



PREFETCH

TECHNICAL CAPABILITY SHEET



open
DRIVES



Capability Definition

Prefetch is a specific type of data retrieval and caching function that Atlas provides by leveraging ZFS file system capabilities along with our own tuning and optimization code. Prefetch takes advantage of both the contextual awareness of requested data and the system's cache to fetch more data than is explicitly demanded. It eliminates associated lag time and disruptive performance during information retrieval because Atlas can read data much faster system cache than from slower storage media types.

Note: The ZFS file system is not inherently tuned and optimized for video- and audio-based media. OpenDrives spent considerable R&D tuning Atlas specifically for large-file media, such as video files, so that working with this type of data is as fast as possible. Lastly, OpenDrives' prefetch tuning was arrived at through research and industry experience, therefore unique to Atlas.

Outcome & Value

OpenDrives' ZFS tuning expertise and quicker data retrieval from cache results in the enhanced, accelerated performance and responsiveness we provide. By maintaining awareness of user/application requests and continual pattern analysis, Atlas can anticipate the flow of requested data. We automatically predict data that will be needed after the explicit request, then use that prediction to prefetch a window of data that anticipates the user's requirements. End users can expect little-to-no lag time within enterprise applications when accessing data.

How It Works

The prefetch capability is based on Atlas's ability to detect any linear access patterns, including forward or backward access. Once it detects an access pattern (contextual awareness), such as scrubbing back and forth through a timeline, the prefetch capability then leverages system

cache to store the most likely subsequent data to be requested around the currently accessed information. If the access pattern continues as anticipated, data is aggressively pre-staged to a much faster cache layer from slower storage media (such as spinning disk) in preparation for future requests. If the access pattern shifts to another portion of data, the whole cycle repeats with the new information.

Premiere-Specific Prefetch: With specific applications like as Adobe Premiere Pro, we make accessible an application-level cache that uses NVMe to provide faster data access for the prefetch capability. For example, we typically allocate 4 of 8 NVMe SSDs specifically to act as cache for prefetching. This area can store all cache files, working directories, and project files. This is a recommended best practice to accelerate Adobe Premiere Pro performance, especially in workflows where users often move workstations and need to rebuild projects in local cache often.

Characteristics

The snapshots capability has the following characteristics:


· The prefetch capability has the following characteristics:

- Block-level (not file-level) caching
- Contextual awareness senses an access pattern, initiating prefetching data in a window around explicitly requested information, such data forward and back from a current timeline read-head location
- Window size from 4K to 2MB per fetch cycle (2MB by default)
- Up to 8 maximum prefetch streams per file
- Jumping around without establishing a clear access pattern deteriorates prefetch performance

Further Reading

As previously stated, prefetch is a type of predictive data caching within Atlas Core. For more information about caching in general, please refer to the Caching Capability Sheet which explains in more detail the mechanics behind our data caching capabilities which we've implemented within Atlas Core.

The following diagram illustrates the prefetch process. The diagram displays a requesting application, the storage media, the requested information (R), the 2MB window around R, and prefetched data (PF).



When the user or application requests information stored on disk, Atlas Core/ZFS begins caching to memory the explicitly requested data and a 2MB window around it. As the access pattern becomes established, all data is read from the much faster cache rather than directly from storage.