ATEME

BEST PRACTICES FOR BUILDING AND OPTIMIZING DVR PLATFORM USING THE PUBLIC CLOUD

WHITE PAPER

White Paper February 2025

Growing Expectations for Cloud DVR Recording

Since its introduction at the 1999 CES show in Las Vegas, Digital Video Recording (DVR) has technology has evolved significantly. Although DVR is not a new innovation – Ateme for instance, has been offering DVR solutions since 2016 - the transition to OTT streaming has brought about a fundamental transformation. Viewers now record more content than ever and expect their recordings to be accessible on any device, from anywhere.

When the era of digital television began, hard drives were integrated into set-top-boxes (STBs) to enable recording. Many of those legacy STBs remain in-use today. However, Service Providers quickly realized that adding DVR capabilities directly within STBs was both highly inefficient and expensive. Providing millions of users with extensive video storage capacity meant managing fleets of STBs equipped with expensive, failure-prone, and rapidly aging hard drives.

To address these challenges, Network attached DVR (nDVR) technology emerged. Instead of storing recordings locally within STBs, in customer living rooms, operators leveraged improved connectivity of STB to remove the onboard storage and consolidate all recordings in their datacenters on a single unified platform. This improves efficiency and reduces overhead as it enabled the emergence of the "shared copy" model where identical video segments are only stored once and that recording is served to all the end users that saved it. This already enabled significant improvement to the total cost of DVR as it drastically reduces services cost with less STB failure and reduced overhead on storage.

With nDVR, operators achieved operational improvements while end-users enjoyed a superior viewing experience. Recordings became accessible from any connected device, enhancing convenience and usability. Meanwhile, user behavior evolved: viewers started recording more content, driven by the increasing availability of high-resolution formats like HD and 4K.

DVR Unique Storage Requirements

Even with a shared copy model, the ever-increasing volume of recordings remains a significant challenge for operators, with larger platforms still reaching up to hundreds of Petabytes. Managing data storage of such magnitude comes with its unique set of challenges.

Scaling infrastructure remains the biggest challenge as DVR requires high performance storage and reliable network infrastructure to ensure a good user experience.

Viewers expect near instant starting time, smooth playback of high-resolution content without re-buffering. This means that operators must build large arrays of high-performance drives to support the handle coming in and out of the storage. Most Network Attached Storage (NAS) with high-capacity hard drives will not have the read or write capacity required to deliver content to tens of thousands of viewers requesting high-resolution content at the same time.

The second challenge for DVR, after scalability, is availability (and therefore redundancy). Viewers now expect their content to be available instantly, even months or even years after the recording. To reach this level of availability, low-cost platforms use large Redundant Arrays of Independent Disks (RAID), while more premium platforms use distributed storage systems with dedicated erasure coding mechanisms for better performance and density. The challenges to meet these requirements for availability and performance compound exponentially when scaling to Petabyte scale, as overhead and complexity adds up.

The third major challenge for large DVR platforms is the exponential growth of maintenance cost with scale. Large recording platforms require hundreds to thousands of MWh per year just in power and cooling to keep hundreds of servers running 24/7/365. This, in turn, requires dedicated engineers to maintain the hardware platform, taking care of frequent disk failures and hardware issues across those servers. Engineers must regularly roll out operating system and software updates to correct bugs and keep up with cybersecurity threats to avoid ransomware attacks (remember the LOG4J Zero-Day vulnerability).

The Case for Cloud Storage

Migrating long-tail DVR content to the cloud provides an elegant solution to the problems listed above, namely scalability, availability, and operational complexity / cost.

Cloud Service Providers (CSP) offer virtually unlimited high performance storage capacity with built-in redundancy across multiple Availability Zones (AZ). This makes cloud infrastructure inherently better suited for VOD and DVR compared to on-premises solution as CSP handle the undifferentiated heavy lifting of scaling hardware maintenance and security.

	Live & Time shift	DVR/VOD
Usage	Live & up to 8h	Months to Years
Latency	Important	Does not matter
Storage	Terabyte	Petabyte
	Static	Grow over time
	Volatile	Persistent
	Can be tied to a session	Must be portable
Operation	Mission critical: Cannot be interrupted or impacted	Not critical: Can be degraded/suspended for short period
Infrastructure Fit	Better on-premises	Better in the Cloud

Table 1: Live and non-linear video storage characteristics

One of the additional benefits of migrating long term storage to cloud storage is native storage tiering, which allows operators to choose from different storage classes based on data access patterns. Less frequently accessed video files can be moved to lower-cost storage tiers, optimizing costs without compromising accessibility. This is especially relevant for video recordings as the usage patterns change drastically with time, with older recording being much less requested than more recent ones.

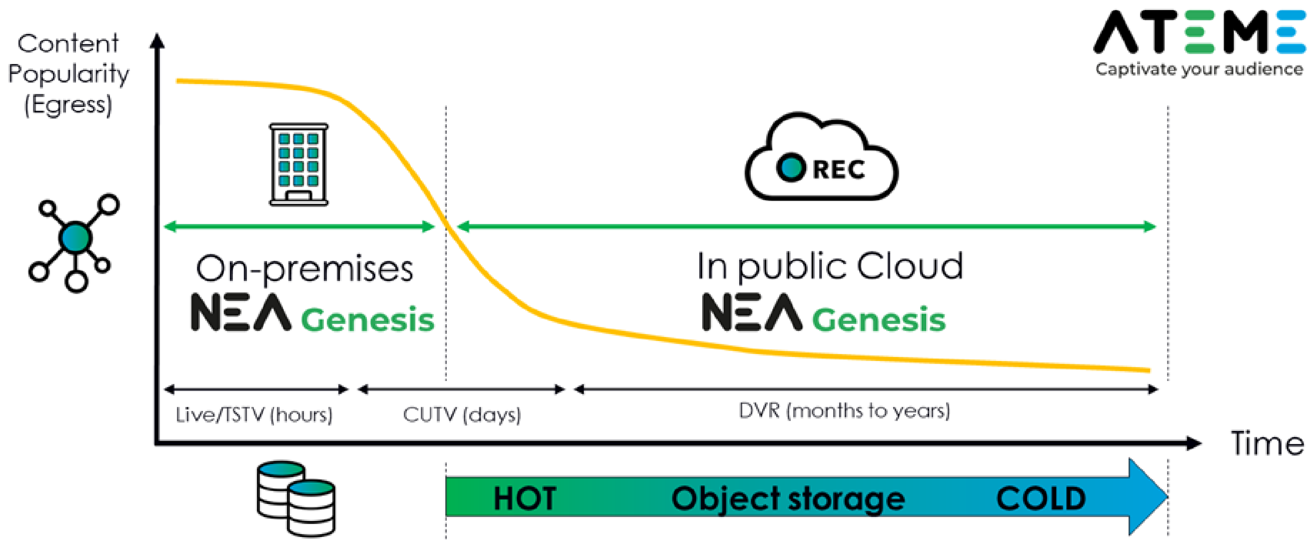

Figure 1: Evolution of video storage performance requirement over time

For streaming OTT platform, a recording will first appear on premises as part of a live channel rolling buffer for time-shifted TV (TSTV) to reduce cloud egress cost as over 90% of the request for this content will come in the first few hours of its lifecycle. When the asset comes out of the TSTV window, it can be sent to the cloud to be stored on a “hot” storage tier for “Catch-up”, or “backwards EPG” services. As a recording ages out of the backward EPG, the number of requests for this asset decreases drastically, and it will automatically be transferred to colder and colder storage tiers. The coldest tier available on most cloud platforms (Archive types) are not suitable for video as they don’t allow for instant retrieval and would result in a poor user experience.

Leveraging a Micro-services Architecture for a Cloud-based Workflow

To best leverage the benefits of cloud infrastructure, the best practice is to split the various components of the workflow into separate services that scale independently based on their unique requirements. In the case of DVR, this means separating content ingest, storage and playout in separate services as they scale independently from each other. Ingest scales linearly with content, with more channels to record requiring more resources. Storage scales with recording patterns, which in turn are influenced by content type and viewer profiles. Sport content and movie channels tend to have higher recording rates than news for example, and some populations are culturally more inclined to use DVR functionalities. Playout (i.e. egress) requires the most elasticity as TV consumption is highly cyclical with strong daily and weekly variations.

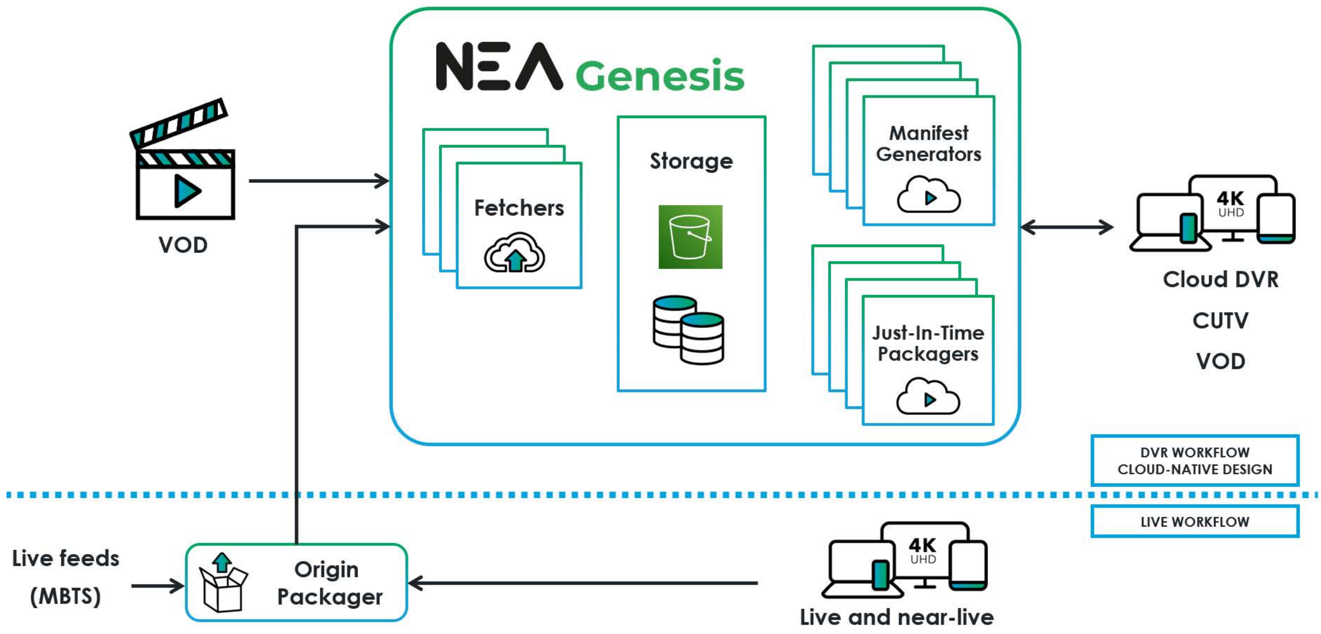

Figure 2: Ateme NEA Genesis high level architecture

Scaling those services independently not only optimizes resources usage but also allows more flexibility to operators. Scalable ingest allows to add large batches of VOD or additional event-based channels. The combination of automated storage tiering and Just-in-Time-Packaging (JITP) enables a "set and forget" mode of operation optimizes resources consumption while handling spiky usage patterns.

This design with separate ingest and JITP also brings more modularity in the design of the overall platform. Content can be decrypted and de-packaged at the ingest to be stored in a secured pivot format (ex: CMAF) for sake of efficiency and future proofness. The content can then be re-packaged and re-encrypted on demand in any format and DRM that is relevant at the time of the request, which can be years after the initial recording.

Storage Tiering

Tiered storage is a terrific way to reduce the total cost of infrequently accessed assets such as video recordings and VOD assets. As explained above, as assets age, they typically generate less traffic and can be shifted to more cost effective, colder tiers of storage. As an example, the tables below show an example of potential savings when applying a tiering strategy over 3 PB of storage.

Cloud storage pricing example for 3PB*	Price per GB/Month	GigaByte	Cost
First 50 TB / Month	$ 0.0230	51,200	$ 1,177.60
Next 450 TB / Month	$ 0.0220	460,800	$ 10,137.60
Over 500 TB / Month	$ 0.0210	2,633,728	$ 55,308.29
	Total	3,145,728 GB	**$ 66,623.49**
Cloud storage pricing example for 3PB using storage tiering*	**Price per GB/Month**	**GB / month**	**Cost**
“HOT“ Tier, First 50 TB (USD/Month)	$ 0.0230	51,200	$ 1,177.60
“HOT“ Tier, Next 450 TB (USD/Month)	$ 0.0220	460,800	$ 10,137.60
“HOT“ Tier, Over 500 TB (USD/Month)	$ 0.0210	117,146	$ 2,460.07
“WARM” Tier, All Storage (USD/Month)	$ 0.0125	209,715	$ 2,621.44
“COLD” Tier, All Storage (USD/Month)	$ 0.0040	2,306,867	$ 9,227.47
	Total	3,145,728 GB	**$ 25,624.17**
		Potential saving (USD/Month)	**$ 40,999.32**
		Potential saving (%)	**61%**

Table 2: Storage price for 3 PB using multiple cloud storage tiers

Note that while “WARM” and “COLD” storage tiers have much lower monthly storage cost per GB, they have costs for requests and retrieval rates that architects need to consider.

Storage Class	**Data retrievals (per GB)**
“HOT”	n/a
“WARM”	$0.01
“COLD”	$0.03
Automated tiering	n/a

Table 3: Requests & data retrievals pricing
** Example based on Major CSP public pricing in the US as of August 9 2024.*

To achieve maximum benefit from reduced costs when using a class tiering strategy, it is important to understand the data access patterns of your service. For situations where access patterns are unpredictable, such as when building a new service without historical data services such as Amazon “S3 Intelligent-Tiering” or Google “Autoclass” can be a good choice. These services can help by automatically moving data to the most cost-effective access tier based on access frequency, without performance impact, retrieval fees, or operational overhead. For a small monthly automation charge, these services monitor access patterns and automatically move objects that have not been accessed to lower-cost storage tiers. Leveraging life-cycle policies allows for a more granular tiering which can be more efficient, but requires more fine-tuning and is more suited for operators with a good understanding of their content’s access patterns. In this white paper we will focus on a green field example using CSPs automation services, but the same concept applies to manual tiering using life-cycle policies.

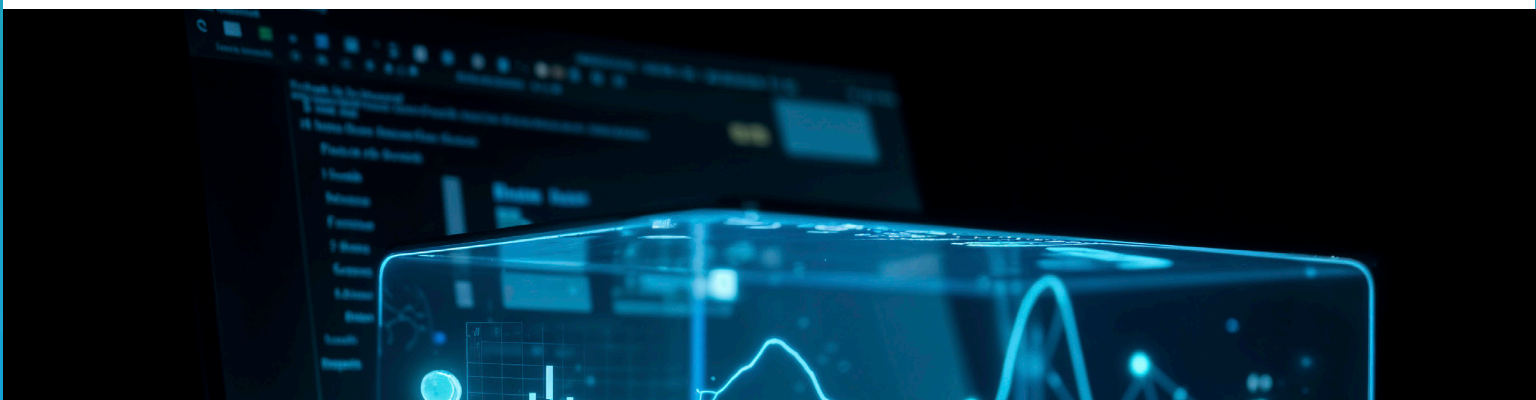

Amazon S3 Intelligent-Tiering and Google Autoclass will automatically move objects that have not been accessed for 30 consecutive days to the warm tier. After another 60 days without access, objects are moved to the cold tier. For both of those services there are no retrieval charges from any tier, and if an object in the *warm* or *cold* tier is accessed later, it is automatically moved back to the *hot* tier.

Automated Storage Tiering pricing example	Price per GB*	Monitoring and Automation (per 1,000 objects)*
"HOT" Tier, First 50 TB (USD/Month)	$0.0230	$0.0025
"HOT" Tier, Next 450 TB (USD/Month)	$0.0220	$0.0025
"HOT" Tier, Over 500 TB (USD/Month)	$0.0210	$0.0025
"WARM" Tier, All Storage (USD/Month)	$0.0125	$0.0025
"COLD" Tier, All Storage (USD/Month)	$0.0040	$0.0025

Table 4: Automated tiering pricing

** Example based on Major CSP public pricing in the US as of August 9 2024.*

NEA Genesis Chunk Containerization

An Adaptive Bit Rate (ABR) live streaming encoder generates multiple video renditions with different resolutions and bitrates needed to enable video players to gracefully transition between different video qualities when network quality is degraded to avoid buffering. An ABR stack can also contain multiple tracks, with different audio channels and subtitle tracks going along the various video renditions. Each track is time aligned and split in chunks, typically, with a duration of two seconds, six, or ten seconds. This process generates multiple files per chunk, each treated as an individual object by cloud storage services which leads to higher ingestion and automation costs NEA Genesis aggregates all those related objects into a larger container, resulting in a significant reduction in the number of objects managed by the storage services.

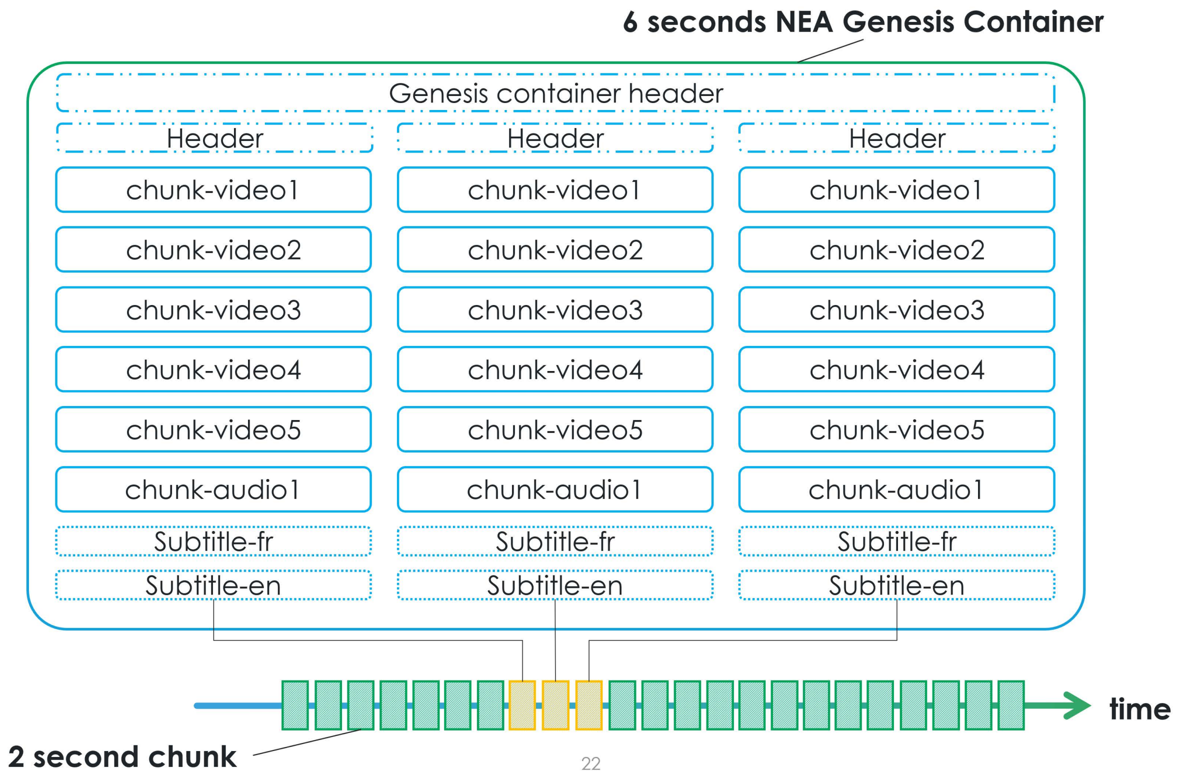

Figure 3: NEA Genesis chunk containerization mechanism

To see the impact of this containerization, let's take the example of a typical HD channel with two second chunks, five video profiles, two audio tracks and two subtitle tracks. Assuming a shared copy model with a 40% recording rate, the impact of chunk containerization is as follows:

Object Monitoring & Automation Costs	Number of objects /channel /month	Automation Cost /channel /month *
Without Chunk containerization	388800	$29.16
With Chunk containerization (2 minute container)	8,670	$0.022

Table 5: Monitoring and automation cost reduction with chunk containerization

** Example based on Major CSP public pricing in the US as of August 9 2024.*

When content playback is requested, NEA Genesis uses Byte-Range Fetches to cloud strorage to retrieve only the needed elements within the Genesis container. This helps retain the benefits of the container while also avoiding unnecessary egress costs from pulling entire container from storage.

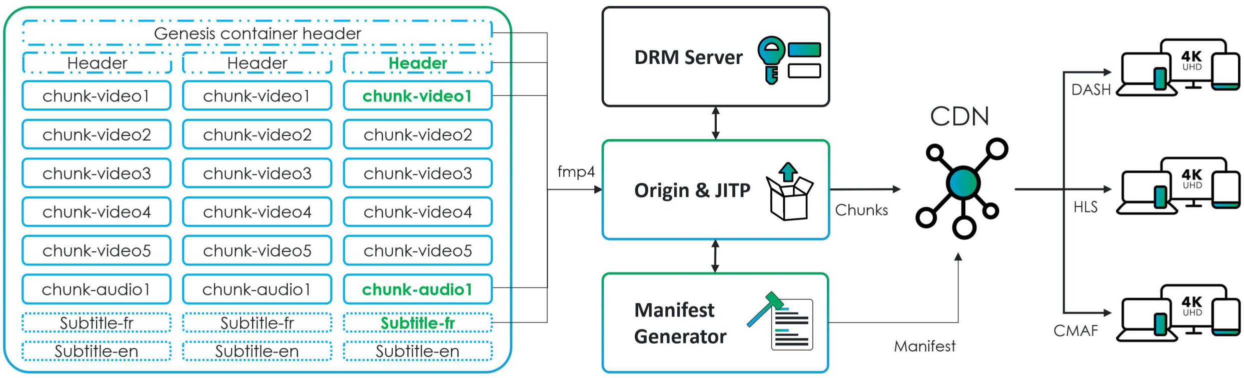

Figure 4: Byte-Range Fetch to Genesis containers

Accurate Trimming and Instant Availability

The last element of cloud storage pricing that we have not covered are the API requests cost to add (PUT request) and retrieve (GET request) data from the storage. Those API call fees, although very small, can add up to a non-negligible cost over time. The intuitive answer to this would be to use the higher possible container size to reduce both the number of API calls to write to storage and the quantity of object for storage tiering automation. In reality, this is a tradeoff where too large containers result in excessive overhead in storage.

By default, NEA Genesis is configured to ingest live content by pulling it from the origin/packager using HLS or DASH as if it were an end-user. It is stored in a channel buffer with two parameters:

- a minimum retention time, it can be seen as a traditional rolling buffer where the entire content is kept for backward EPG services
- a maximum retention time aka "infinite", where when viewers create DVR assets, all media elements for the duration of the assets are tagged as "DVR" and kept until deletion is triggered.

Beyond the minimum retention time, the channel buffer is simply cleaned. To optimize storage, all media elements that have not been tagged as a DVR by at least one subscriber are automatically deleted. With this mechanism an asset does not have to be duplicated for recordings and backward EPG and can be shared with multiple users if a shared-copy model is available.

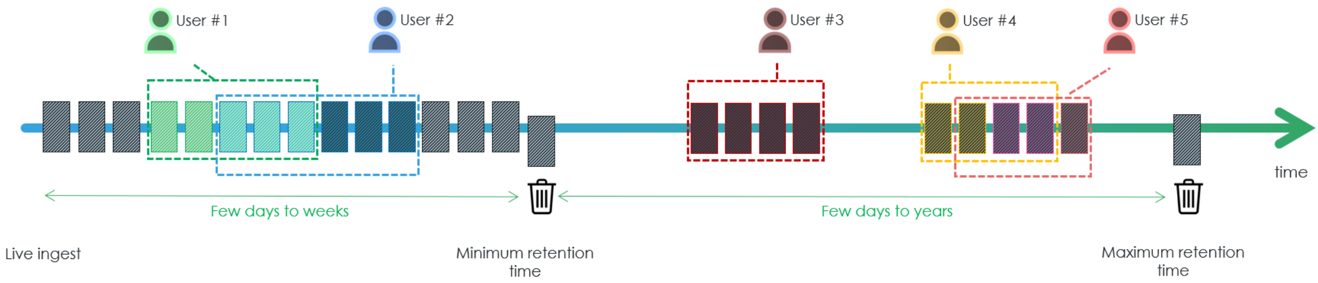

Figure 5: Ateme Infinite buffer technology

This technology has several advantages, for example it enables recordings to be immediately available with no transition period after the end of the program. It also enables significant savings by efficiently deleting all the non-recorded assets after the configurable duration for the backward EPG. Ateme's Infinite buffer enables the use of shorter media segments, ensuring exceptional precision in content trimming. This advanced approach facilitates the virtual movement of content between minimum and maximum retention time buffers.

To illustrate this point let's look at a popular channel with a 40% recording rate and compare the savings produced by Ateme's infinite buffer. In this example, NEA Genesis is configured with automated cleaning policies based in 2 minutes segments versus 5h segments still found on many popular DVR solution in the market.

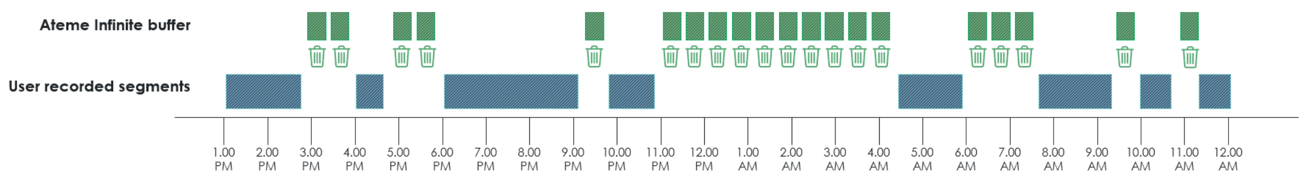

Figure 6: improved recording efficiency with Ateme Infinite buffer accurate trimming

As illustrated in figure 6, the likelihood of being able to remove more than a single media segment with a 5h segment duration is extremely unlikely with a traditional 5h segment duration. With a typical recording distribution, Ateme will remove in average 431 two-minutes segment, which results in much lower overhead, as shown in the table below:

Container length	No container	2 seconds	2 minutes	5 hours
Number of video segment per day	NA	43200	720	4.8
Deleted segment per day at 40% recording rate	NA	25920	**431**	**1**
New managed elements per day*	388800	17280	289	2.8
Retained content (%)	40.0%	40.0%	40.1%	79.2%
Overhead (%)	0.0%	0.0%	0.1%	39.2%
Retained content (GB/Month) *	1540	1540	1564	3048
PUT/ month (one time cost)	**$29.16**	$2.59	$0.04	$0.0006
Storage cost (recurring cost) **	$11.39	$11.39	$11.43	$22.55
Automation/ month (recurring cost)	$29.16	$1.30	$0.022	$0.0003
Total recurring cost /channel /month	**$40.55**	**$12.73**	**$11.46**	**$22.55**

Table 6: Improvement in recording overhead
** Assuming 12Mpbs aggregated bitrate , 5 video profiles, 2 audio & 2 subtitles tracks, 2 second chunks*
*** Assuming 15% "HOT" storage (over 500TB), 10% "WARM", and 75% "COLD", example based on Major CSP public pricing in the US as of August 9 2024.*

By using shorter, 2-minutes segment Ateme's infinite buffer reduces the overhead drastically, while still reducing the number of objects written and managed by the cloud storage. This 2-minutes container size generally provides a good compromise between managing overhead and reducing API and automation cost, but each deployment will have should be optimize for its specific requirements.

Future Optimizations

While NEA Genesis is already a breakthrough for cloud DVR, Ateme is still working on further optimizing the solution to bring more intelligence to the platform.

Ateme is at the forefront of this transformation, bringing video/recording specific optimizations' advanced functionalities to the cloud with .NEA Genesis. The future of DVR is in the cloud, Ateme is leading the charge with the world's first large-scale deployment, recently announced at Swisscom

If you want to learn more about how Ateme can help you optimize your cloud DVR or VOD service, contact representatives from Ateme today.

Francois Guilleautot, Director of Cloud Solutions, Ateme