Olympic Streams

USP performance during the Dutch 2012 Summer of Sports.

UNIFIED STREAMING BV

Stadhouderskade 54hs
1072AB Amsterdam
The Netherlands

T +31 6 411 92 024
info@unified-streaming.com

www.unified-streaming.com

Table of Contents

INTRODUCTION

For broadcasters, audio and video content delivery in a multi-platform and consistent way is hard: there are many devices - each with different operating systems, formats and content handling. This divide runs from handheld to desktop to connected TV in living rooms.

Next to a multitude of devices and formats, broadcasters face the concern to support multi-platform consumption without compromising content security or network efficiencies.

Some legacy solutions have difficulty scaling, RTP/RTSP for instance, needs a per stream connection, which impairs any large scale use. Others lack basic web delivery functions.

USP wraps the platform specific solutions available from Apple (HTTP Live Streaming), Adobe (HTTP Dynamic Streaming), Microsoft (Smooth Streaming) and MPEG DASH into one, within existing web server architectures using use generic HTTP servers, caches and proxies.

This whitepaper will discuss the *Unified Streaming Platform* (USP) as used for the 2012 Summer of Sports (Olympics e.a.) by the NOS/NPO in the Netherlands.

About NOS
The Netherlands Broadcasting Foundation (Nederlandse Omroep Stichting, NOS) is the largest news organization in the Netherlands and provides independent and accurate coverage of news, sports and international events - in realtime using multimedia channels as www.nos.nl, radio, (connected) tv, mobile applications, social networks, game consoles and other internet connected devices.

website: www.nos.nl

About NPO
The Netherlands Public Broadcasting (Nederlandse Publieke Omroep, NPO) is known for television channels Nederland 1, 2, 3 and Radio 1, 2, 3FM, 4, 5, 6 and responsible for nationwide delivery of high quality programming in the areas of culture, education, information and entertainment.

website: www.publiekeomroep.nl

About Unified Streaming
Unified Streaming is a leading provider of cross-platform video streaming technologies. Unified Streaming's products are in operation around the world with customers ranging from broadcast networks and online content distributors to small companies and web casters. The ease of use and reliability of Unified Streaming's solutions allow customers to shorten their time-to-market. Our solutions fit into existing frameworks (Apache, Lighttpd, Nginx, IIS5/6/7) thus allowing for greater return on existing investment.

website: www.unified-streaming.com

SUMMER OF SPORTS

The summer of 2012 not only had the Olympics in London, but also the European Football Championship in Poland and Ukraine and the Tour de France – the worlds's biggest annual cycling event.

All of these are prime time with massive audiences, on TV as well as online – and a lot of content is only available online.

In the Netherlands, the NOS/NPO is responsible for delivery of public, nationwide televised and online sport events to Dutch viewers.

As an online event, the Olympics had up to 14 LIVE HTTP streams which should be viewable by 150.000 people simultaneously on all possible play-outs going over 150Gbps.

For this a new platform based on Unified Streaming (USP) was developed in cooperation with the NPO.

Features of the the new USP based online platform are:

- many devices with different capabilities, many play-out formats
- layered caching of HTTP streams (meeting the > 150Gbps target)
- content protection and geoip
- DVR for LIVE streams
- realtime statistics of streams

The NPO has been using Unified Streaming (USP) since 2010 for several new media projects such as the HLS video stream for ‘UitzendingGemist’ (catch up TV) on iPhone/iPad. This service provides the public on demand access to recent and archived public TV broadcasts. Secondly, USP is used for Radio events as the Top2000 and the VisualRadio stream.

The NOS/NPO effort brings the innovative use of existing, well known technologies as HTTP and the latest content delivery formats to the widest Dutch audience possible.

SETUP

This section will discuss used hardware and software.

HARDWARE

The hardware used for the newly built 'HASP' platform (HASP - HTTP Adaptive Streaming Platform) is commodity x86 hardware.

Two enclosures were used, each with 14 blades.

Each blade had:

- 16 cores
- 92 GB RAM
- 10 Gb/s network interface

Each enclosure had 100Gb/s out connected to the high bandwidth uplink at the NPO premise.

One of the blades from each enclosure was used as ingest/origin - this is show in figure 1.

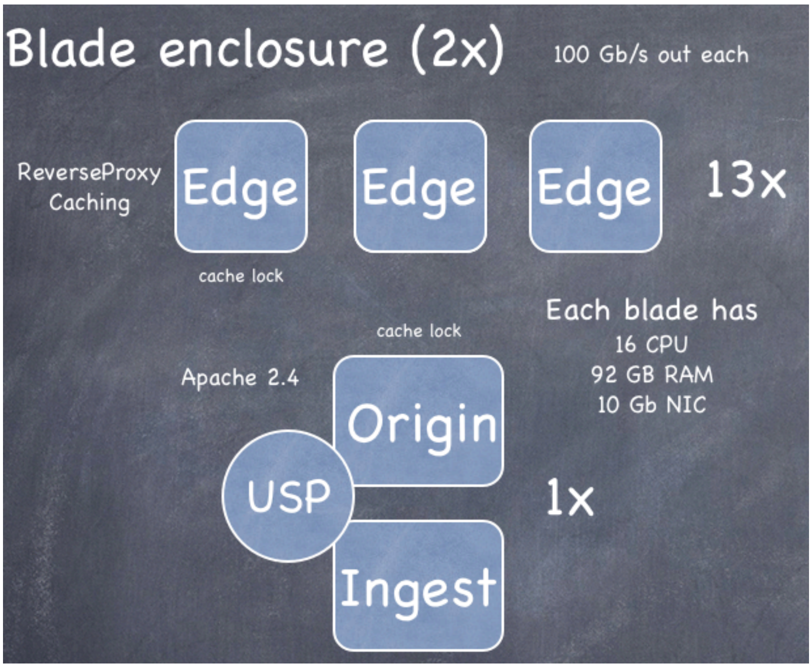

figure 1: Enclosure

SOFTWARE

As software was used:

- Linux OS
- Squid 2.0
- Apache 2.4
    - mod_lua (for scripting)
    - mod_disk_cache
    - mod_proxy
    - mod_smooth_streaming (USP 1.4.45)

Python and Bash were used for some external scripting tasks.

STREAMING
For streaming USP 1.4.45 was used throughout the events.

As encoder was used Digital Rapids Flux type 2200 (4 in total, with dedicated hardware used).

STATISTICS
For NPO's in-house existing statistics and tracking tool ('statscollector') a custom collector was developed.

The NPO statscollector is used to gather information about load on the edges, bandwidth used (per edge/stream and total) and the number of viewers per stream (simultaneous and aggregate).

This information then is used for instance realtime to direct new viewers to the least used edge so an even load per stream can be accomplished.

There were two options to gather these statistics from the edge:

- the server status info from the edge Apache server
- the log generated by the edge Apache server

After initial testing it turned out the best option was to use the generated log files and parse these every n seconds to gather total bandwidth per edge, bandwidth per stream and (unique) viewers per streams.

Full real-time log parsing (for instance by using a fifo or queueing log streams for secondary processing) was not needed in this solution, but for future events and subsequent load on the edges this might be a development to consider.

Each edge uses a python script to parse log files, taking into account log rotation and keeping track of appearing and disappearing viewers. Viewers can be identified by their unique session id which the receive after being checked for their geo-location on the main website before being directed to a particular edge.

Edges keep track of statistics data using a small sqlite database and a second process reads this database every n seconds to create a HTTP post to the existing NPO statscollector (which then could be queried by the load balancing algorithm mentioned before). Figure 2 exemplifies this.

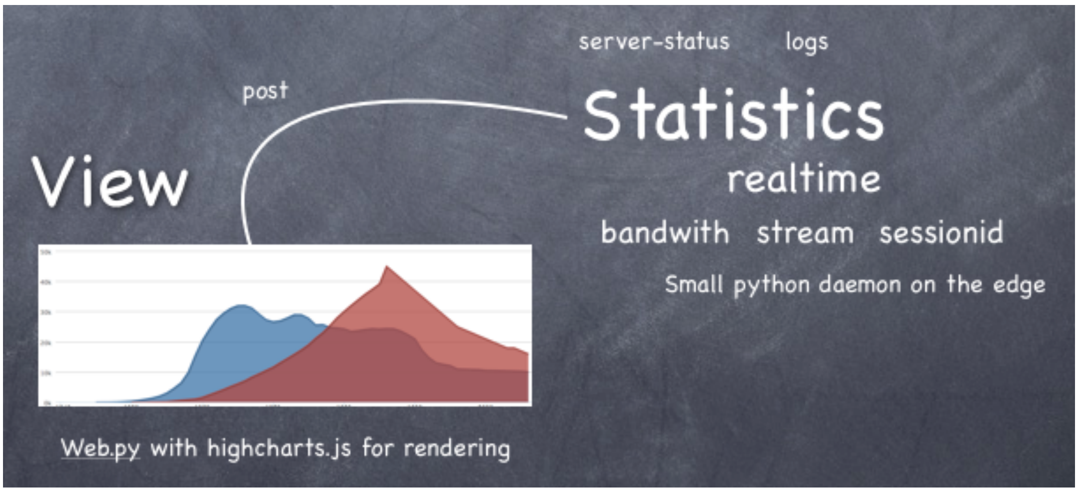

figure 2: statistics collection

THROUGHPUT

BITRATES

The bitrates used for the streaming platform were the following:

VIDEO	PROFILE	BITRATE
H264	Base	350k
		500k
		700k
		1000k
	Main	1500k
		2000k

AUDIO	BITRATE
AAC	128k

The live streams (14 max with the Olympics, 1 with the Tour the France and 2 simultaneous with the European Football Championships) were streamed from the encoders to both ingests and duplicated. This ensures robustness (the ingest will receive 2 inputs for each stream so when one stream should stop, ingest continues from the other).

The aggregate incoming data per stream on each ingest thus was 6050Kb/s (or about 6Mb/s, which is about 0.756 MB/s - under one MB per multibitrate stream).

In total each ingest had about 169.4 Mb/s incoming for the maximum case of 14 live Olympic streams (14 * 6.05 = 84.7 Mb/s but duplicated so times two, which is 169.4 Mb/s)

PEAK

The peak moment recorded was slightly over 100 Gb/s as aggregated output of all edges to 70.000 simultaneous viewers:

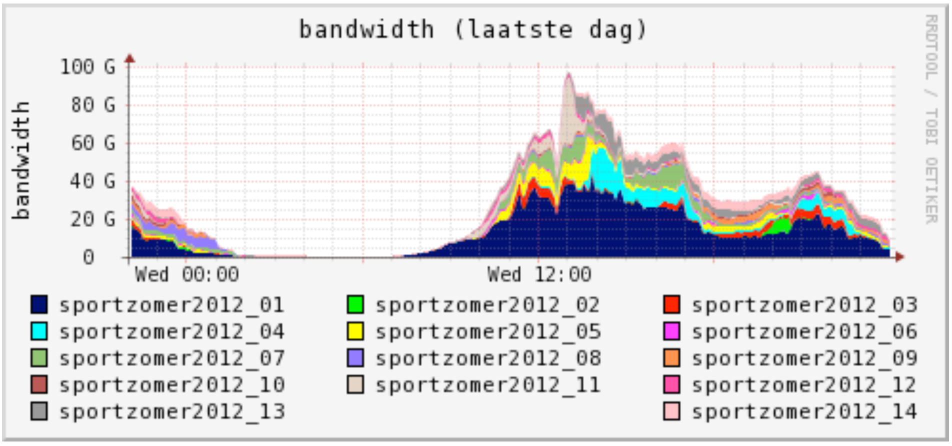

figure 3: peak recorded bandwith

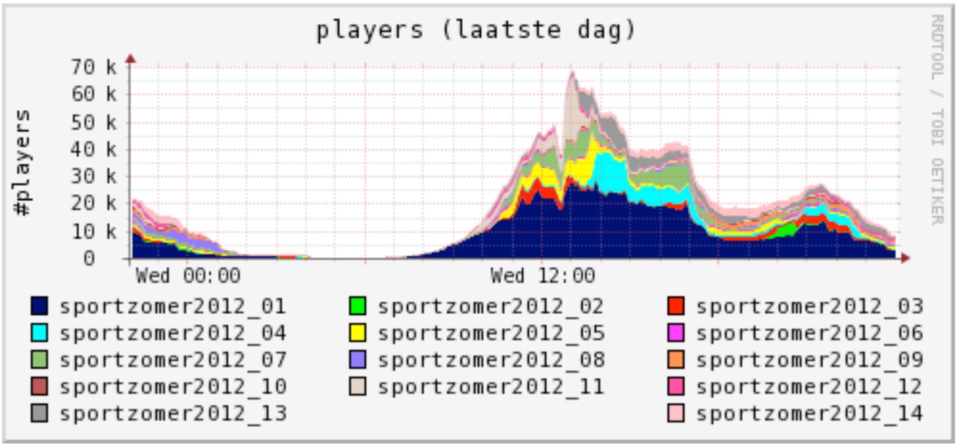

figure 4: peak simultaneous viewers

The different colors represent the different streams watched.

FURTHER INSIGHTS

From the gathered statistics we can learn more, for instance about the bitrates viewed or trends over a longer period.

MOST VIEWED BITRATE

Another measurement indicates that the number of players at around 23:55 on the 6th of August of 2012 was around 16.000, the bandwidth streamed was as pictured in figure 5 - about 26 Gb/s. It is reasonable to assume most of these players were viewing the streams at the highest bitrate (2 Mb/s): if all players would watch at 2 Mb/s the total bandwidth would be 32 Gb/s but it is in fact a little lower: about 6 Gb/s lower.

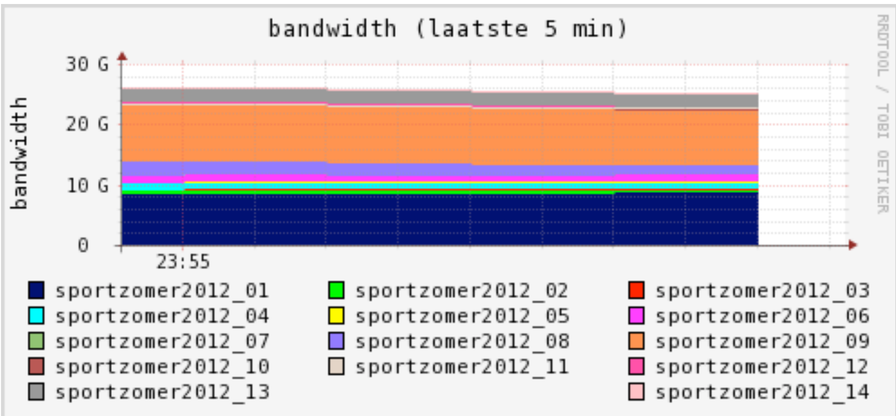

figure 5: bandwith closeup

TIME OF WATCHING

From the day and week overview it can be learned there were two peaks daily: one in the afternoon and one in the early evening, dropping off to almost zero after 01:00 - and that this behaviour was consistent.

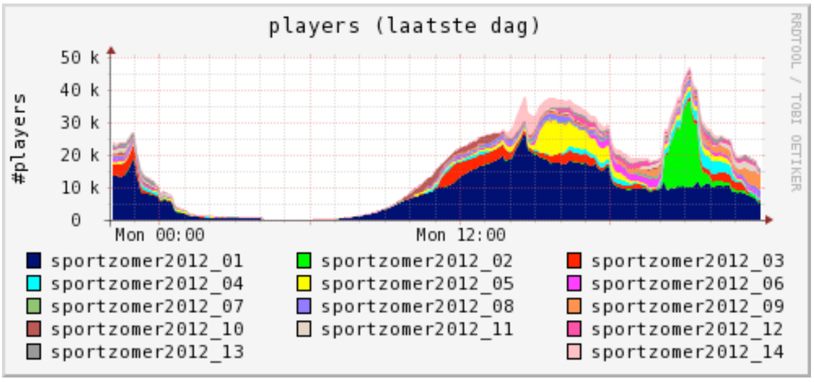

figure 6: viewers per day

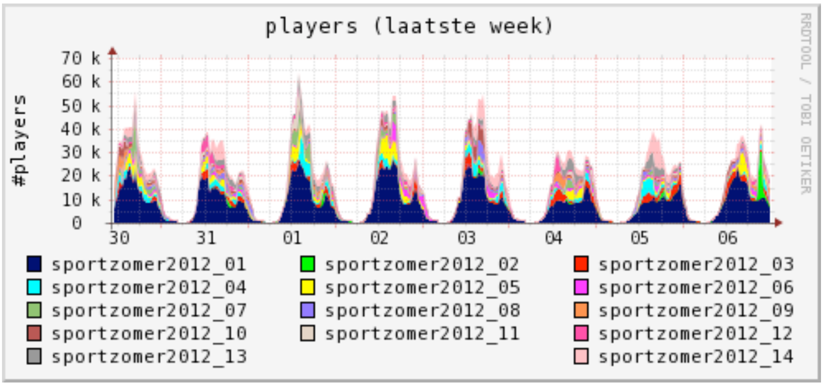

figure 7: viewers per week

STABILITY (4 WEEK VIEW)

From a four week period it can also be learned that the the platform behaved stable and predictable (the single color streams in week 28/29 where of the Tour de France):

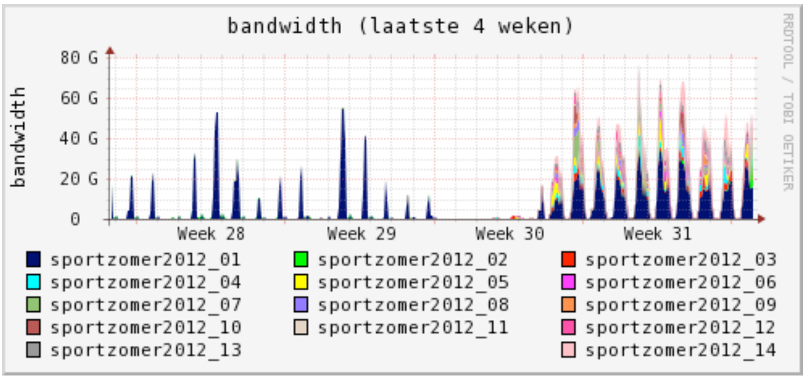

figure 8: bandwith four weeks

CACHING

As can be seen from figure 1, caching was used heavily. In effect, only one blade in each enclosure was used as ingest/origin. The thirteen other blades were used as caching edges.

Caching was setup using Apache 2.4 and mod_disk_cache as well as mod_proxy for the reverse proxy (so when fragments were not in the cache they could be fetched at the appropriate origin).

Relevant proxy part of the configuration file (the vhost config):

```
<Proxy balancer://origin>
    BalancerMember http://origin.example.com
    ProxySet lbmethod=byrequests
</Proxy>

ProxyPass / balancer://origin/
ProxyPassReverse / balancer://origin/
```

And for the caching part:

```
CacheRoot /tmp/http_edge_cache
CacheEnable disk /
CacheDirLevels 5
CacheDirLength 3
CacheHeader on

CacheQuickHandler on
```

LOCKING

Unfortunately, cache locking in Apache is not working very well with either the process or thread worker model. The event worker was deemed to too immature to rely on it, even though its locking implementation seems to be better.

To address this issue Squid 2 was used. Initially only on the origin, but it was found out that this was not sufficient: Squid 2 is single threaded and the Squid process was maxing out the one CPU it could use on the origin.

Therefor, Squid was also put in place on each edge, effectively employing a ‘double locking strategy’. For each cache miss (not finding a HLS, HDS or SS fragment) the edge would do only one fetch at the origin and as the origin also used Squid only one edge effectively fetched from the origin. Other edges requesting the same fragment would be served from Squid on the origin.

This double locking lead to a very stable situation where the origin served each fragment only once. Load on origin was very low: each fragment ingested was originated only once.

CACHE HIT/MISS RATIOS

No statistics were kept of hit/miss ratios on the edges with the actual events - during the testphase were edge performance was deemed to be sufficient for the expected load.

However, given de double locking strategy as mentioned above combined with the data from figure 5 it can be understood that the average load per viewer at that moment was 1.625 Mb/s served so on a total load of 1.13 Gb/s per edge a minimum of 1.625 Mb/s would be fetched from the origin per stream type (HDS, HLS of SS) to a combined load on the origin of 4.875 Mb/s.

The calculated hit ratio therefore is 0.9956% and the miss ratio is 0.0044%.

CONCLUSION

USP performed very well and proved to be a stable and scalable streaming platform.

During the 2 Olympic weeks over 5 million online viewers visited the NOS Olympic site and at peak moments, as for intance the Gold medal of Dutchman Epke Sonderland, the combined traffic was over 100 Gb/s on the HASP platform.

An addional overload was created for subscribers of KPN (a Dutch telco) pushing the total amount of data streamed during the Olympics to over 130 Gb/s

This has made the NOS/NPO conclude to add more servers for the next big event.

Thus asked, dr. ir. Egon Verharen, manager NPO R&D and responsible for the setup of the new streaming platform, said the following “USP has proven to be the best choice for flexible and scalable streaming platform able to deliver quality streams to all relevant client apps and devices that our audience use at the moment.”