

Snapshots and Automated Storage Tiering

Introduction

Snapshots are an essential feature of modern SAN/NAS technologies and are A MUST for primary virtual machine storage because all virtual machine backup applications use Logical Unit (LU) snapshots as the first stage of creating VM backups. Nevertheless storing all the snapshot history on the Tier 1 may be challenging.

Problem


All the conventional storage systems that utilize snapshots have two problems: wasted space and degraded performance.

While data (i.e. VM images) doesn't consume that much space, snapshot history takes up a lot of space.

When talking about having the need to run snapshots, the disk space on the primary storage can be used in the following ways:

- RAW – total capacity that storage system has;
- Usable – space available to user after all the space reductions due to one or another reason (RAID, Synchronous Mirroring, Erasure Coding etc.);
- Space for snapshots – part of usable space reserved for snapshots.

This results in a great waste of expensive disk space that turns critical when going from spindle hard drives to flash storage.



All data is evenly distributed on flash

The problem is partially resolved by using Information Life-Cycle Management (ILM). One of the examples of ILM is Automatic or Manual Tiering. In some systems there is a separate lower cost pool, where the snapshots and other "cold" data reside. However Tiering takes time, no matter if it's automated or manual. Tiering also creates an additional load on the hardware resources because data

gets copied from one tier to another constantly, which will significantly decrease the performance and flash lifetime.

Snapshots are usually implemented by either Copy-on-Write (COW) or Redirect-on-Write (ROW). They are characterized as follows:

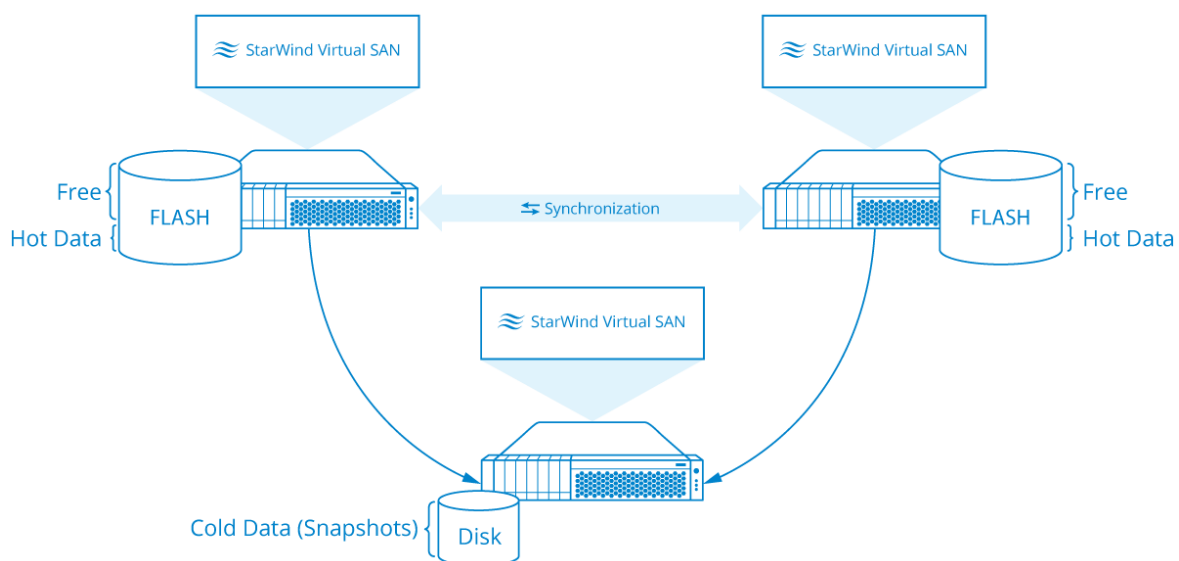
- COW snapshots utilize read-modify-write sequence. To put it simple, when new data blocks come, the system moves the original data block to another location, and after, the new block is written instead of the original one. The system has to perform 3 operations (one read and two writes), instead of one write. Obviously that creates additional load on the hardware involved in the snapshot related processes, which results in performance decrease.
- ROW creates fragmentation since all the new data is written in another place on the disk, instead of putting it sequentially. That results in random reads and writes, and performance decrease. Also, ROW stores all the snapshot history, thus consuming the disk space.

To make a long story short, the COW kills performance and ROW boosts performance using some more disk space.

Solution

StarWind brings an extra node with inexpensive, high capacity SATA spindles to serve as dedicated Snapshot Storage. The “cold” data is automatically moved to the snapshot tier, represented by the separate storage node. Thus, the primary storage keeps only “hot” data which allows increasing the amount of free usable space on primary expensive highly performable storage.

Instead of keeping “cold” data on all-Flash Primary storage, it is simply offloaded to the cheap SATA dedicated storage.



Extra node for Snapshot Storage with inexpensive SATA

Conclusion

As a result of using StarWind solution:

- The usable space on primary high performance storage system is increased since all the “cold” data resides on dedicated inexpensive storage;
- Absence of performance degradation achieved by offloading to the disk subsystem;
- Life time of SSDs is prolonged significantly since all major decrease in amount of write operations;
- Redundancy and disaster recovery is improved by storing snapshot history physically separately from primary storage.

Implementing StarWind Virtual SAN as the dedicated snapshot storage helps to avoid the problems that conventional solutions have: waste of space and a decrease in performance. StarWind allows increasing the effectiveness of the disk space usage on the primary storage, stop performance from going down and increase the life time of the flash drives.

In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms”.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.