



NetApp™
Go further, faster

PROFESSIONAL SERVICES TECH NOTE #052

Understanding snap list Output

The Concept

It is easy to determine which snapshots are in a volume using the **snap list** command. However, what do the four columns of numbers mean? Which one tells me how much space an individual snapshot is using?

Examples

This 100GB volume is nearly empty.

```
fas3020a> df -h PaulG
Filesystem                total      used    avail capacity  Mounted on
/vol/PaulG/                80GB       17MB     79GB      0%  /vol/PaulG/
/vol/PaulG/.snapshot       20GB       61MB     19GB      0%  /vol/PaulG/.snapshot
```

```
fas3020a> snap list PaulG
Volume PaulG
working...
```

%/used	%/total	date	name
31% (31%)	0% (0%)	Mar 08 00:00	nightly.0
47% (29%)	0% (0%)	Mar 07 20:00	hourly.0
57% (32%)	0% (0%)	Mar 07 16:00	hourly.1
64% (29%)	0% (0%)	Mar 07 12:00	hourly.2
69% (32%)	0% (0%)	Mar 07 08:00	hourly.3
73% (29%)	0% (0%)	Mar 07 00:01	nightly.1
76% (30%)	0% (0%)	Mar 06 20:00	hourly.4
78% (31%)	0% (0%)	Mar 06 16:00	hourly.5

The %/used column shows space consumed by snapshots divided by currently used blocks in the volume, regardless if that used space is in the Active File System (AFS) or snapshots. The first number is cumulative for all snapshots listed. The (second number) is for the specified snapshot alone. The snapshots in /vol/PaulG hold 78% of all the used blocks and each snapshot link to 29 – 32% of the used blocks.

The %/total column shows space consumed by snapshots divided by total disk space in the volume. This includes data space + snapshot reserve and includes all used and free blocks available. At 100GB, the 61MB of snapshots consume < 1% of the total volume space.

The combination of new, changed, and deleted blocks in a volume determines how closely the numbers in these columns track with each other. If the same set of files or amount of data turns over then % total will remain the same but %used will vary widely. If the rate of change remains constant, but the volume grows or shrinks then the %total will change while %used remains the same.

Watch how the numbers in the columns change after the volume size is decreased from 100GB to 1GB.

```
fas3020a> df -h PaulG
Filesystem                total      used    avail capacity  Mounted on
/vol/PaulG/               819MB       17MB     801MB      2%  /vol/PaulG/
/vol/PaulG/.snapshot      204MB       61MB     143MB     30%  /vol/PaulG/.snapshot
```

```
fas3020a> snap list PaulG
Volume PaulG
working...
```

%/used	%/total	date	name
31% (31%)	1% (1%)	Mar 08 00:00	nightly.0
47% (29%)	1% (1%)	Mar 07 20:00	hourly.0
57% (32%)	2% (1%)	Mar 07 16:00	hourly.1
64% (29%)	3% (1%)	Mar 07 12:00	hourly.2
69% (32%)	4% (1%)	Mar 07 08:00	hourly.3
73% (29%)	4% (1%)	Mar 07 00:01	nightly.1
76% (30%)	5% (1%)	Mar 06 20:00	hourly.4
78% (31%)	6% (1%)	Mar 06 16:00	hourly.5

The snapshots now consume a larger percentage of the total volume space, but the percent of used blocks in snapshots compared to active filesystem blocks remains the same.

What It Means To You

So how should you use these numbers? When you have a lot of snapshot data and available free space is high, increase the snapshot reserve. When free space is low, then increase the volume size.

According to all the documentation, % used is good for determining if snapshot reserve is too large or small. A large reserve can act as a built in buffer when the volume fills up. A reserve that is too small will bleed snapshots into data space and risks unexpectedly filling up the volume (if a volume fills up *on the hour*, most likely a snapshot is to blame)

The % total is an indication when the snapshot reserve is too small and has overflowed into data space and by how much. If you want to keep your snapshots, use these numbers to determine how much larger to grow a volume. As the volume grows, so does the snapshot reserve until it is large enough to hold 100% of snapshot blocks.

If you decide to delete snapshots, you should use the **snap delta** and **snap reclaimable** commands to figure out which snapshots to delete. Be sure to run these commands immediately before deleting a snapshot. The number extensions on snapshot filenames increments when new snapshots are taken. For instance, if you take snapshots every hour, then the hourly.1 snapshot at 9 AM will have incremented to hourly.5 by 1 PM (Same applies to nightly.# and weekly.# snapshots. Named snapshots --SnapMirror, SnapVault, and manually created -- do not increment.)

More Information

Refer to the Data ONTAP Online Backup Guide to understand how the values in **snap list** are calculated:

http://now.netapp.com/NOW/knowledge/docs/ontap/rel732_vs/html/ontap/onlinebk/GUID-841C0685-93D0-4B11-8D63-4B8E2B4C2419.html

A great Knowledge Base article explaining how to use and interpret the different snap commands:

<https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb15495>

Brought to you by NetApp Americas Professional Services

Send comments, corrections & suggestion to Michael Cope at mcope@netapp.com
Tech Note archive: http://communities.netapp.com/community/products_and_solutions?view=documents