

# Understanding Snapshot Consumption

Source : <http://broadcast.oreilly.com/2009/02/understanding-snapshot-consump.html>

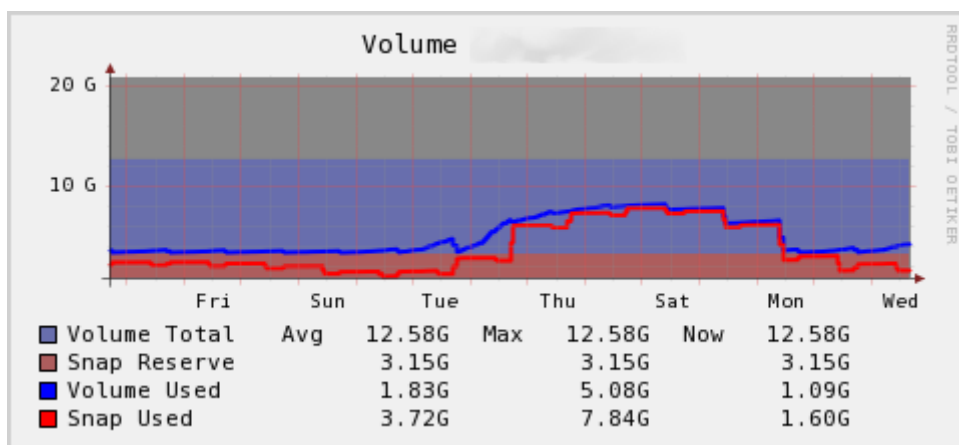
The following is an experiment to visualize snapshot consumption on a Netapp filer. Snapshots are references to how a filesystem appeared at a certain point in time. It's an easy way to retrieve files that may have been deleted or changed accidentally by referring to an online copy of the file. Snapshots record this data by keeping track of the blocks that have changed in between snapshots and the current filesystem. The more changes you introduce to a filesystem can correlate to increases in the snapshot size.

There are other snapshot technologies out there, some that work strictly at the block level, or others that integrate the snapshots into the live filesystem. It may be possible to run this experiment against those filesystems, but the process will probably be different.

A Netapp has two types of volumes. The traditional volume is a filesystem that is based on a hard allocation of disks, which can be increased in size by adding more disks to it, but it cannot decrease in size. The flexible volume is a filesystem that is overlaid across a series of disks known as an aggregate. It can be increased or decreased at any time, as long as it doesn't exceed the available space on the aggregate. This allows multiple volumes to share the same physical disks, increasing the ability to provision all available storage.

All volumes start out with a reserve area for snapshots, typically twenty percent of the volume size. Snapshots can exceed the reserve, but regular files cannot consume the snapshot reserve area. It is possible to change the amount of space set aside for the reserve area. The reserve space also adjusts itself if the volume is resized.

The following graph tracks a filesystem and shows the total space of the volume (light blue area), the amount of the reserve space (light red area), the storage in use by the main filesystem (blue line), and the storage used by all the snapshots for that volume (red line).



The volume shown in this graph handles database log files, so the normal usage involves continuous updates to a series of files. The graph covers a two week period.

The RRD file was created by a perl script that runs SNMP polls against the filer. For a given volume, it takes the total space and the used space, derived from a pair of high and low 32 bit counters, and merges them into a 64 bit integer. It does the same calculations against the snapshot area tied to the same volume. That's eight SNMP data points, brought together into four variables in a single RRD file.

Since we're dealing with a resizable filesystem, the graph will adjust the light blue background, and the red background if the size of the volume changes, or if the snapshot reserve value changes. This allows the administrator to have a record of any resizing operations against the volume over a period of time.

Given what we know about the behavior of snapshots, we can assume the following:

1. Increases in snapshot consumption (the red line), indicate changes in blocks. If the amount of space used by snapshots increases, even though the amount of used storage by the files remains unchanged, we could assume that existing files in the volume are undergoing a high rate of change. *Remember, the red line can cross into the blue area, but the blue line will not cross into the red area.*
2. A large increase in the size of snapshots could also indicate the deletion of files, since a deleted file isn't really deleted until every snapshot referencing the file is removed. The increase in the snapshot size should have a correlating decrease in the main filesystem when a file is deleted.
3. It could be possible for the snapshot storage used to exceed the filesystem storage used based on the rate of change in files, or depending on how old the accumulated snapshots are. Monitoring the size of the snapshots can be a good way to determine how large your snapshot reserve should be.
4. Decreases in snapshot consumption are likely an indication of snapshots being rotated out, or removed from the volume.

The software doing this is all in serious beta, and still a little klunky. I'm also trying to see if this graph layout is the best way to represent utilization, and the color scheme could probably use some work. Otherwise, I'd appreciate feedback.