HD LIVE TO VOD AND PLAYOUT AUTOMATED WORKFLOW


ABOUT RTL NETHERLANDS

RTL Netherlands is a trendsetting multimedia company with a leading position in the Dutch consumer and advertising market. RTL Netherlands has seven television channels and a radio station and are the market leader in the Dutch media world. The company is characterised by an informal culture, professionalism, creativity and passion for serving the needs of its customers: advertisers, viewers and users of interactive media.

OBJECTIVE

With this initiative, RTL was looking to streamline their content processing, editing and delivery workflow in terms of:

- Preserve the highest video quality from ingest to playback.
- Maintain end-to-end Quality of Service.
- Enhance compliance and audience reach.
- Revolutionize content creation and editing workflows.
- Minimize resource use and leverage Just-in-Time automations.
- Streamline content delivery and load balancing operations.
- Integrate Live and VoD workflows.
- Maintain full visibility and control of the entire operation.

The following explains how RTL team was able to achieve these objectives with the help of their partners: Media Excel (Multi-screen transcoding experts), Garland (Media Excel technical partner in Europe), Solarflare (application-intelligent networking I/O software and hardware) and Unified Streaming Platform (streaming and wrapping from one unified source)

OVERVIEW

In short, RTL initiative involved the following aspects:

- HD Live encoding (including subtitle burn-in) and creation of multi-screen/bitrate media assets in a fully 1:1 redundant architecture.
- Dynamic and frame-accurate content clipping/stitching for fast turn-around, without re-transcoding.
- Live multi-screen/bitrate transcoding/streaming and Just-in-Time Live & VoD content repackaging and encryption.
- Streamlined CDN and load balancing infrastructure to sustain projected demand.
- A unified monitoring and control architecture for both Live and VoD workflows.
- Extensive Quality Control and QoS reporting throughout the fully redundant architecture.

1. HD LIVE CAPTURE

(WITH MEDIA EXCEL HERO ENCODER):

RTL Netherlands maintains content acquisition, processing and broadcast facilities across multiple locations (and even across countries). Transporting content fast, efficiently and reliably among those locations can be challenging when the SLAs and overall service availability constrains of an organisation such as RTL are factored in.

Media Excel HERO encoders ingest HD Live content over SDI and produce multi-screen/rate Smooth Streaming assets in a fully redundant and timecode synced manner. Assets are fully interchangeable between each workflow path (A or B) and any disruptions of service one of paths can be fully contained (see Section 3 for more details).

For each input channels, a total of 8 bitrate/resolution variants are produced by HERO and are linked in the Smooth Streaming archive.

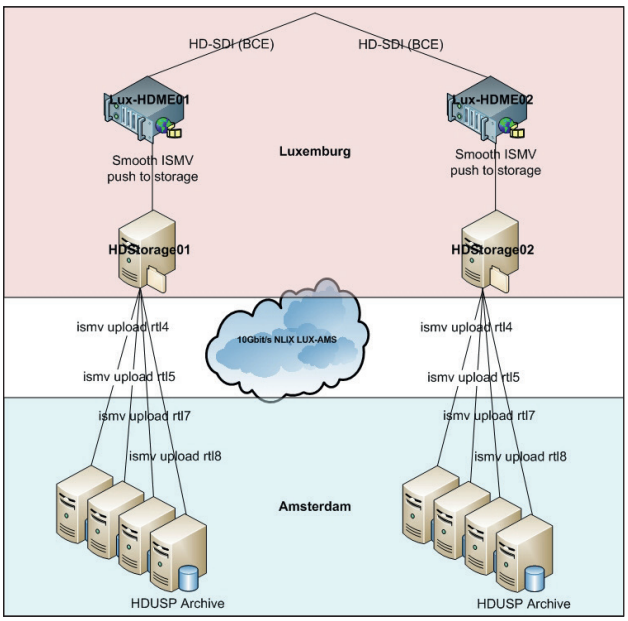

2. SUBTITLE BURN-IN, OPEN CAPTIONING

RTL typically broadcasts content from abroad in English, with the requirement that the Live-to-VoD workflow have the subtitles 'burned' into the streams across all bit rates, instead of having the subtitles done by the playout system (or the player). This means that the SDI source is a clean feed without the subtitles present in the video. Subtitles are inserted in the SDI as OP47-compliant payload by TCube systems.

Media Excel HERO encoders are integrated with RTL's existing broadcast workflow of OP47 subtitles that are embedded into the SDI feed. The generated multi-screen/ bitrate assets contain the subtitles as part of the video frame, this helps reduce content rights costs and also alleviates compatibility concerns with the older playback CE devices.

The solution has given RTL a very clear subtitling on lower bit rates and the ability to fully customise the settings for the burn–in (i.e., position, font metrics, etc.).

3. UNIFIED STREAMING PLATFORM

The delivery from Luxembourg to Amsterdam is done with curl: 1-minute chunks created by the HERO encoder are POSTed with curl to the USP ingest (the Unified Streaming webserver). For this to work, the USP ingest has been updated to allow for chunk posting (instead of the more common workflow of having the encoder POST to the ingest).

Separation of encoder and ingest by means of storing chunks locally and then 'uploading' them to the USP ingest is an extra redundancy step. Content will always be stored locally, so the local archive could be reloaded to the ingest in case of network interrupts. Also, in case of an interruption on Path A, Path B can feed Path A further down the line and vice versa. This is possible because both encoders create exactly the same chunk based on the frame-accurate time code. In other words, chunks are interchangeable between independent paths.

4. STITCHING PROCESS: UNIFIED CAPTURE

The stitching process allows for RTL to create VoD assets post-encode, ensuring fast turnaround and storage efficiency. The stitching workflow allows for VoD assets to be created in a frame-accurate manner by one of two USP servers — should one network path go down or become corrupt. Prior to the stitching process, RTL perform a number of validation checks upon the ISML segments to ensure integrity at the a/v and wrapper layers.

The Media Excel HERO encoders take the embedded SDI time code on mirrored platforms to ensure that all segments across those encoders are absolutely identical, in terms of both content and segment indexing. This allows RTL to safeguard the integrity of the archive. The video segments hosted on the USP server are then available for Just-in-Time stitching/clipping through a simple API call.

As soon as a programme is finished, RTL is able to create a customised VoD package with snippets removed (adverts for example), starting on any frame and ending on any frame. RTL are able to get this customised session on the website in less than 6 minutes from the end of the programme, where most of the delay is introduced via the inline validation checks.

```
<?xml version="1.0" encoding="utf-8"?>
<smil xmlns="http://www.w3.org/2001/SMIL20/Language">
  <head>
  </head>
  <body>
    <seq>
      <video
        src="http://usp-test/video.out/rtl8/rtl8.ism/Manifest"
        clipBegin="wallclock(2014-01-30T15:02:45.960Z)"
          clipEnd="wallclock(2014-01-30T15:02:50.960Z)"
      />
      <video
        src="http://usp-test/video.out/rtl8/rtl8.ism/Manifest"
        clipBegin="wallclock(2014-01-30T15:17:52.680Z)"
          clipEnd="wallclock(2014-01-30T15:17:57.680Z)"
      />
    </seq>
  </body>
</smil>
```

USP and Media Excel have collaborated to ensure that a seamless handoff between the two systems is in place. Unified capture has been extended in two ways:

- The capability to create frame-accurate clips has been added. This means re-encoding first and last chunk to start on the exact frame as specified in the frame accurate EPG RTL provides.
- The ability to seamlessly combine different clips into one has been added. Clips are defined by start and end times in ISO8601 format and presented in the form of a SMIL file to unified capture. A new clip is then created.

5. HPC AND TRANSFORM MANAGER

RTL has created a customer HPC farm (Microsoft Windows-based compute farm) that gets its task from Microsoft Transform Manager. Depending on the target group, it allocates and delegates jobs to Compute Nodes in 4 simple steps:

1. It launches the unified capture application.
2. It downloads all smooth streaming assets from an USP archive.
3. It inserts a key frame at the beginning of the archive (where one might not be).
4. It posts the result to RTLs central storage and publishes the video on the website.

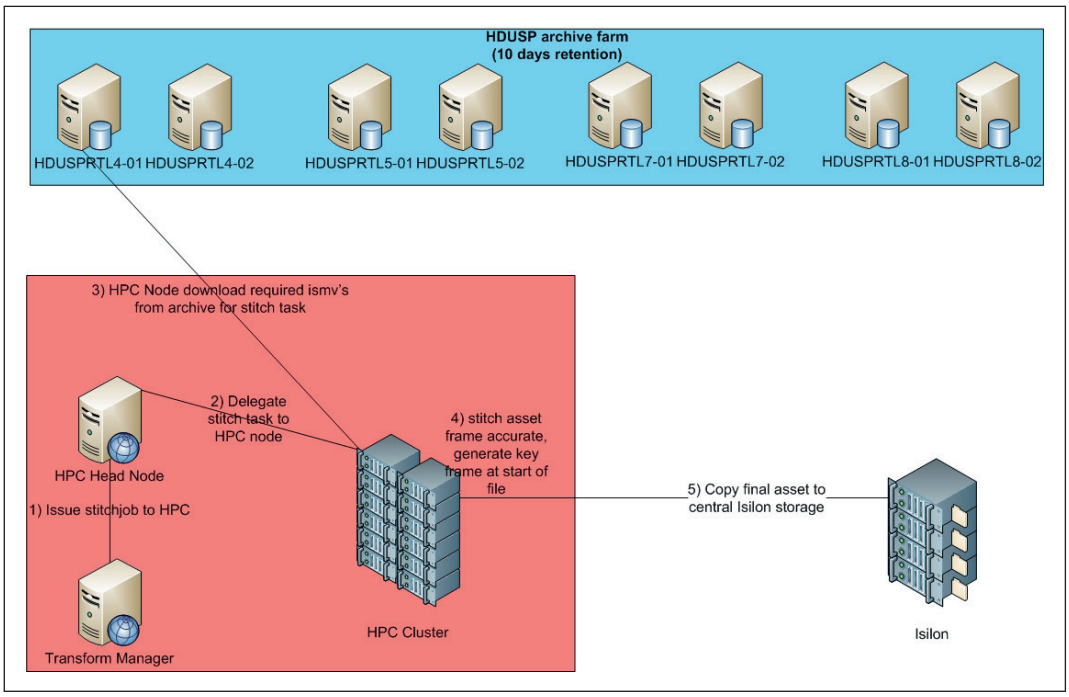

6. LIVE STREAMING

Due to the great success with Media Excel for the VoD platform, RTL Nederland decided to use the same encoders for their live streaming platform as well to replace the older Inlet encoders. The density of the Media Excel encoders (up to 8 x HD SDI per 1RU server) makes it possible to replace the five Inlet encoders with just two Media Excel units.

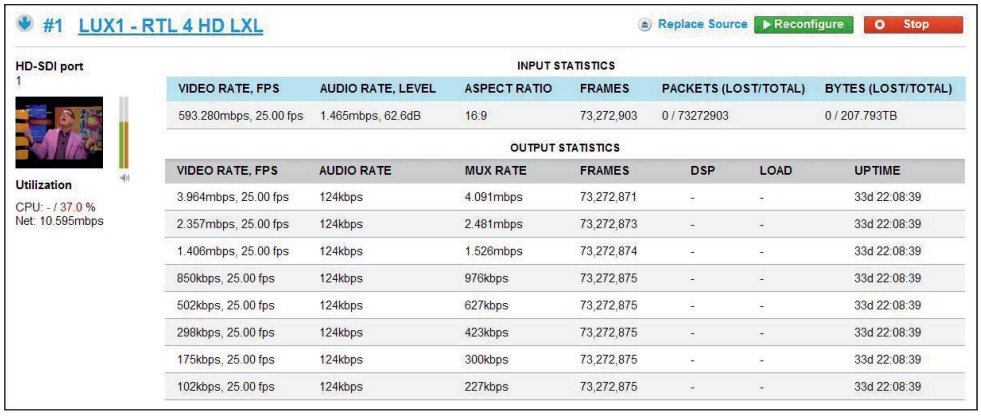

As an extra benefit, those encoders can also do file-based encoding if and when needed. An additional advantage is the centralised HMS management system where all encoders (Live and VoD) can be monitored and controlled.

7. JUST-IN-TIME REPACKAGING & DRM

RTL has been using Unified Streaming Platform for their on–the-fly repackaging and DRM of videos for several years. This gives RTL the ability to store the VoD assets in one format and apply the DRM and repackage output format based on what the device can handle. The CDN servers cache the result, and share a common shared cache. Due to this shared cache, a specific response from the backend can be shared between CDN servers in case it is not available locally on the CDN server. This offloads the backend thereby greatly minimizing the need for a big and costly backend to cope with high demand. For HD, RTL migrated from USP v.1.4.5 to the latest version 1.6.6. The same upgrade will be done for the Live streaming backend as the latest version introduces additional enhancements for HLS sequencing.

8. MANAGEMENT AND MONITORING

Encoder Management

The HMS platform allows RTL to manage all of their encoding platforms across a geographically diverse network. HMS monitors and controls all the HERO transcoders via a single web interface with rich error reporting to the RTL SNMP monitoring system, which allows status updates anywhere in the RTL broadcast chain.

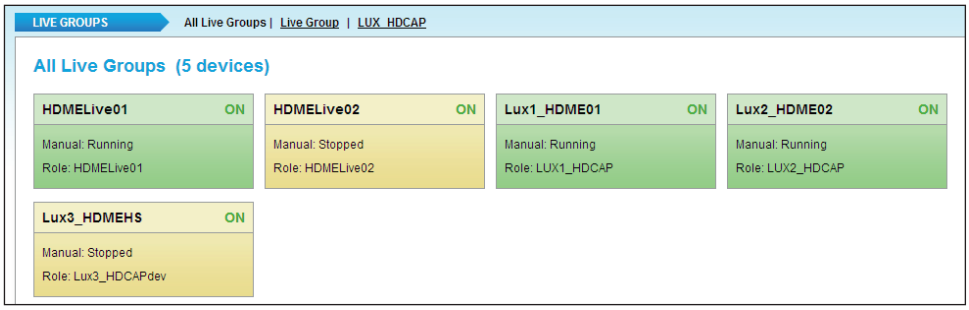

Quality Control

Various quality control points are introduced throughout the flow and are designed so that interruptions in either the A or B path are recovered from the other path, where possible. With the in-house developed monitoring and dashboarding platform, RTL is able to see any loss of frames instantly.

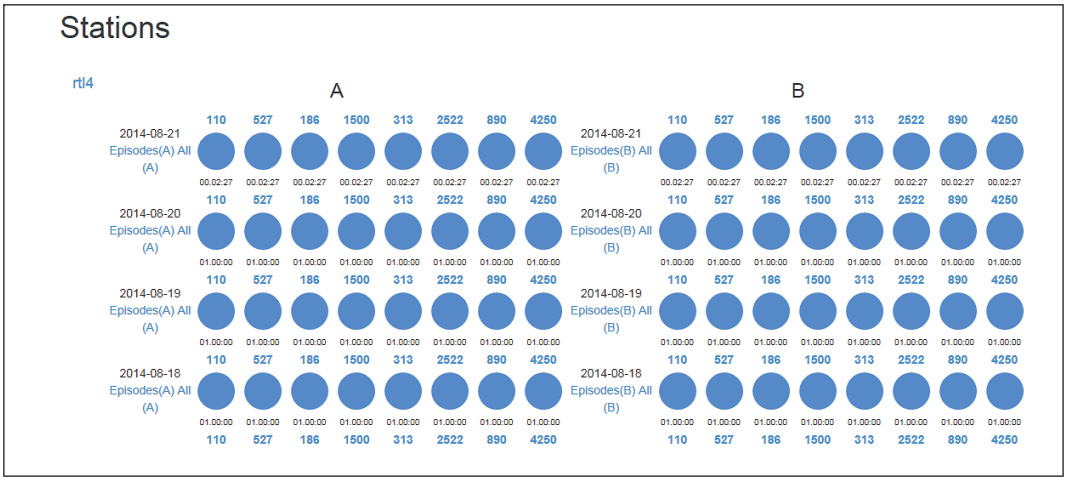

When all frames are accounted for in the archive, the corresponding pie chart on the dashboard is blue. When a single frame is lost, it turns yellow. When multiple frames are lost, it turns red. RTL's automation maintains a slight margin of tolerance for loss of frames before deciding that an asset cannot be created from the archive.

9. PLAYOUT CDN

With the introduction of HD for VoD, the load on the playout capacity has tripled. So far a typical VoD video on XL had a maximum of 1.5Mbit/s. The current highest bit rate for HD is 4.25 Mbit/s. To accommodate the additional demand for network capacity required rethinking the entire CDN from the ground up. Scaling out with the use of the existing setup would become unmanageable. Also turnkey solutions were not available or were too costly on the long term.

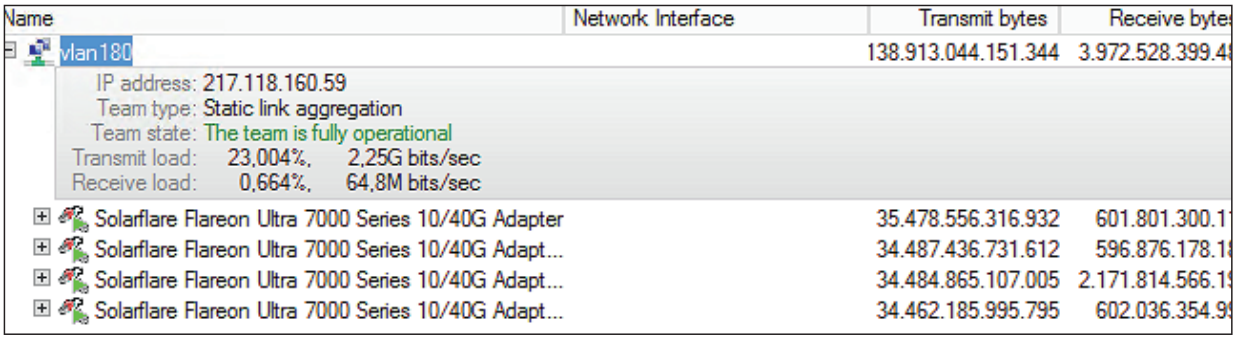

Up until late 2012, the maximum capacity per CDN server was about 4Gbps. At that throughput, the first bottleneck was interrupts not spreading evenly across all available CPU cores. This turned out to be a driver problem of the current NIC vendor. RTL searched for other vendors and started testing Solarflare SFN6122 10GbE Server Adapters. The improvement in throughput was immediately visible, but also revealed a new bottleneck: IO on the local SSD PCI-E cards. The maximum throughput for these cards was around 5,5Gbps with RTL's typical chunk length and size. To overcome this, RTL started tweaking the Application Request Routing (ARR) and Internet Information Server (IIS) settings and found the solution in the way IIS uses Kernel cache. With the correct registry settings, a 128GB kernel cache can be achieved. With the benefits of memory caching, the maximum wire speed of 10Gbit/s per server can easily be achieved. To benefit the most from each CDN server, a 20Gbit/s team was created and CDN servers are field tested to 19Gbit/s in production.

With the additional demand for bandwidth fuelled by the HD rollout, Solarflare joined RTL to test the latest 40Gbit/s (SFN7000 series). RTL needs 40Gbits from one single nic is because each CDN server can only accommodate two PCI-E cards. One is used for the SSD PCI-E storage card and the other is used by the Solarflare Server Adapter. By using Solarflare's latest SFN7142Q 40GbE Server Adapter, each CDN server is connected to the core routers with 4 x 10Gbit/s and should be able to achieve 30Gbit/s. One 10Gbit/s port will be used for failover purposes as it cannot actively be used due to bus/IIS limitations. Thanks to Solarflare, RTL are able to get more out of their existing CDN infrastructure and enhance capacity by using the SFN7142Q cards.

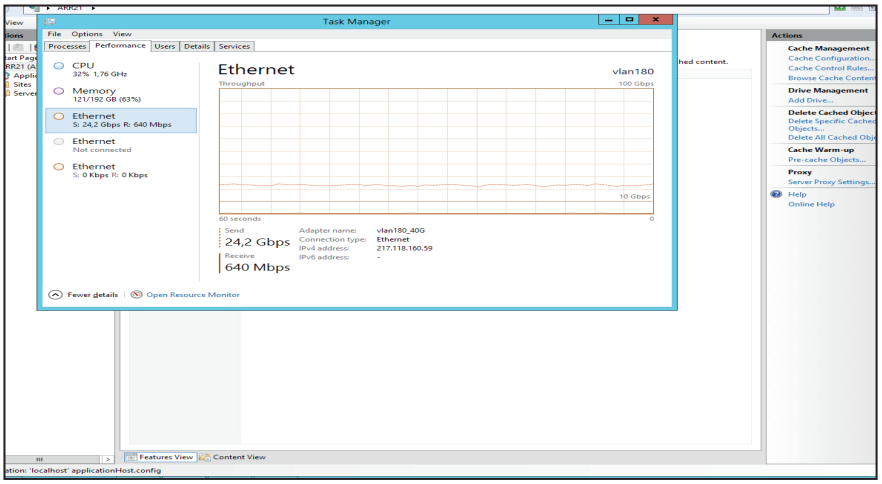

In field testing during the off-season summer months, RTL was able to push a single server up to 24Gbit/s. This was accomplished with the new SFN7142Q and without the use of the extremely expensive PCI-E SSD card. RTL has replaced the PCI-E with four Intel SSDs. By doing so, the cost per CDN server is halved, storage capacity has doubled and throughput per server has increased by 50%.

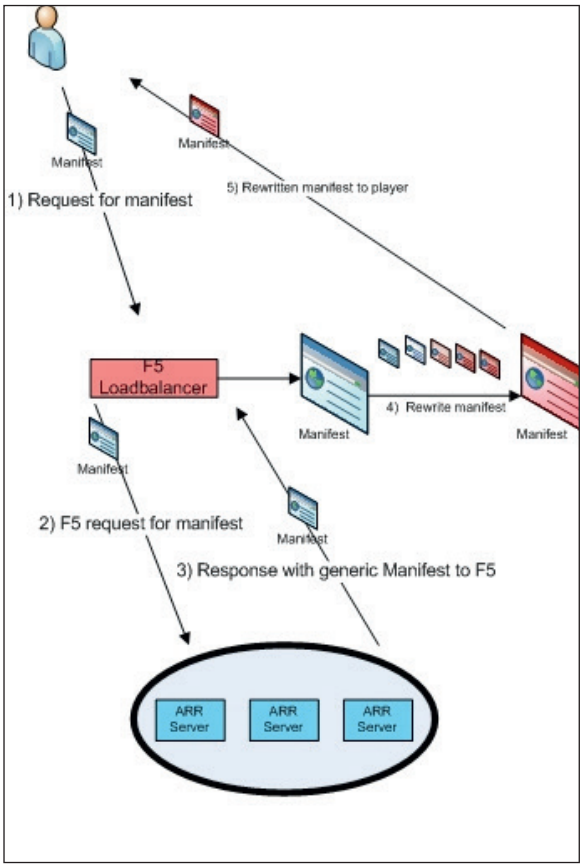

With the throughput per CDN server covered by Solarflare, the IO bottleneck needs to be addressed. To do this RTL use F5 load balancers to cleverly target CDN servers based on the requested content. To minimise IO per CDN server, the content needs to be available on the CDN server chosen by the load balancer. This was done by introducing URL hash-based balancing with the use of custom-build iRule code running on the F5 load balancer. This enhances the cache hit rate to nearly 99% for popular content (less than 30 days old). Further, for long tail content, the URL-based hashing is of great benefit in achieving best cache hit ratio. To prevent a single specific (viral) video from overloading a single CDN server, URL hashing can also target a specific pool instead of one server (this is of great use for the daily popular soap 'Goede tijden slechte tijden').

This solution enables RTL to balance between content duplication on CDN servers and throughput between CDN servers.

CONCLUSION

"Thanks to these new systems and workflow, RTL is able to increase video quality by 300% and at the same time reduce publication time in halve," Says Fardau van Neerden Sr. System Engineer at RTL Nederland. "Expensive post processing and transcoding farms are not needed anymore. Because CDN capacity does not need to scale out with the higher demand of HD, costs can be kept in line with prices on the market, increasing profitability. And with redundancy and auto recovery mechanisms implemented throughout the workflow, downtime and human intervention are kept to a minimal. With better quality and faster publication time, customer satisfaction is increased."