DELL EQUALLOGIC

SNAPSHOT BASICS – DATA COPIES CREATED IN THE STORAGE ARRAY

Snapshot implementations are not all the same – be sure to understand your needs and the features and restrictions of any snapshot solution.

Copying data is a routine part of IT operations – administrators copy data for backups to protect it, to share it, to deploy replicated systems, to test applications, and to move data among servers when redeploying applications. Various tools are used to both create and maintain these data copies, but common challenges include the time it takes to create data copies, the disk space required, and the impact on server operations during the copy process. Many storage arrays support the ability to create data copies within the storage array – this can dramatically speed the copy process, improve the integrity of the data, and reduce the disk space required for temporary copies, and in SAN configurations provide the ability to change data access between servers without performing a copy operation at all.

Storage arrays support many methods of creating and maintaining copies, some of which depend on whether the copy will be permanent (a full copy) or temporary (usually a differential copy). Their intended use also defines the amount of storage needed – full copies need additional disk space equal to the data object being copied, while temporary copies can operate with much less space (depending on the IT function and storage array features).

In storage arrays, copies can also be maintained at the same or different location than the primary data. Remote copies are commonly used to enhance data protection – having a copy at another location protects against a physical disaster. Remote copies can also be used for application and data migrations between datacenters.

TYPES OF SNAPSHOTS AND HOW THEY ARE CREATED

Snapshots are volume (LUN) level copies of data made at particular points in time. As snapshots can be applied in various manners, for this discussion we'll reference them as one of two families of snapshots: differential or full copy.

Differential snapshots allow for fast creation and reduced disk space consumption. Some implementations of differential snapshots include copy-on-write or allocate-on-write; these methods are based on sharing common data and storing the differences unique to the snapshot or volume. Better implementations create these copies instantly, allow the copies to be used read-write, permit many copies to co-exist and be active at the same time, and can return the snapshot data to the original volume instantly. Backup, disaster protection, and testing are common uses of differential snapshots.

Full copy snapshots consume the same amount of storage as their original volumes and are implemented in three variations. The "break-the-mirror" method creates a copy by breaking off a disk mirror in a RAID environment. The clone method creates a complete copy of the data in a background operation at the time of snapshot creation. In both cases, the creation is instant, but full copies have the disadvantage of requiring additional work – either beforehand to create the mirror or afterwards to copy and create the clone. The replica method is similar to a clone, except that the data is copied to a remote system over a network, usually geographically remote for disaster protection.

CAUTION: NOT ALL SNAPSHOT IMPLEMENTATIONS ARE ALIKE

Your snapshot capabilities will depend on your storage, backup software, and OS/applications working together. Before selecting a snapshot solution, identify your needs and understand the solution's details and restrictions. Below are some criteria for evaluating snapshots.

1. Disk space efficiency. Differential snapshots consume only the differences from the original volume, while full copy snapshots use as much space as the original.

Snapshot Criteria	EqualLogic PS Series Array Features
Disk space efficiency	Uses differential for maximum space efficiency
System performance impact	Minimal due to caching and allocation mechanisms
Read only or read-write	Supports multiple read / write snapshots per volume
Remote copy	Supports both local and remote snapshots
Cost	No licensing fees - snapshot capability standard in all arrays
Number of snapshots	Supports up to 512 snapshots per volume, 10,000 total
Copy frequency	Unrestricted for both differential and full copies
Reinstating	Immediate reinstatement with no impact to other snapshots
Space allocation	Guaranteed, and customer can change allocation at any time
Convert snapshots to volumes	Supports converting differential copies to full copies

2. System overhead (performance impact). Snapshot implementations affect system performance, and some are dependent on the number of active snapshots.
3. Read only vs. read-write. Some snapshot implementations are read-only – these cannot be used for testing or parallel processing.
4. Remote storage. Remote snapshots are helpful for disaster protection or application movement between data centers, but not all solutions offer that capability.
5. Cost. Snapshot functionality is not always affordable, as many vendors charge high licensing fees.
6. Restrictions. Many products support a limited number of simultaneous snapshots.
7. Frequency. Implementations that support only full copy snapshots restrict the frequency of copy operations.
8. Reinstating. Not all solutions enable you to place a snapshot back into the original volume to immediately resume production. In addition, some automatically delete all other snapshots when you reinstate, causing loss of data that is needed for data protection or diagnosis.
9. Guaranteed space allocation. Snapshot disk space should be guaranteed and controllable. Some products over-subscribe space – for operations that depend on the snapshot, this can result in unpredictable operations, and future snapshots may fail due to lack of space.
10. Ability to convert differential copies to full copies. The ability to convert temporary copies to full copies provides the flexibility in test environments to have data copies be made permanent.

A well-implemented storage solution can efficiently create, manage, and maintain snapshots without causing significant additional overhead in management or performance, making it a less expensive copy process. Be certain that the solution you select matches your requirements.

Snapshot implementations are not all the same – be sure to understand your needs and the features and restrictions of any snapshot solution.

EQUALLOGIC

110 Spit Brook Road, Building ZKO2, Nashua, NH 03062
Tel 603.579.9762 / Fax 603.579.6910 / www.equallogic.com

CB109_USA_022205–012808