

INDEXLESS PARTIAL FILE RESTORE


Restore the clip, not the whole master.

Optimized for Quantum ActiveScale and other S3-compatible media archives where restore time, bandwidth, and egress charges drive costs up.




INDEXLESS PARTIAL FILE RESTORE FOR S3 COMPATIBLE ARCHIVES

Typical File Restores



Prepared Workflows



Indexless Workflows




Table of Contents


A concise sales narrative for turning existing cloud archives into clip-addressable media archives.

01	Customer pain	Full-object restores waste time, bandwidth, storage, and budget.
02	Prepared and Indexless PFR	Proxy/index workflows for new ingest; indexless workflows for existing archives.
06	Quantum ActiveScale fit	S3 Glacier restore range access plus Drastic media intelligence.
07	Proven PFR	High-volume production restore workflows in use for over a decade.
08	Pilot business case	Compare bytes read, time to usable file, compatibility, and cost impact.

Core message: *Object storage can move byte ranges - Drastic turns that into media-aware partial file restore.*

The Customer Pain

Typical Partial File Restores



Archived media is massive.

Restore requests are small.


A producer needs a few minutes. The archive holds a full master file. Traditional restore pulls the whole object first, then trims it later.

THAT MEANS UNNECESSARY:

- Cloud/object restore time
- Network bandwidth
- Temporary storage
- Egress or retrieval cost
- Delay before editorial can start

Why restore the whole file when the user only needs the clip?

Prepared Proxy / Index Workflow




Best for new ingest, MAM workflows, review/approval, and browse/search.

DRASTIC CAN CREATE PROXY / INDEX FILES WITH

- Multichannel audio
- Metadata
- Closed captions
- Frame-accurate timecode
- Accurate start/end selection for PFR

Prepared workflows remain ideal when you control ingest and need browse/search capability.

Indexless Workflow for Existing Archives



Immediate PFR value without pre-indexing the archive.

Drastic reads only the minimum headers, footers, tables, and media structure needed to understand the source file, then restores only the required media ranges.

PRIORITY WRAPPERS


- MXF
- MOV
- MP4
- LXF
- GXF

Additional file types can be added for qualified workflows.

Turn an existing S3 archive into a clip-addressable archive without restoring and indexing every file first.

The New Breakthrough: Move Less Data


Read only the media structure that matters. Restore only the media ranges requested.



Subsequent requests for often used media can be served faster because the required structure and retrieved ranges may already be known or cached by the workflow.

We support the traditional proxy/index model — but no longer require it to get PFR value from existing archives.

What the Customer Gets Back



A real production file, not a broken fragment.


Drastic can create a new wrapped file with the same file timing while copying audio, video, metadata, timecode, and captions losslessly where supported.

OR TRANSCODE ON THE FLY TO

- MOV ProRes
- MXF DNxHR / DNxHD
- AVC-I
- XAVC
- XDCAM
- Other house formats

The editor gets a usable file immediately — not a partial object that still needs engineering work.

Quantum ActiveScale and S3: Why It Fits



Object storage can deliver byte ranges. Drastic knows which byte ranges matter.


Quantum ActiveScale and other S3-compatible systems can support range-based access or restore behavior.

DRASTIC MEDIA INTELLIGENCE

- Understand the wrapper
- Locate the requested timecode range
- Restore only required media data
- Rewrap or transcode into a usable deliverable
- Avoid costly full-object restore wherever possible

S3 can move ranges. Drastic turns that into media-aware partial file restore.

FastClip BitCache Feature



One Request Caches the Bits.

When users require many short clips from a game, only the first request gets the file metadata. Subsequent requests use the BitCache to find and retrieve just their clips. This provides the efficiency and speed of pre-indexing.

ON-DISK CACHE KEEPS METADATA LOCAL

Elements prioritized based on:

- How old the elements are
- How large the assets are
- How long the assets took to retrieve

Save time and money with faster restore requests.

Proven PFR at Production Scale



This is not a lab feature.


Drastic PFR has been used in high-volume production for over a decade.

Thousands of restores every day

For 11 years at a high visibility U.S. sports-media archive.

Your archive should be making you money, not causing your staff headaches.

Close: The Business Case



Drastically Simple: swift restores that save budget.

Indexless PFR helps large media archives avoid the usual blocker: "We would love PFR, but we cannot afford to pre-index petabytes of existing media."

Start with the archive you already have.

Compare full restore versus Drastic indexless PFR on real assets and quantify the savings.

Give us the archive you already have. We will show how much less you need to move.

Pilot Proposal

A low-risk proof using real archived assets and measurable outcomes.

10–20

real archived assets

2

restore methods compared

4

business metrics

1

go / no-go decision

TEST DESIGN

- Select representative archive files and requested timecode ranges
- Run full restore and Drastic indexless PFR against the same assets
- Validate output compatibility in the production toolchain

MEASURE

- Bytes read
- Time to usable file
- Output compatibility
- Cost impact

The pilot will prove the savings before broader deployment.

Next step: Prove it on your archive

Indexless PFR for Quantum ActiveScale and S3-compatible media archives

10–20

assets

4

pilot metrics

1

usable deliverable

Measure bytes read, time to usable file, output compatibility, and cost impact. Use the results to quantify the business case before broad deployment.

Restore the clip. Move less data. Lower archive costs.