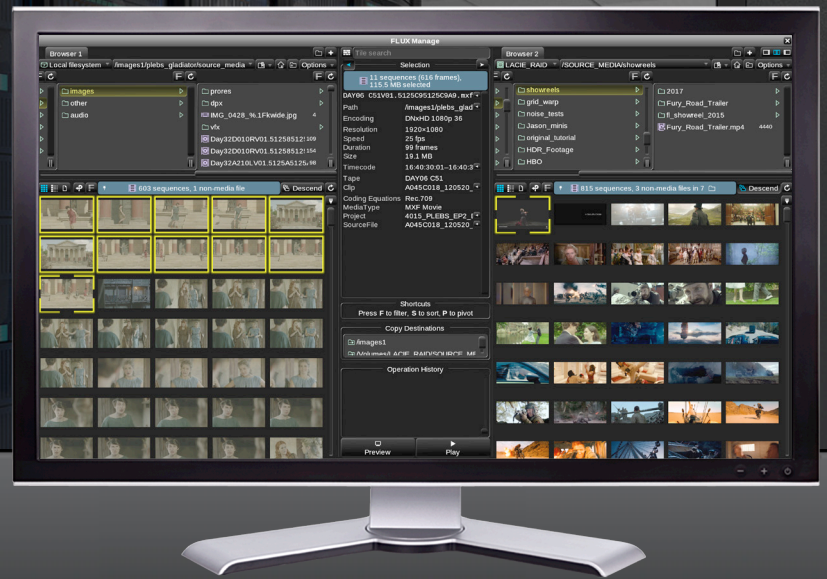# FLUX Index

Standalone server for fast metadata searches

FilmLight

More storage should not mean slower access.

The advent of high capacity SAN and NAS storage systems, combined with modern shoot ratios, has meant that the amount of material to sift through in post-production has itself become a bottleneck.

Even the simple task of running a well-defined EDL conform against a standard filesystem can itself be laborious - with possibly hundreds of RAW or compressed movie files that each require a third party camera manufacturer SDK (RED, ARRI, Sony...) to be invoked in order to extract the necessary metadata.

One solution, used by many manufacturers, is to create a proprietary filesystem where all of the metadata is extracted as assets that are checked in to the storage pool. This works well for the specific use case they are designed for, where this 'index' is accessed by their own applications to quickly retrieve metadata for operations like conform.

But what happens when you need a collaborative storage solution? Although these systems may present with standard filesystem semantics, using them collaboratively across multiple vendor applications typically acts through a 'spoofed' software layer, resulting in poor performance.

Enhanced collaboration through the metadata index

A more collaborative approach is to maintain a metadata index of any type of filesystem, whether it is proprietary or open. FilmLight has developed just such an agnostic approach for its own Baselight and FLUX Store products.

Introducing FLUX Index

FLUX Index is a standalone product that utilises the same logic FilmLight uses to maintain and index its own XFS filesystem-based products, but can be run against any filesystem. Physically, FLUX Index is a 1U Linux rack server with optional fibre channel or Ethernet interfaces to the storage architecture of your choice, typically NAS or SAN.

The server maintains a relational database on SSD of all mounted volumes. It contains all relevant post-production metadata relating to image sequences; for example, timecode, tape name, frame rate, resolution, ISO and so on. FLUX Index also decodes and stores a poster frame from each discreet sequence to aid fast visual browsing.
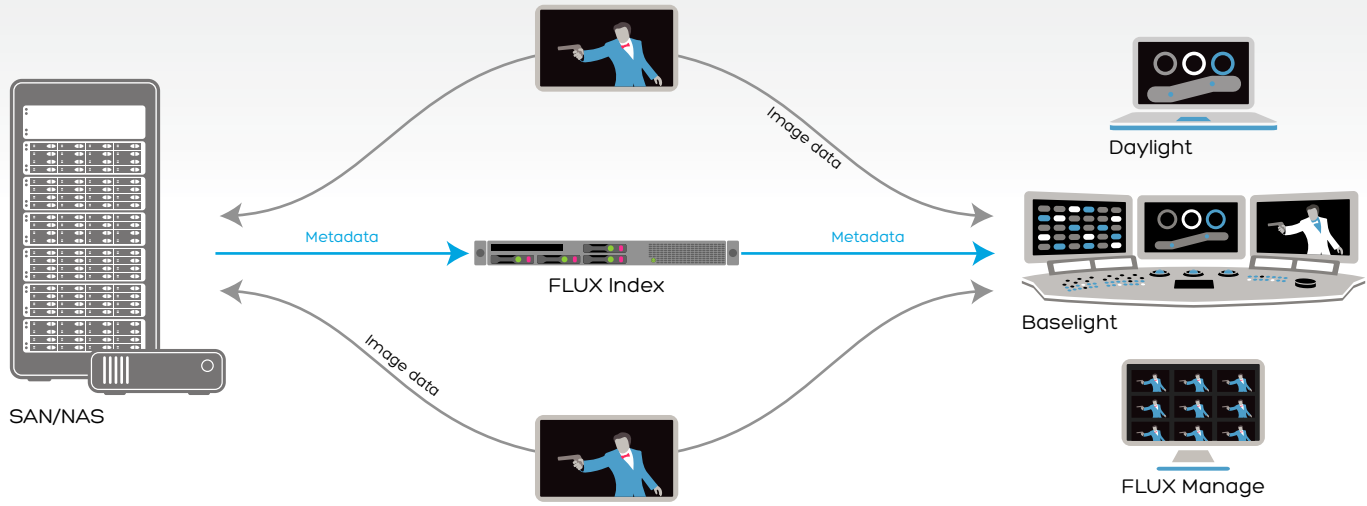
Access to third party storage is now branched at the FilmLight application level, with Baselight, Daylight and

FLUX Manage all able to access metadata and thumbnail images quickly through the FLUX Index server. This is completed without accessing any files on the primary storage, whilst standard filesystem operations, like reading and writing image data, are performed through the normal direct path (e.g. locally connected Ethernet or fibre channel).

The result is almost instantaneous searches across multiple volumes containing petabytes of data, with EDLs conformed rapidly using the same underlying architecture.

Low indexing overheads

Although FLUX Index is a FilmLight product designed to accelerate FilmLight applications, the filesystems themselves remain open and unchanged and can be accessed and modified by third party applications normally as if the index did not exist. The index maintains itself by comparing directory hierarchy fingerprints (uid/gid/mode/size/mtime) and rebuilding a subtree lazily on demand when these have changed. In real post-production - where a good portion of the online storage is largely static for the duration of a project - this small indexing overhead typically occurs quite infrequently.

FLUX Index

SAN/NAS

Metadata

FLUX Index

Image data

Image data

Metadata

Daylight

Baselight

FLUX Manage

Scalable implementation cost - without per seat licensing

FLUX Index runs on a FilmLight-supplied and -specified HP DL20 1U server. There is a one-off upfront hardware cost that includes installation and setup, coupled with an annual subscription for the index service along with the normal FilmLight support contract. The subscription cost is capacity tiered in bands as follows:

»   < 0.25 PB

»   < 1.0PB

»   < 4.0PB

There is no additional cost for indexing multiple volumes under the total storage capacity nor for number of client seats.

A second FLUX Index server can be configured to mirror the index database for redundancy in high availability environments such as broadcast.

FLUX Index server

Specifications

»   HPE ProLiant DL20 Gen9 server

»   Dual 900W hot plug redundant power supplies

»   2 x 1Gb network ports

»   2 x PCIe 3.0 slots for storage HBA

»   Operating system: 512GB SSD

»   FLUX Index: 1TB SSD

»   CentOS 7.5

»   Dimensions (HxWxD): 4.32 x 43.46 x 38.22 cm

»   Weight (approximate): 8kg

FilmLight