

ION VIDEO LIMITED

Technical Whitepaper

Video Virtualisation Architecture

How ION transforms rendered video into an addressable, orchestratable data layer for intelligent systems.

CLASSIFICATION

Technical Deep Dive

AUDIENCE

CTOs, Principal Engineers, Cloud Architects

STATUS

Confidential

EXECUTIVE SUMMARY

ION's video virtualisation architecture transforms rendered video from static, file-bound media into an addressable, time-indexed, orchestratable data layer that intelligent systems can assemble, personalise, and deliver progressively in real time, without re-encoding or duplicating source media.

Traditional video infrastructure treats compressed media as monolithic assets. Even within ISO Base Media containers where samples are technically indexed, the assembly map is static, generated once at encode time, binding structure, sequencing, and timing permanently to a single playback experience.

ION generalises and externalises this assembly layer.

Rather than treating the MP4 container as the final authority of structure, ION virtualises encoded video samples into a programmable reference space. Each compressed frame or segment becomes independently addressable, with timing, ordering, and inclusion controlled dynamically by an orchestration engine rather than a fixed sample table.

ARCHITECTURE

How the Architecture Works

PROGRAMMABLE ASSEMBLY

This enables intelligent systems to construct video experiences on demand. AI models can select, sequence, and time video components in real time based on user context, intent, behavioural signals, or inference outputs. Instead of rendering new video files, the orchestration layer computes a personalised assembly graph that maps virtualised references to underlying encoded media.

PROGRESSIVE DELIVERY

Progressive delivery occurs as the assembly is resolved. The system can begin streaming immediately while subsequent segments are dynamically selected and scheduled, enabling real-time personalisation, adaptive narratives, contextual highlights, summaries, and AI-driven content composition without re-encoding or duplicating media.

DECOUPLED INFRASTRUCTURE LAYERS

From a systems perspective, ION decouples the following into independent, programmable layers:

- Encoded media storage
- Structural assembly
- Temporal orchestration
- Delivery sequencing

This collapses rendering pipelines, eliminates exponential storage growth, and enables AI-native interaction with video as a composable data substrate rather than a static asset.

Core architectural principle

ION decouples four previously bound layers: encoded media storage, structural assembly, temporal orchestration, and delivery sequencing. Each becomes independently programmable. This collapses rendering pipelines, eliminates exponential storage growth, and enables AI-native interaction with video at infrastructure cost profiles not practically achievable with traditional rendering models.

Crucially, the approach remains compatible with existing codec ecosystems and streaming infrastructure. Underlying media remains standard compressed samples. What changes is the control plane that governs how those samples are referenced, assembled, and delivered.

TECHNICAL FOUNDATION

HOW ISO BASE MEDIA FUNCTIONS AS AN ASSEMBLY FRAMEWORK

Understanding ION requires understanding what MP4 and ISO Base Media Format actually are. They are not simply file formats. They are time-indexed assembly frameworks for compressed media.

The encoding pipeline

When a modern video encoder runs — whether H.264, HEVC, AV1, or similar — it does not create a movie file. It produces a continuous stream of compressed data units. Each unit represents the information needed to reconstruct a single moment in time. These units are called samples.

In almost all modern video workflows, one sample corresponds directly to one frame of video. Frame one becomes sample one, frame two becomes sample two, and so on. Each sample is simply a block of compressed bytes: no timing, no structure, just encoded data.

The movie file itself comes later.

Container architecture

MP4, MOV, and the wider ISO Base Media File Format are not codecs. They are structured containers built from a hierarchy of boxes, sometimes called atoms, that organise raw encoded data into a navigable, time-aware system.

Two are foundational:

mdat

The media data box. All encoded video data lives here. Inside mdat, samples are written sequentially, back-to-back: sample one, sample two, sample three. There is no inherent meaning inside mdat. No timing. No frame boundaries. No playback logic. It is simply a warehouse of compressed data chunks stored in order.

`moov`

The intelligence of the file lives here. This is where the sample tables live. The sample tables describe exactly how to interpret the raw bytes stored in mdat.

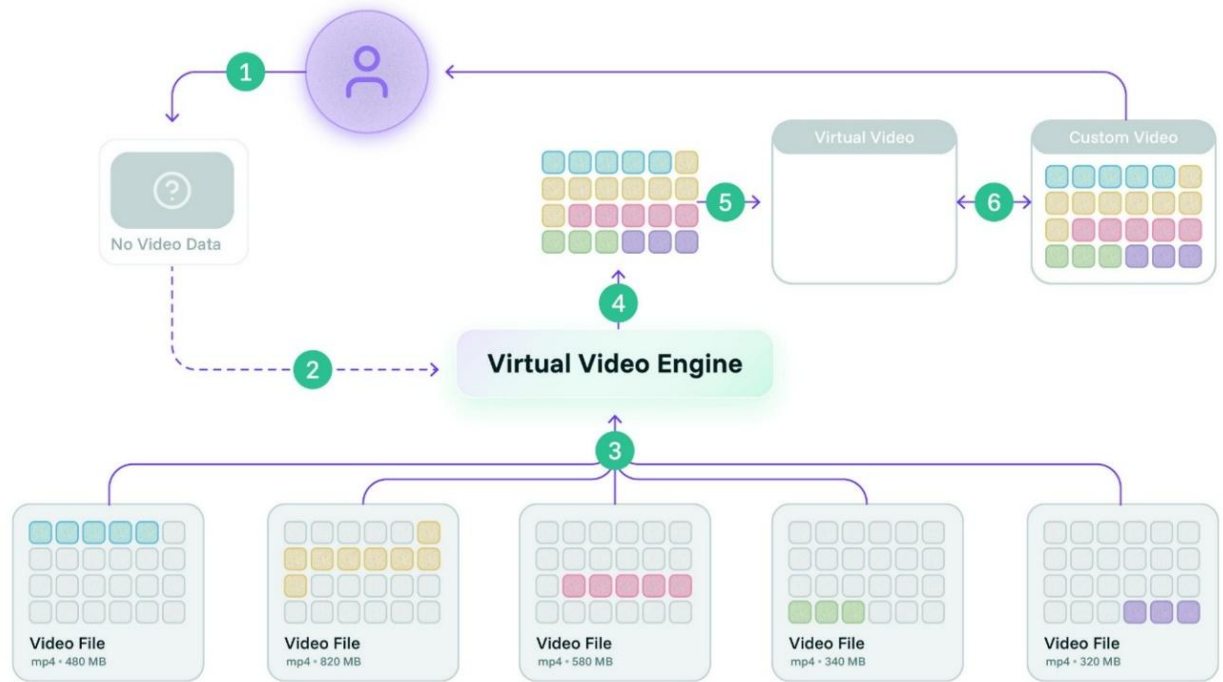
The sample tables answer the critical questions every player needs:

- Where does each sample start in the data stream?
- How large is each sample in bytes?
- How long should each sample be displayed?
- Which samples are keyframes?

Several tables work together to provide this map. One records the size of each sample, another records the byte offset of each sample inside mdat, another defines how long each sample lasts in time, and another flags random access frames.

Together, these tables define how encoded samples are located, timed, and presented.

<code>stco / co64</code>	Byte offset of each sample inside mdat
<code>stsz</code>	Size of each sample in bytes
<code>stts</code>	Duration of each sample in timescale units
<code>stss</code>	Flags identifying random access keyframes
<code>ctts</code>	Composition offset for decode vs. presentation reordering



- 1 User Plays Virtual Video
- 2 Virtual Video Calls ION
- 3 ION Extracts Samples from external video files and assembles sample order for delivery
- 4 ION Delivers Organized Sample Blocks
- 5 Sample blocks assembled into Virtual Video
- 6 Video becomes complete

The timescale clock system

MP4 does not store time in frames per second; instead, it uses a clock system based on time units called a timescale. The media declares how many ticks equal one second.

EXAMPLE — 25 fps at timescale 25,000

$$\begin{aligned} \text{timescale} &= 25,000 \\ &= 25,000 \text{ time units} = 1 \text{ second} \end{aligned}$$

Each video sample is then assigned a duration in those units. At 25 frames per second, there are 25 samples per second, so each sample must last:

$$25,000 \div 25 = 1,000 \text{ time units per sample}$$

The timing table simply records that each sample has a duration of 1,000 ticks. After 25 samples, the accumulated duration reaches 25,000 ticks — exactly one second of playback.

EXAMPLE — 25 fps at timescale 90,000

Many professional systems use a timescale of 90,000. The math works the same way.

$$90,000 \div 25 = 3,600 \text{ ticks per sample}$$

Each frame lasts 3,600 time units. Accumulate 25 of them and you have one second.

The file never stores “25 fps”. It stores durations. Playback time is calculated by summing sample durations.

STANDARD VIDEO — FRAME RATE TO SAMPLE RATE

In standard video, one sample equals one frame:

Frame rate	Samples per second
24 fps	24 samples per second
25 fps	25 samples per second
30 fps	30 samples per second
60 fps	60 samples per second

Audio works differently, but video is almost always one compressed sample per frame.

THE ENCODING PIPELINE

Putting the whole system together, the pipeline looks like this:

1. Raw frames go into the encoder.
2. The encoder compresses each frame into a sample.
3. The muxer writes every sample sequentially into mdat.
4. Simultaneously, it builds the sample tables inside moov — recording for each sample its byte position, byte size, duration in time units, and whether it is a keyframe or B-frame.
5. Once complete, the MP4 file becomes a fully indexed, time-aware assembly of media.

THE BOOK ANALOGY

mdat is the entire book with all pages glued together into one long scroll. The sample tables are the table of contents: they tell you where each section begins, how long it is, and how long it should last in time. The content itself remains raw and untouched.

GROUPED RUNS

Timing inside the file is often stored in grouped runs. Instead of listing every frame individually, the timing table might specify eight samples, each lasting 3,600 ticks. In a 90,000 timescale system, that equals:

$$3,600 \div 90,000 = 0.04 \text{ seconds per sample}$$
$$= 25 \text{ frames per second}$$

Those eight samples together represent 0.32 seconds of timeline.

Again, the clock emerges purely from accumulated durations.

DECODING TIME VS PRESENTATION TIME

The format includes two related timing concepts: decoding time and presentation time.

Because modern codecs reorder frames internally, the order frames are decoded is not always the order they are displayed. The container handles this using an additional composition offset table.

But even here, everything is still expressed as time deltas in ticks. The same clock model applies.

In summary

mdat holds the compressed frames. Sample size tables define how big they are. Offset tables define where they live. Timing tables define how long they last.

Playback time is never stored as a master timeline. It is computed by summing the durations of samples as the player moves forward.

ISO Base Media is not just a file format. It is a time-indexed assembly framework for compressed media.

ARCHITECTURE AT A GLANCE

Component	Role in ION Architecture
mdat box	Warehouse of compressed frame data. Raw, sequential, untouched by virtualisation.
moov / sample tables	Assembly map. Defines byte positions, sizes, durations, and keyframe flags for every sample.
Timescale system	Clock mechanism. Playback time is computed by summing sample durations, never stored as fps.
ION reference layer	Externalises and programmes the assembly map. Samples become independently addressable across any source file.
Orchestration engine	Computes personalised assembly graphs dynamically. Delivers progressive streams without re-encoding.
Virtual Video File	The resulting primitive: a programmable, AI-addressable video data structure compatible with standard players and CDNs.

ION'S ORCHESTRATION CAPABILITY

ION treats the static sample table not as a fixed output but as a programmable reference space. Rather than treating the MP4 container as the final authority on structure, ION virtualises encoded video samples so their timing, ordering, and inclusion can be dynamically controlled by an orchestration engine rather than a fixed sample table generated at encode time.

Each compressed frame or segment becomes independently addressable. The orchestration layer computes a personalised assembly graph that maps virtualised references to underlying encoded media across one or more source files, resolving SMPTE timecodes (H:M:S:F) to sample locations, managing the container structure, format, sample mapping, and reference file locations simultaneously.

This is the technical nucleus of ION's architecture: assembling content from distributed encoded sources into a single virtual, orchestratable stream without re-encoding or modifying the underlying media.

The virtualisation principle

ION does for video what virtual memory did for compute and what object storage did for cloud: it introduces an abstraction layer that separates the logical structure from the physical storage. Rendered media becomes addressable, composable, and dynamically orchestrated by intelligent systems.

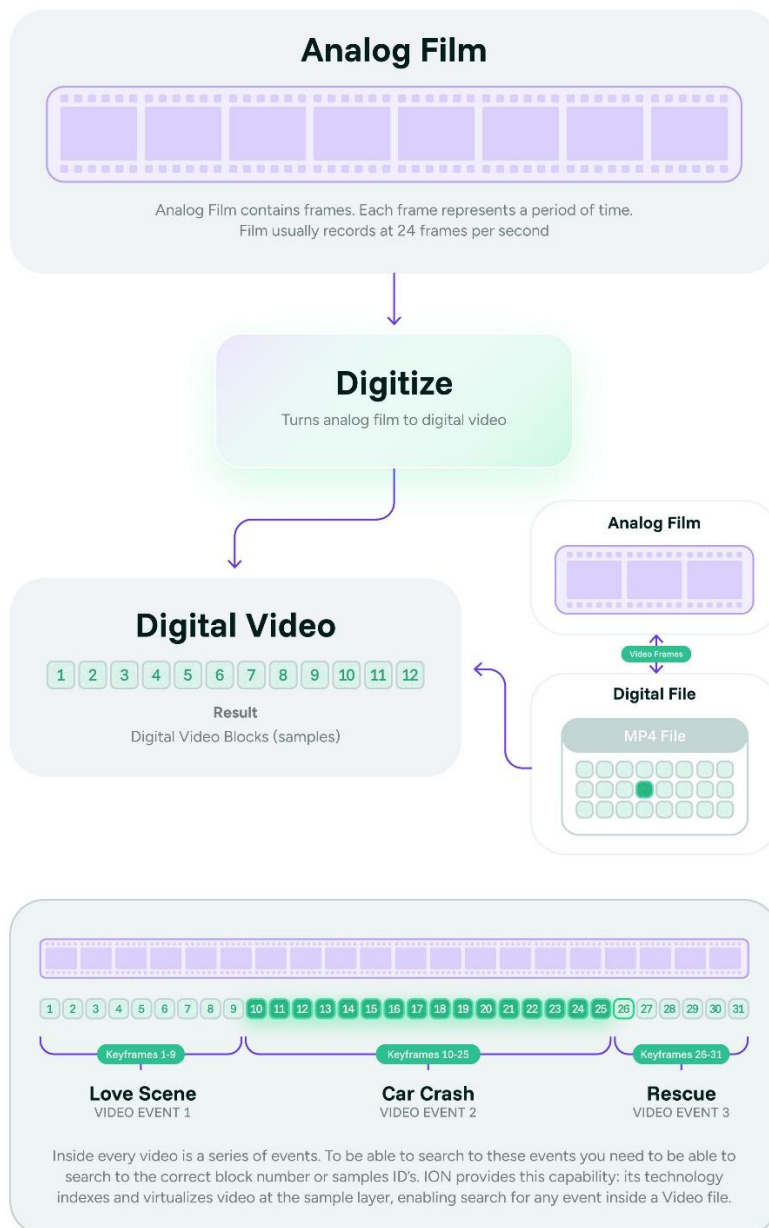
Four decoupled layers

The architectural consequence of ION's virtualisation is that four previously bound operations become independently programmable:

Encoded media storage	Underlying compressed samples remain standard. No re-encoding required. Existing codec ecosystems and CDN infrastructure are compatible.
Structural assembly	The sample tables governing which frames appear and in what order are computed dynamically by the ION orchestration engine at request time.
Temporal orchestration	Timing, sequencing, and narrative structure are controlled by AI inference outputs, user context, or behavioural signals rather than fixed timecodes.
Delivery sequencing	Progressive delivery begins immediately while subsequent segments are dynamically selected and scheduled, enabling real-time personalised streaming.

AI-NATIVE IMPLICATIONS

For AI platforms, ION's architecture creates a new primitive: video as an orchestratable, real-time computable medium. Intelligent systems do not generate new video files. They generate assembly instructions.



AI models can select, sequence, and time video components based on user context, intent, behavioural signals, or inference outputs. The orchestration layer resolves these instructions into a valid virtual assembly graph referencing existing encoded media. Progressive delivery of the assembled stream begins immediately.

This architecture unlocks capabilities that are not practically achievable with traditional rendering models:

- Scalable personalisation at infrastructure cost profiles, with no per-variation encode cost
- Contextual highlights, summaries, and adaptive narratives assembled in real-time
- AI-driven content composition across distributed source media without duplication
- Continuous adaptive video experiences that respond to evolving inference outputs
- Agent-driven content creation where AI orchestration replaces traditional editorial pipelines

The infrastructure consequence

Because no new files are generated and no re-encoding occurs, the storage and compute cost profile of video at scale transforms fundamentally. A single virtualised source can produce an unbounded number of distinct assembled experiences. The exponential growth in storage associated with file-per-variation rendering models is eliminated.

COMPATIBILITY AND INFRASTRUCTURE INTEGRATION

ION's approach remains fully compatible with existing codec ecosystems and streaming infrastructure. Underlying media is standard compressed samples stored in standard containers. What changes is the control plane that governs how those samples are referenced, assembled, and delivered.

Existing CDNs, players, and delivery infrastructure operate without modification. ION sits upstream as the data layer: it produces valid, standards-compliant virtual assembly outputs that downstream infrastructure handles identically to conventionally encoded streams.

For organisations operating at scale across content, AI, media delivery, or personalised digital experiences, this integration profile is critical. ION does not require infrastructure replacement. It adds a programmable assembly layer above existing encoded media stores, extending the capability of the entire downstream stack without disrupting it.

PATENT PROTECTION

The video virtualisation architecture described in this document is protected by four foundational granted patents, with an additional 2026 patent filing extending protection into the runtime control and governance layer. The foundational patent estate covers the core method of virtualising encoded video samples into an independently addressable reference space, enabling dynamic assembly without re-encoding.

US 8,893,203 B2	Core video virtualisation method. Generalisation and externalisation of the ISO Base Media assembly layer.
US 10,721,507 B2	Continuation. Extended claims covering orchestration and dynamic assembly workflows.
US 9,516,392 B2	Continuation. Additional claims covering progressive delivery and reference-based streaming.
US 9,544,657 B2	Content blockchain. Distributed ledger integration for virtualised media asset provenance and rights management.

In 2026, ION filed a further patent application in Australia and the United States, extending this protection into the control layer of Virtual Video. The filing covers token-governed runtime resolution, including authorisation, licensing, territorial conditions, consent parameters, transaction rules, and device or session validation at the moment of assembly.

This extends ION's protection beyond video virtualisation and dynamic assembly into the governance and commercialisation layer required for AI-native and agentic media environments.

CONCLUSION

ISO Base Media is not just a file format. It is a time-indexed assembly framework for compressed media. The separation between the mdat warehouse and the moov assembly map is a foundational architectural design choice. ION extends this design by making the assembly map programmable, dynamic, and AI-driven.

The result is a virtualisation layer for video that follows the same architectural logic as virtual memory for compute or object storage for cloud: an abstraction that decouples logical structure from physical storage, enabling a new class of infrastructure capabilities that the original system was not designed to provide.

For any organisation building AI systems, cloud infrastructure, or intelligent media workflows at scale, ION represents a foundational infrastructure primitive: the layer that makes rendered video programmable by intelligent systems.

ION: Video Superintelligence.

<https://ion.video>
