


Cache as a Service: Leveraging SDN to Efficiently and Transparently Support Video-on-Demand on the Last Mile

Motivation

- Globally, **Internet video traffic was 57% of all consumer Internet traffic** in 2012 and will be 69% in 2017 ^[1]
- Mobile video traffic** exceeded 50% for the first time in 2012 ^[2]
- High-Definition VoD surpassed Standard-Definition VoD** in 2011. By 2016, 79% of Internet VoD will be HD ^[1]
- Trend to improve video quality even more** as we move to Ultra-HD (4K-8K) and 3DTV that support 4 times higher resolution than HD



Observations


- Video streaming is fast becoming an **essential part of consumers' lives**
- The network has now to transfer an **enormous amount of video traffic** (~45.000 Petabytes per month in 2016); big strain on the network
- We need a solution that :
 - Ensures high-throughput end-to-end** (especially with HD)
 - Minimizes distance between source video content server and user** so that data transfer occurs quickly & reliably to the user -> increased Quality of Experience (QoE)

Challenges with Video-on-Demand

- Currently VoD requests are **handled naively**; there is an independent flow per request
- These flows are **duplicated** minutes, hours or days later (by same or different user)
- We observe identical delivery of media objects through the same network segments
- Consequently, the end-to-end capacity of network infrastructure must grow continuously to match the increasing number of Internet video users
- The increasing popularity of VoD and especially of HD content worsens this

OpenCache Architecture

OpenCache is an **OpenFlow-assisted in-network caching service** that provides efficient, transparent and highly configurable caching and distribution of VoD in the last mile




OpenCache is

Transparent	Extensible (API)
- Cache delivery undetectable to client - No need for new client h/w or s/w - Retains underlying delivery mechanism - Works with existing caches and CDNs	- Exposes an interface for third parties - Enables part or all of the caches to be provisioned at will by cache owner - Simple integration of new instances
Adaptive	Efficient
- Uses real-time metrics from OpenFlow switching hardware - Supplemented with live cache hit and miss metrics - Enables informed decisions by operators	- Reduces the repeated delivery of identical content; satisfy requests locally - Reduces inter-domain traffic and cost - Retains unified point of control - New revenue stream for network operator

Evaluation on OFELIA


- OFELIA** is an OpenFlow pan-European experimentation testbed
- Topology** : Deployed OpenCache on three OFELIA islands distributed geographically
 - Switzerland : ETH Zurich
 - Italy : Create-NET
 - Spain : i2CAT
- Tests** : Over 120 inter-island (federated) VoD experiments using an adaptive video streaming technology (MPEG-DASH)




- Evaluation Criteria** :
 - Startup delay** (QoE metric)
 - External link **network utilization**
 - Video quality (bitrate)** requested (QoE metric)

Key Results

- Reduced startup delay up to 35% -> increased QoE for end-user**
- External **link utilisation reduced to virtually zero** (only background traffic remained)



- Increased video quality**; requested bitrate 8 times higher -> **increased QoE for end-user**



Panagiotis Georgopoulos*, Matthew Broadbent*, Bernhard Plattner*, Nicholas Race*

* Communication Systems Group, ETH Zurich, 8092 Zurich, Switzerland

◇ School of Computing and Communications, Lancaster University, Lancaster, LA1 4WA, UK

[panos@tik.ee.ethz.ch, m.broadbent@lancaster.ac.uk, plattner@tik.ee.ethz.ch, n.race@lancaster.ac.uk]