

Monitoring Snapshot Space Consumption with FlashArray™

Snapshots are *incremental* data copies, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data.

Snapshots are designed to be used as recovery points during changes to an underlying volume. Most often they will be deleted based on a 7-day retention schedule. It is important to understand the risk of having them live longer: Not only do they consume space, but they also tend to cause a performance hit, if not designed properly.

Snapshots are not intended to be used as backup copies, since they do not have a file system or database level consistency check.

Coordinated snapshots can be taken across multiple storage volumes, allowing you to take exact point-in-time, data coordinated and crash-consistent snapshots across multiple volumes.

Snapshot Space Consumption

As with conventional disks, a volume's storage capacity is presented as a set of consecutively numbered 512-byte *sectors* into which data can be written and from which it can be read. Hosts read and write data in *blocks*, which are represented as consecutively-numbered sequences of sectors. FlashArrays do not allocate physical storage for volume sectors that no host has ever written, or for *trimmed* (expressly deallocated by host or array administrator command) sector addresses. FlashArrays allocate only the exact amount of physical storage required by each host-written block after reduction. In FlashArrays, there is no concept of allocating storage in "chunks" of some fixed size.

The Purity//FA storage software applies pattern elimination, duplicate elimination, and compression techniques to data as it enters an array, as well as throughout the data's lifetime. FlashArray snapshots occupy physical storage only in proportion to the number of sectors of their source volumes that are overwritten by hosts. The net physical storage consumed by host-written data is a fraction of the number of unique volume sector addresses to which hosts have written data.

Snapshot Space Reporting

Snapshot Space is defined as the amount of space that would be freed by deleting snapshots. The figure below is a screenshot of the space used by the vSHNData 01-03 volumes and their snapshots. The space usage of a volume/snapshot as reported in the GUI is defined as follows:

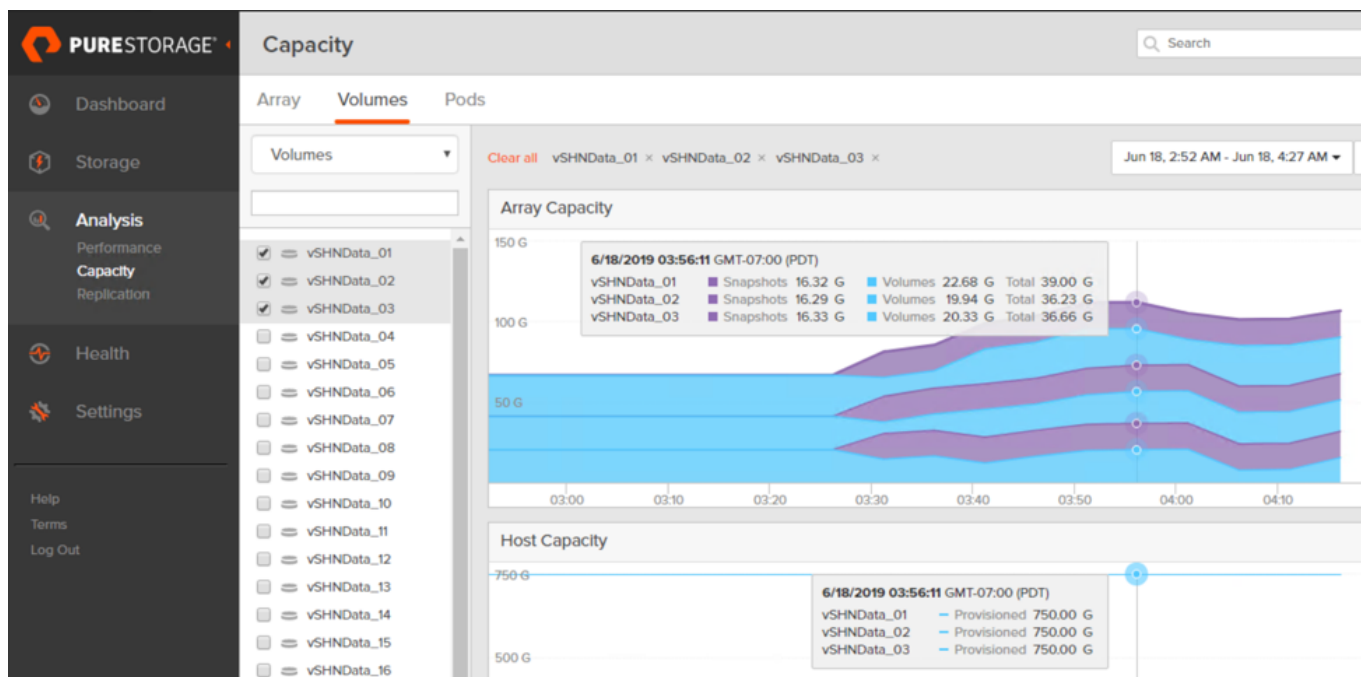
1. Space attributed to the vSHNData_01 volume:

The amount of space that is unique to the volume (i.e. which is not duplicated on another volume or snapshot).

2. Space attributed to a snapshot of the vSHNData_01 volume:

The amount that would be freed if that snapshot was deleted. This is the amount of space in the volume that has been overwritten on the volume (or a newer snapshot) but not de-duped since the point the

snapshot was taken. If there have been no writes to the volume since the time the snapshot was taken, the size of the snapshot will be 0.



Note that space reporting does not report what the differences between the snapshots and volumes are, only the amount of space that could be reclaimed by deleting the volume or the maximum that could be reclaimed by deleting a particular snapshot.

Snapshot Space Monitoring/Alerting

As storage administrator, it is important to set usage alerts, and forward them to those that need to see them. A production storage array should never consume more than 80% of capacity and performance (CPU). At 75% of the storage system space consumption, we want to set a yellow alert, at 95% a red alert. It is important to address storage consumption level issues before they occur to avoid running into major problems. If you use snapshots properly, these alerts will hopefully never show.

To manage snapshots capacity consumption, quotas can be defined in the form of a tuneable parameter, based on snapshot size per volume or per system. In case of space crunch, it is possible to limit the snapshot space in the system.

Because data stored in a FlashArray is virtualised, thin-provisioned, and reduced, volume storage is monitored, managed, and displayed from two viewpoints:

1. Host view: Displays the virtual storage capacity (size) and consumption as seen by the host storage administration tools.
2. Array view: Displays the physical storage capacity occupied by data and the metadata that describes and protects it.

Capacity on the array is monitored as per the following example:

```

root@sn1-x50r2-f04-33-ct0:~# purevol list --space
Name          Size Thin Provisioning Data Reduction Total Reduction Volume Snapshots Shared Space System Total
vSHNData_01  750G 97%                1.8 to 1      52.0 to 1     11.83G 15.96G - - 27.79G
vSHNData_02  750G 97%                2.2 to 1     77.1 to 1     7.15G 15.93G - - 23.09G
vSHNData_03  750G 97%                1.4 to 1     48.0 to 1    14.40G 15.99G - - 30.38G
vSHNData_04  750G 97%                1.4 to 1     43.6 to 1    15.62G 19.65G - - 35.27G
vSHNData_05  750G 97%                2.1 to 1     76.2 to 1     7.50G 15.96G - - 23.46G
vSHNData_06  750G 97%                2.0 to 1     76.9 to 1     7.59G 16.03G - - 23.62G
vSHNData_07  750G 97%                1.3 to 1     41.8 to 1    16.84G 15.72G - - 32.55G
vSHNData_08  750G 98%                1.5 to 1     68.7 to 1     9.57G 19.05G - - 28.63G
vSHNData_09  750G 98%                1.2 to 1     52.8 to 1    13.60G 16.00G - - 29.60G
vSHNData_10  750G 97%                1.5 to 1     59.3 to 1    11.29G 15.93G - - 27.21G
vSHNData_11  750G 97%                2.1 to 1     74.9 to 1     7.49G 16.01G - - 23.49G
vSHNData_12  750G 97%                1.3 to 1     41.8 to 1    16.60G 19.06G - - 35.66G
vSHNData_13  750G 97%                1.4 to 1     45.7 to 1    14.75G 16.25G - - 31.00G
vSHNData_14  750G 97%                1.3 to 1     43.5 to 1    16.06G 16.00G - - 32.06G
vSHNData_15  750G 97%                1.4 to 1     54.6 to 1    12.58G 15.92G - - 28.51G
vSHNData_16  750G 97%                1.6 to 1     61.3 to 1    10.62G 15.71G - - 26.33G

```

Useful alerts, recommended to monitor, are the following:

1. Alert 50 - Volume (or Protection Group) Snapshots are at X% of Snapshot Limit:

Snapshot creation is quick and easy. This alert informs you that the number of snapshots created in a Protection Group are either coming close or has met the limit the number of snapshots that can be made. You will not be able to create any more snapshots within the Protection Group.

2. Alert 142 - Volumes are at X% of Limit:

This alert informs you that the number of volumes created in an array are either coming close or has met the limit the number of volumes that can be made. You will not be able to create any more volumes.

3. Alert 25 - Storage Consumption has Reached X% of Usable Capacity:

This alert has three thresholds at which it will alert: 80%, 90%, 100%. These alerts serve as a warning for the FlashArray Administrator about an upwards trend in space consumption. If the array fills up on space usage, without correction, you may suffer an outage. If phone home is enabled, at the 100% threshold, this alert is also sent to Pure Storage support so they can investigate the issue and reach out to you with any recommended actions.

Example:

Appliance	Updated	Summary	Status	Closed
smf-pure-002	2019-04-05 04:47:54	storage array.capacity high utilization	closed	2019-06-18 14:47:47
lad-pure-005	2019-04-19 06:11:59	storage array.capacity high utilization	closed	2019-06-18 14:47:45

Predictive Planning/Alerting

Pure1® Meta™ applies predictive analytics to a massive collection of storage array capacity and

performance data to enable both a white glove customer support experience and breakthrough capabilities like accurate forecasting.

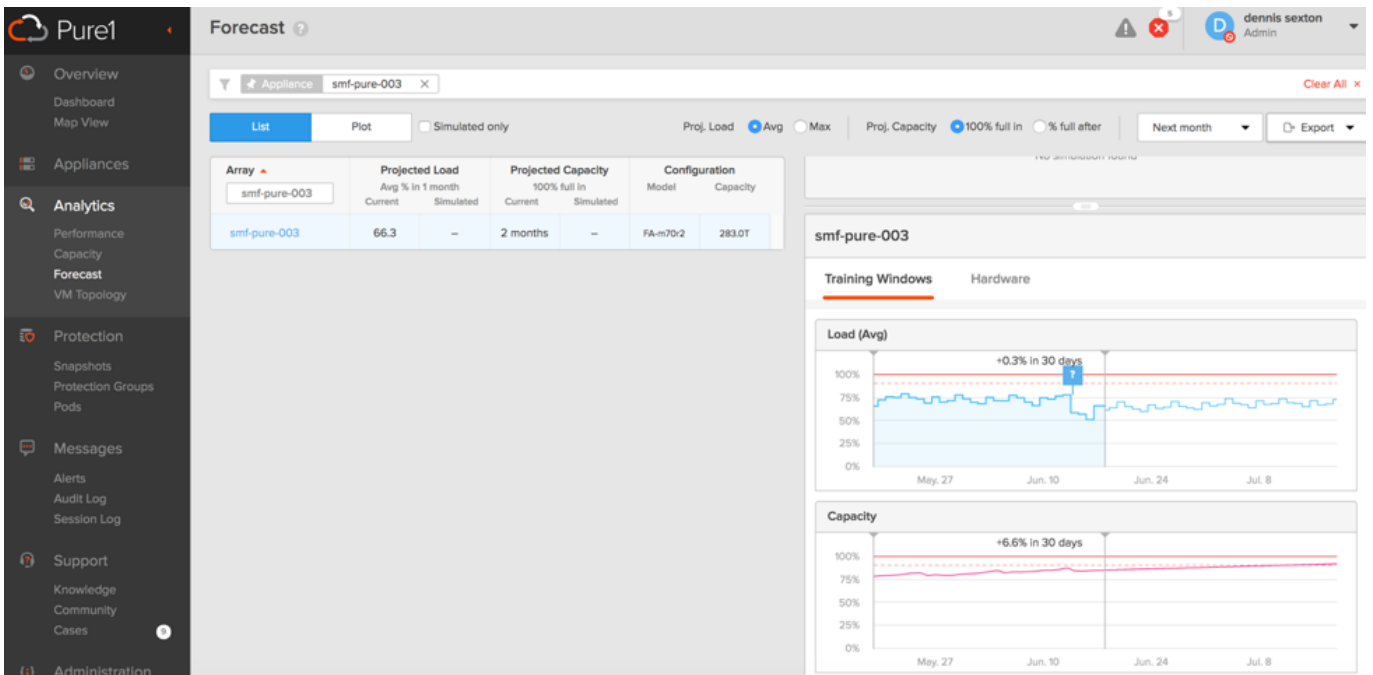
With Pure1® Analyze™, you get a wealth of simple-to-understand metrics on the health, capacity and performance of your infrastructure.

You can see all your snapshot information in one place, as per the following example:

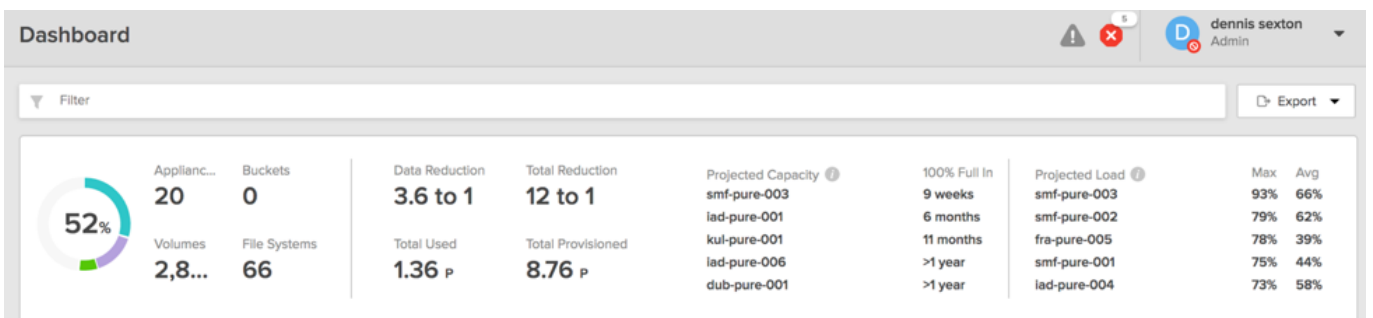
The screenshot displays the Pure1 Snapshots management interface. The top navigation bar includes the Pure1 logo, the 'Snapshots' title, and user information for 'dennis sexton Admin'. The main content area is divided into 'Volumes' and 'File Systems' tabs, with 'Volumes' selected. A filter bar and a date range selector (Last 30 days) are present. The primary table lists snapshots with columns for Last RPO, Array, Volume, Total Size, and a calendar grid for dates from May 20 to Jun 17. Below this is a detailed table of snapshots with columns for Snapshot, Size, Target, Target Type, Protection Group, and Created On.

Snapshot	Size	Target	Target Type	Protection Group	Created On
dub-pure-001-infra.3073.dub-pure-001-L...	637.43 K	dub-pure-001	Local	dub-pure-001-infra	2019-06-13 20:36:30
dub-pure-001-infra.3067.dub-pure-001-L...	393.93 K	dub-pure-001	Local	dub-pure-001-infra	2019-06-13 14:36:30
dub-pure-001-infra.3061.dub-pure-001-L...	538.35 K	dub-pure-001	Local	dub-pure-001-infra	2019-06-13 08:36:30
dub-pure-001-infra.3055.dub-pure-001-L...	1,009.32 K	dub-pure-001	Local	dub-pure-001-infra	2019-06-13 02:36:30
dub-pure-001-infra.3049.dub-pure-001-L...	878.25 K	dub-pure-001	Local	dub-pure-001-infra	2019-06-12 20:36:30
dub-pure-001-infra.3043.dub-pure-001-L...	513.98 K	dub-pure-001	Local	dub-pure-001-infra	2019-06-12 14:36:30

You can predict performance and capacity requirements through Pure1 Meta workload analysis, using trending and forecast to produce accurate estimates of load and capacity growth:



A dashboard allows you to see all Pure arrays, volumes and their growth deltas to allow you to see at a glance what is consuming storage, in case the size of the growth is concern:



The Pure1 Meta Workload Planner predicts both capacity and performance, and provides intelligent advice on workload deployment, interaction and optimisation:

