simple description of the snapshot function we consider files only and do not talk about volume blocks. Also we do describe read-only snapshots.

Description of Open-E Snapshot

File Level Description

8:00 A.M. - Snapshot starts.

Live Data (100Gb), 8:00 A.M.

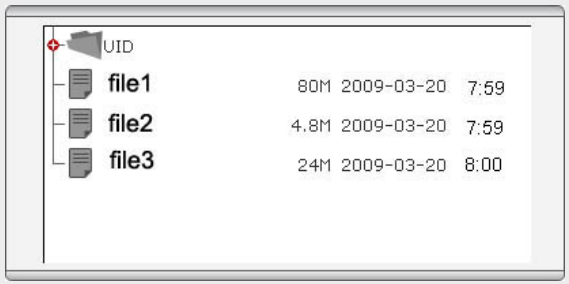

Snapshot (100Gb), frozen in point-in-time at snapshot start time 8:00 A.M.

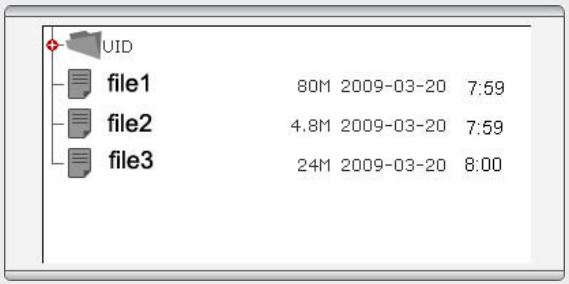

After the snapshot starts, the files system will read the data as usual, but writes will first copy the original data into the snapshot reserved space (copy-on-write) and then write new data.

As a result of every write operation during the active snapshot the process will be using reserved space of logical volume for written data.

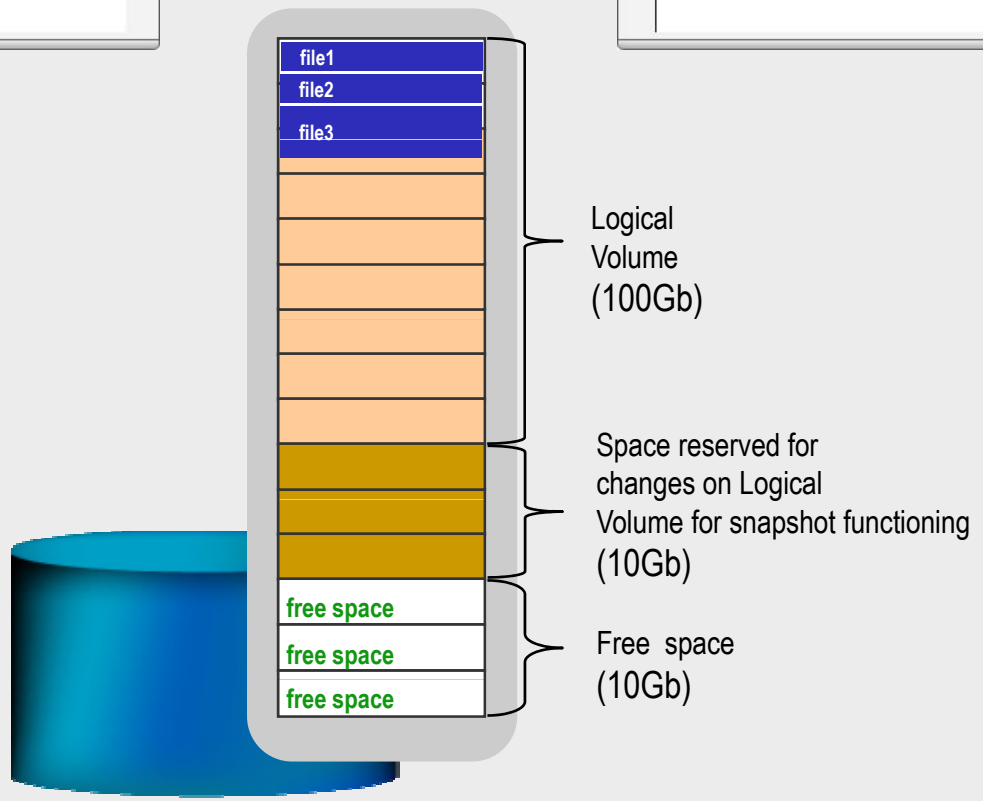

Description of Open-E Snapshot

File Level Description

8:12 A.M. - Snapshot active

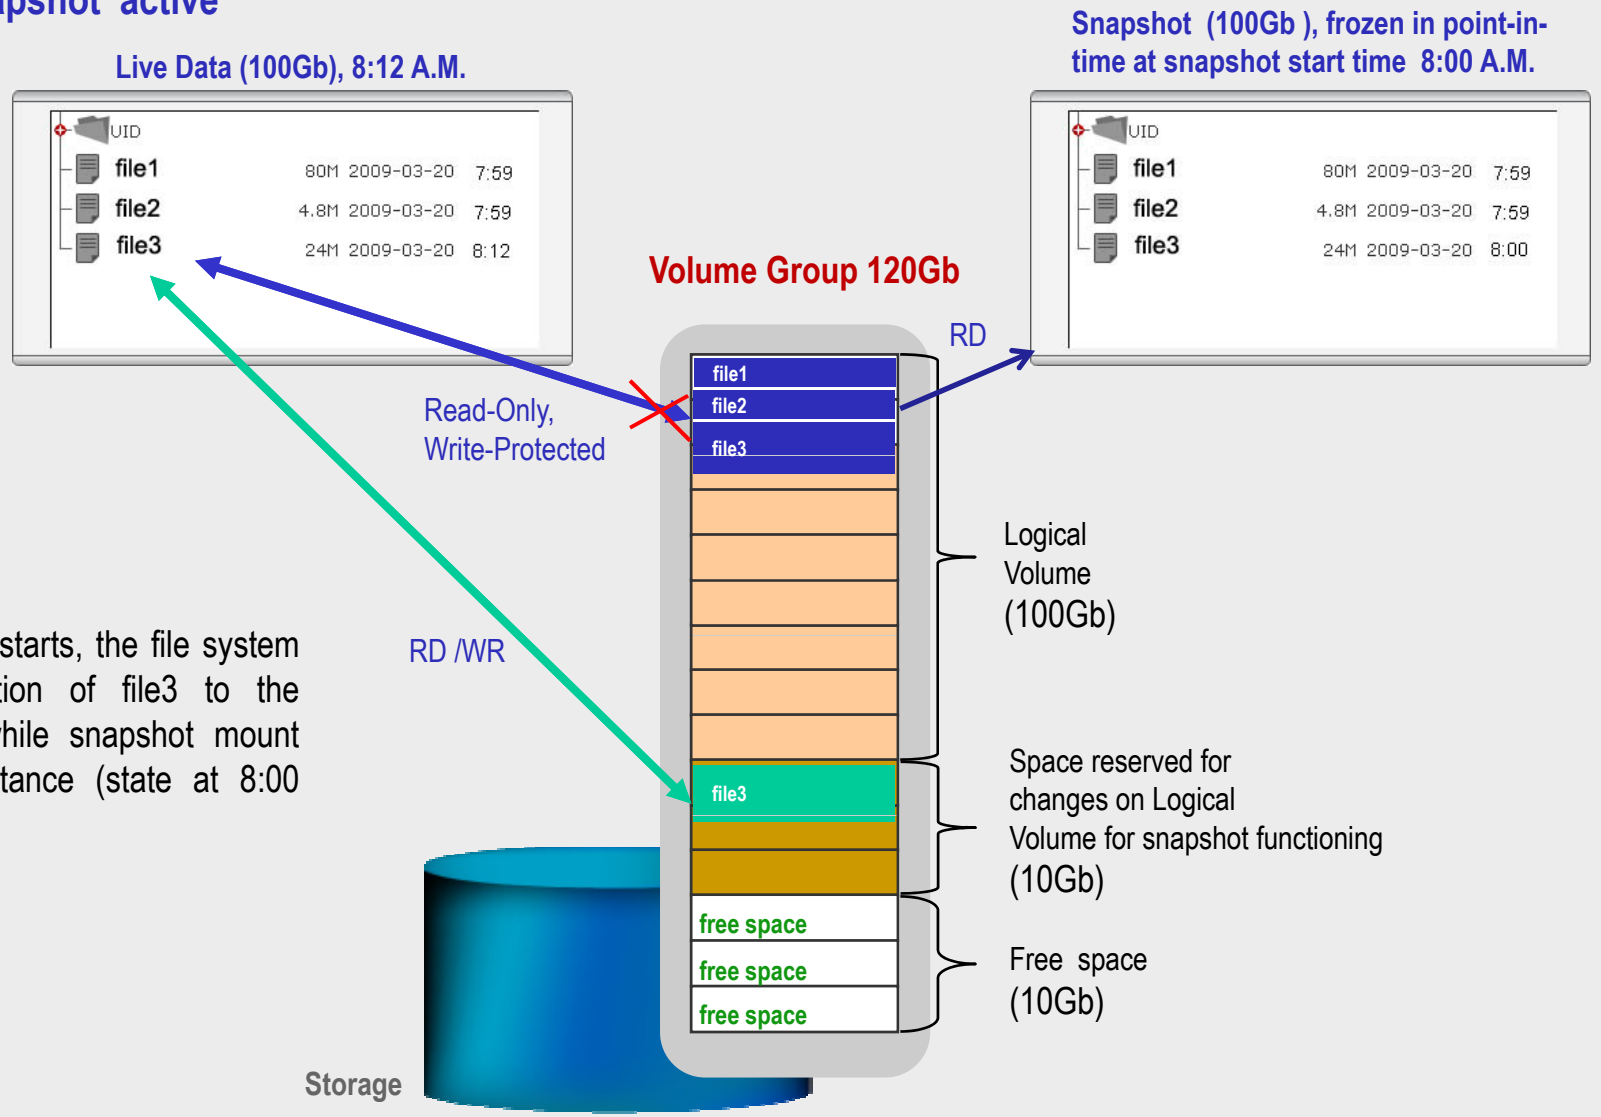

After the snapshot starts, the file system writes a modification of file3 to the reserved space, while snapshot mount show old file3 instance (state at 8:00 A.M.)

Description of Open-E Snapshot

File Level Description

8:19 A.M. - Snapshot active

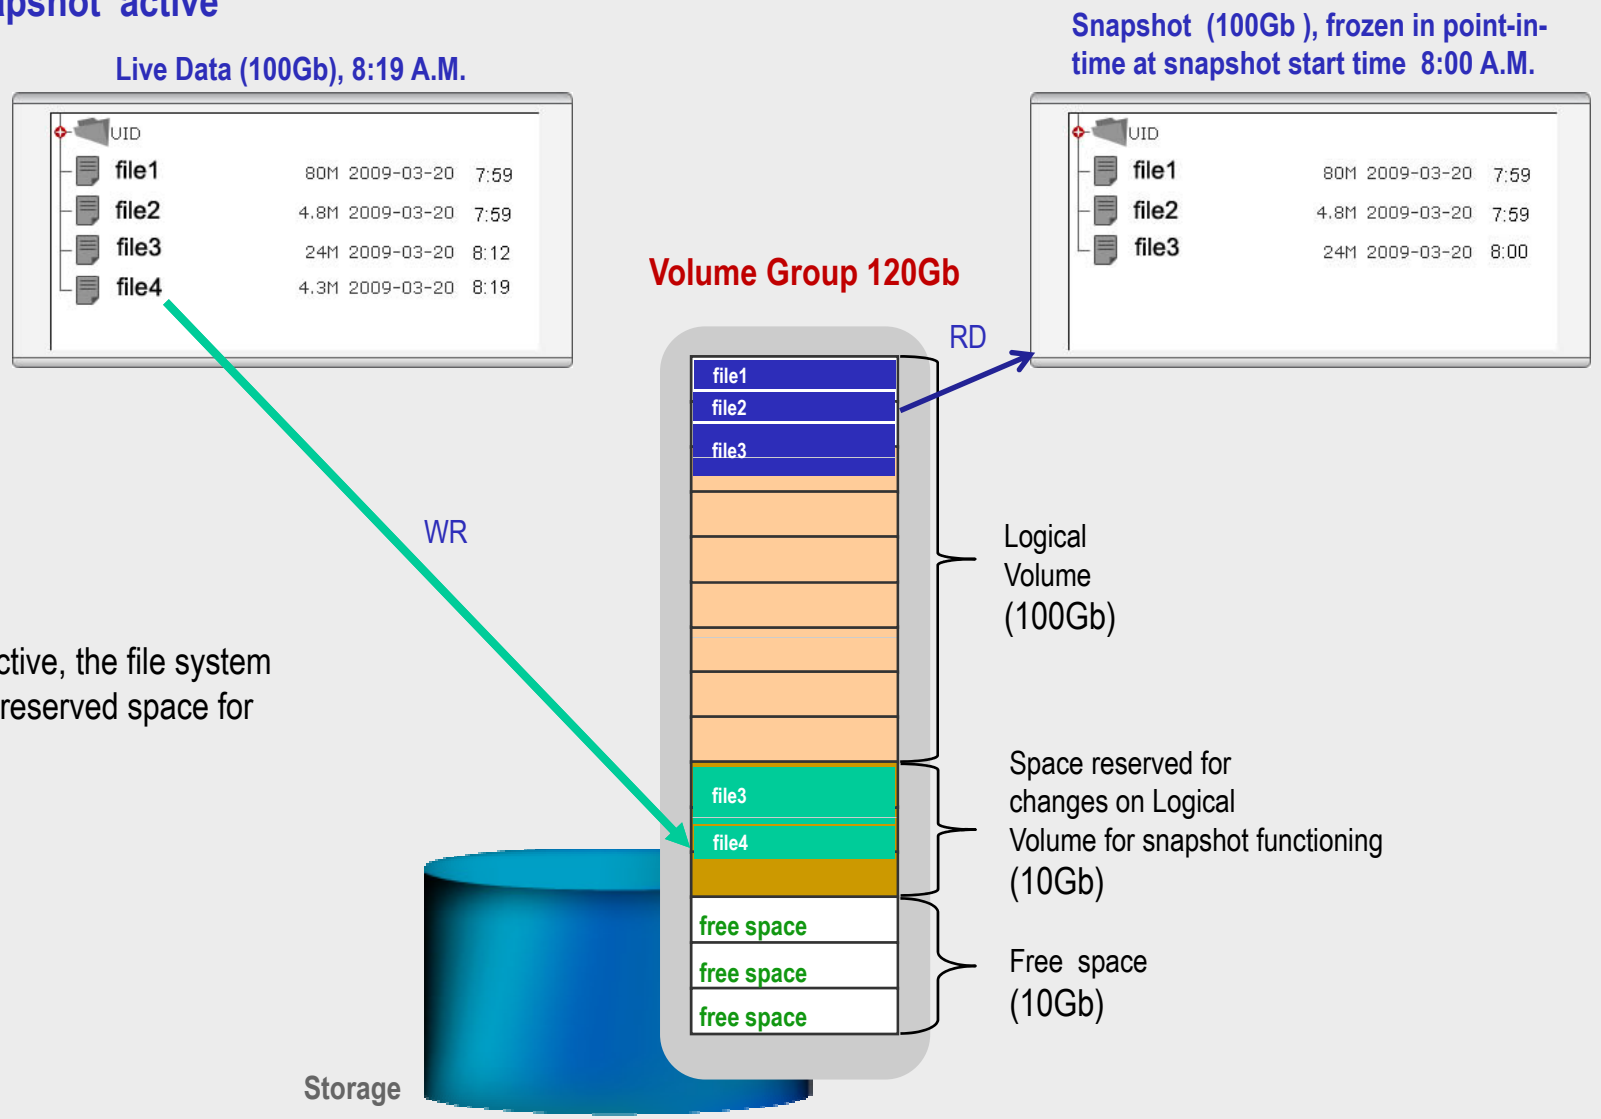

While snapshot is active, the file system forces the usage of reserved space for the new file4.

Description of Open-E Snapshot

Block Level Description

For an easier explanation of how snapshot functions, we consider files only and do not talk about volume blocks.

For example, a file uses 2 blocks and is being modified, if the first part which resides on first block, then the write operation is related to the first block only.

The rest of the file (unchanged block) is being copied into the reserved space while snapshot is active. The (COW) Logical volume will lose the bind with the old block and get a new bind with just written blocks in the reserved space. The copy-on-write operation is shown on next page.

Description of Open-E Snapshot

BLOCK LEVEL DESCRIPTION

8:12 A.M. - Snapshot active - modification of the file3.

For example, the file system is going to modify the file3 which resides on 2 blocks (block1 and 2). The snapshot is active. For example, the file system is going to modify the file3 which resides on 2 blocks (block1 and 2). The modification must happen on the small part of the second block only. Now the block 2' is copied into reserved space and then the file system posts the modification (green color). Now, file3 is located partially on Logical Volume area on block 1 and on block 2 which resides on the reserved space for changes.

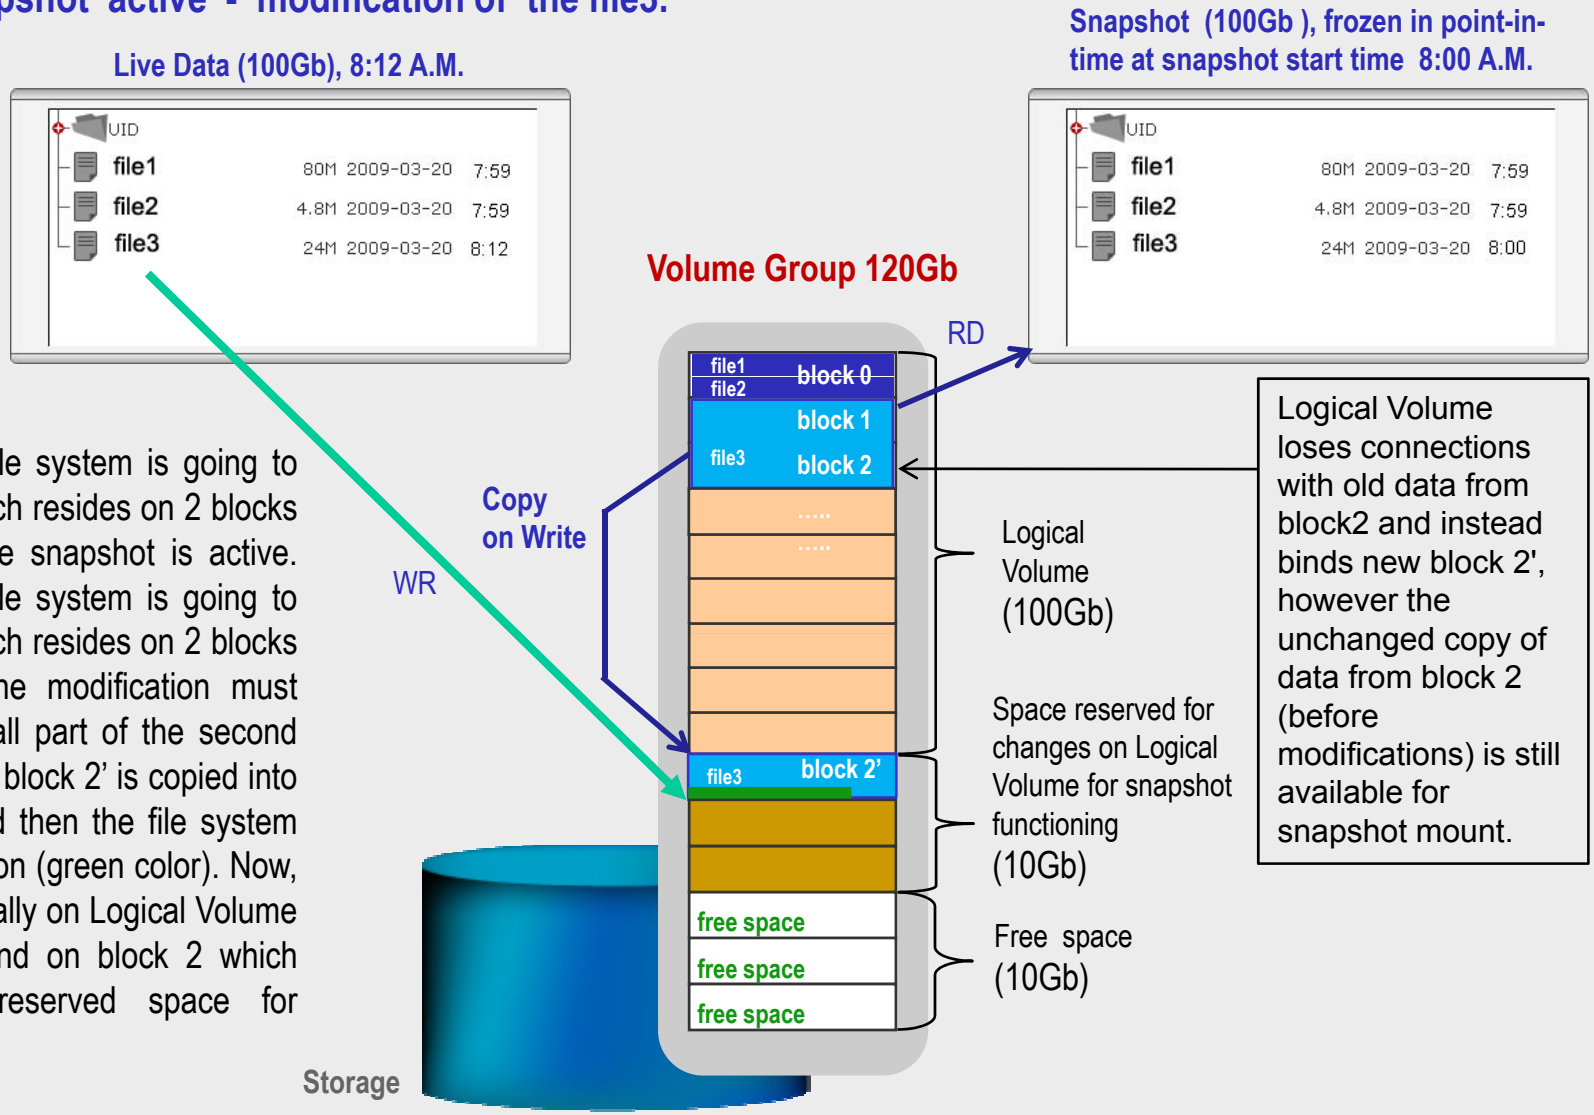

Description of Open-E Snapshot

Block Level Description

File system on the Volume after deactivation of Snapshot – state after 8:12 A.M.

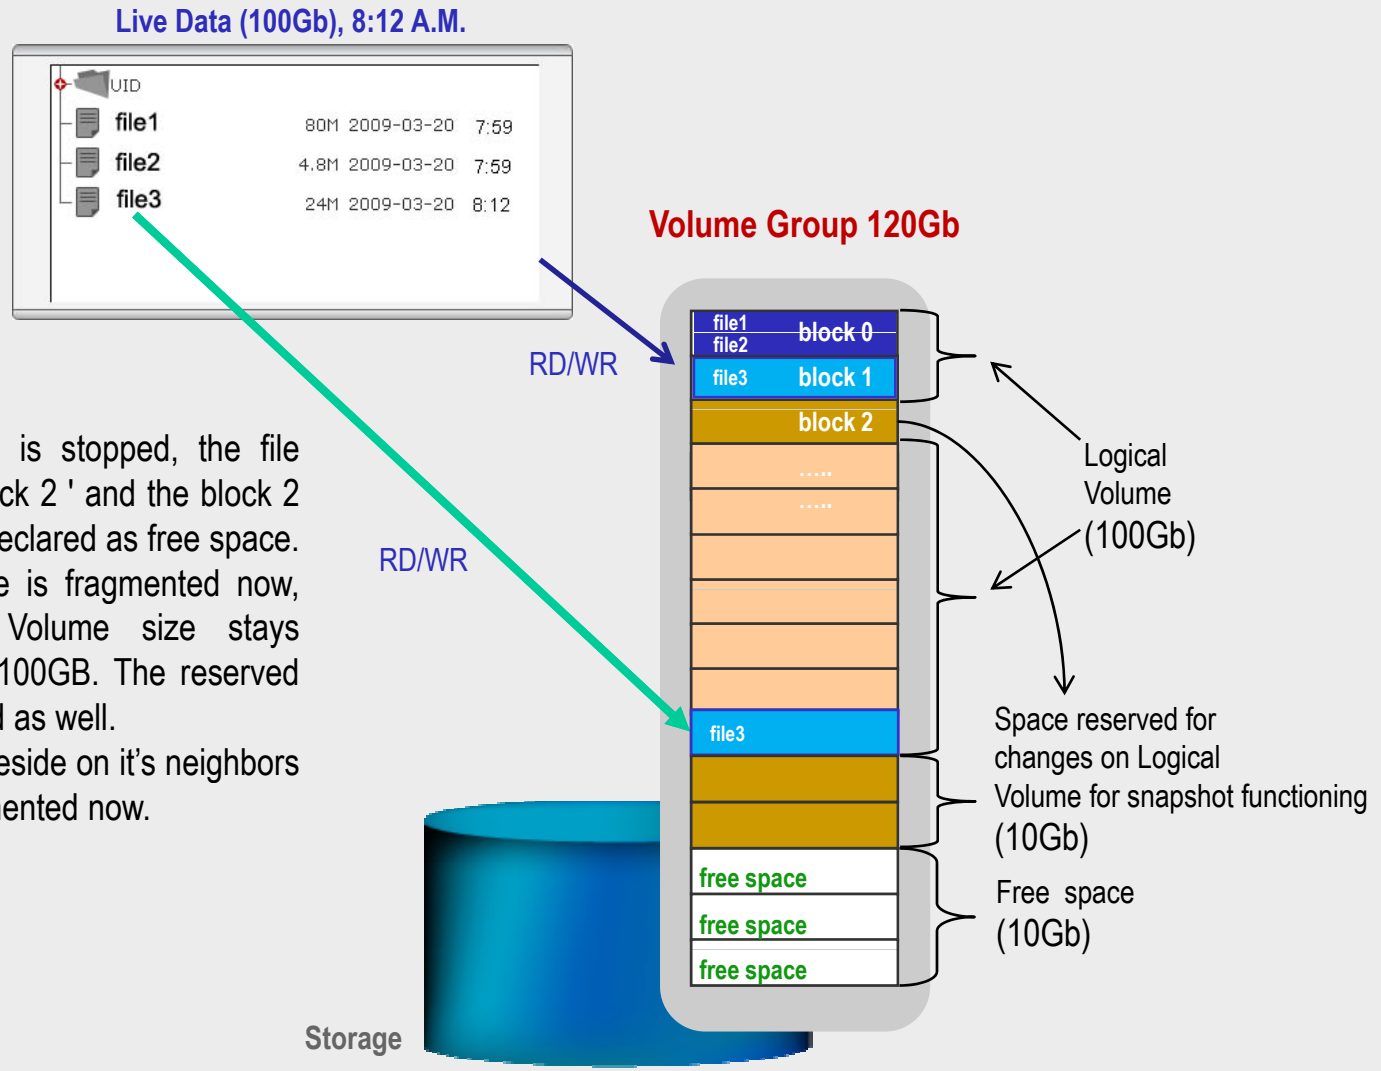

After the snapshot is stopped, the file system binds to block 2 ' and the block 2 is unbound and is declared as free space. The Logical volume is fragmented now, but the Logical Volume size stays unchanged and is 100GB. The reserved space is fragmented as well.
The file3 does not reside on it's neighbors blocks, but is fragmented now.

Description of Open-E Snapshot

Block Level Description

User decides to remove (delete) the reserved space for snapshot functioning. State after 8:12 A.M.

Live Data (100Gb), 8:12 A.M.

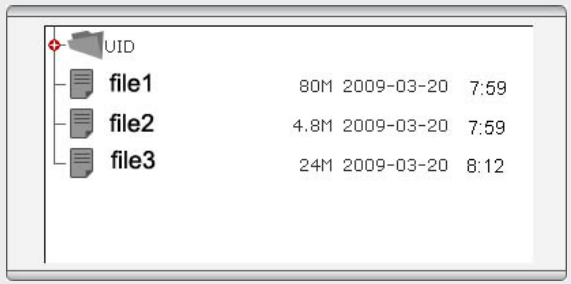

After the deletion of the reserved space for snapshot function, this space is declared back as free space. The free space is 20Gb now.

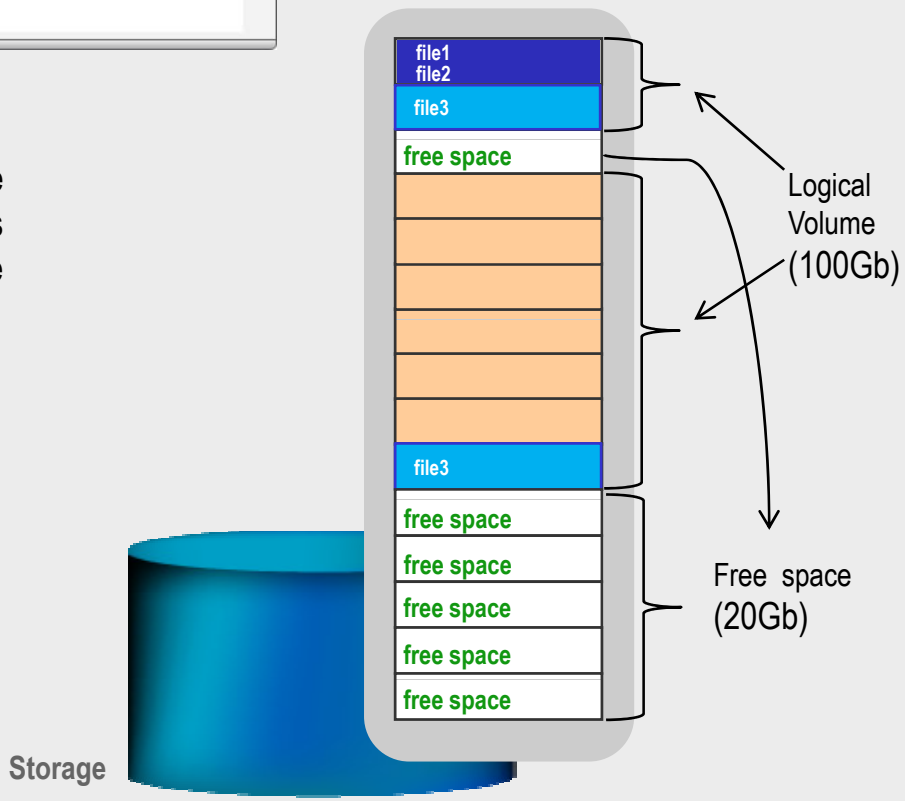

Description of Open-E Snapshot

open-e

ADVANTAGES

- If there are no modifications made to the original data, no data copy will be created.
- Snapshot is used for Backups and Data Replication, snapshot provides an exclusive virtual access to the volume. This exclusive access for Backup or Data Replication can work any time, also during production hours.
- Starting or stopping a snapshot is very fast, this only takes a few seconds even for large amount of data.

Description of Open-E Snapshot

DISADVANTAGES

- Overflow of space reserved for snapshot causes the snapshot volume to become inaccessible and lost access to the point-in-time data.
- Writing speed decreases with growing number of active snapshots (because of copy-on-write).
- Copying data from snapshot mount back to the live mount, may cause overflow of space reserved for snapshot and as result unfinished copy.
- In the case of iSCSI or FC snapshots, reading only the data can cause a modification of last access time of a file. This will cause the whole block chunk to be moved to a space reserved for changes. The disadvantage will appear on file systems with an enabled last access time attribute, as a read operation will cause a change of last access time attribute. It does not take place for NAS volumes as Open-E systems will not store last access time.
- Obviously, an iSCSI or FC target can be formatted by any file system. Most of file systems support the following file attributes: creation, modify and last access time. If last access time is used, any read access will cause change of this attribute and as result will write to the volume. Snapshot works on the volume level and uses a block size of 32 MB. Single file read will result one block change, so a consumption of 32MB is reserved space for snapshot functionality.

Description of Open-E Snapshot

SUMMARY

The Snapshot functions within Open-E software is used for:

- Backups
- Data replications
- Access to accidentally deleted or modified files

Thank You!