



SNAPSHOTS

TECHNICAL CAPABILITY SHEET



open
DRIVES



Capability Definition

Atlas can generate immutable read-only snapshots of the filesystem at any point in time. With a snapshot, the user can roll back to a previous state of the filesystem and can even allow extraction and recovery of individual files from a previous filesystem state. The cadence of snapshots is user-configurable but is limited to a maximum of one snapshot per minute. However, you can retain as many snapshots as necessary depending on operational best practices, though snapshots do become larger the more files change within the filesystem, necessitating a sensible process of purging snapshots.

Note: Snapshots are lightweight and take up very little storage space, usually kilobytes to megabytes in size initially. The reason for this is that snapshots do not contain actual copies of the data they reference but rather point to that data and the delta of change in the filesystem between subsequent snapshots. Furthermore, snapshots require very little computational resources, making a negligible impact on overall storage system performance. The time it takes to create a snapshot is trivial (within milliseconds).

Outcome & Value

Snapshots provide the administrator with peace of mind when upgrading applications, making large-scale environmental changes, or even when dealing with error-prone users who accidentally delete files and need to recover them. While a read-only structure, a snapshot provides the ability to roll back the entire filesystem to a previous state or to restore individual files that existed without error in previous snapshots. For business continuity, snapshots are a sensible precaution by enabling a quicker restore to a previous filesystem state than a full-blown backup-and-restore procedure. Snapshots, however, are complementary to enterprise backups rather than a direct replacement. Working together, snapshots and backups are part of a comprehensive disaster recovery solution for catastrophic situations like a ransomware attacks.

How It Works

A snapshot captures the state of the filesystem at a certain point in time (PIT). A snapshot is both fast and efficient because it does not create a copy of all filesystem data but rather records all the references of and pointers to data which have changed since the previous snapshot's PIT capture. A snapshot essentially is the amount of change (also known as the difference or delta) since the previous snapshot was taken. Keep in mind that as time goes

on and this chain of snapshots grows, the increased amount of information will require more storage space, but each individual snapshot is small and rapidly created. This is the reason that snapshots are an excellent short-term recovery solution, but should work in tandem with larger-scale enterprise backups.

The filesystem within Atlas is basically a tree structure starting with a root block, or uberblock. An initial snapshot of a new filesystem contains virtually no information because no delta has occurred between snapshots. With subsequent snapshots that capture changes at the filesystem level (new directories, new files, modified files, etc.), Atlas records the state of the live filesystem and corresponding tree structure at that PIT, what essentially amounts to a collection or journal of references and pointers that indicate differences since the last snapshot.

This tree structure changes and evolves with a live filesystem. To understand this change over time, you need to be aware that Atlas incorporates a copy-on-write process, meaning that it writes data (for example, changes to an existing file) to new blocks rather than overwriting the existing ones. Once it writes new blocks, it must rewrite all the references and pointers up the tree structure all the way to the uberblock (incidentally recalculating checksums along the way). A snapshot records all these live state changes that are housed in the storage system's RAM memory.

Characteristics

The snapshots capability has the following characteristics:

- Creates an immutable read-only snapshot of the filesystem at any point in time.
- Enables extraction of single files or a complete roll back of the entire filesystem at the point in time of the snapshot.
- Has negligible effect on storage system performance.
- Much faster restore than enterprise backup solutions.
- Contain no actual data, only pointers and references, leading to very lightweight snapshots.
- Using Atlas' routines engine, users can automate configurable snapshot schedules with up to one minute granularity.
- Complements an enterprise backup solution for a more comprehensive disaster recovery plan.


Further Reading

The checksum and snapshot capabilities are closely related because full checksumming means that the entire tree structure recursively checksums branches. Please refer to the Checksum Capability Sheet which explains in more detail the mechanics behind the way in which Atlas Core provides data integrity through checksum calculations.

The following diagram illustrates the basic process of capturing snapshots of the live filesystem at a single point in time—a snapshot only captures the delta (change) from the previous snapshot, which makes it lightweight and not burdensome on storage system performance:

Snapshot 1


Captures initial filesystem – no change



New filesystem

Snapshot 2


Captures changed blocks in green



New directories and/or files,
and/or existing data modifications

Snapshot 3

Captures changed blocks in orange



New directories and/or files,
and/or existing data modifications

Time →

The important fact to remember is that snapshots are sequentially generated over time. Each subsequent snapshot captures only the differences in data change since the previous snapshot. This feature keeps them lightweight and quicker to generate.