

MUSLIN

Simon Da Silva

Delivering video content with a high and fairly shared *quality of experience* is a challenging task in view of the drastic video traffic increase forecasts, as live video traffic will grow 15-fold by 2022. Currently, Content Delivery Networks provide numerous servers hosting replicas of the video content, and consuming clients are re-directed to the closest server. Then, the video content is streamed using HTTP Adaptive Streaming solutions. However, servers and network links often become overloaded during major events, and users may experience a poor or unfairly distributed *quality of experience*, unless more servers are provisioned.

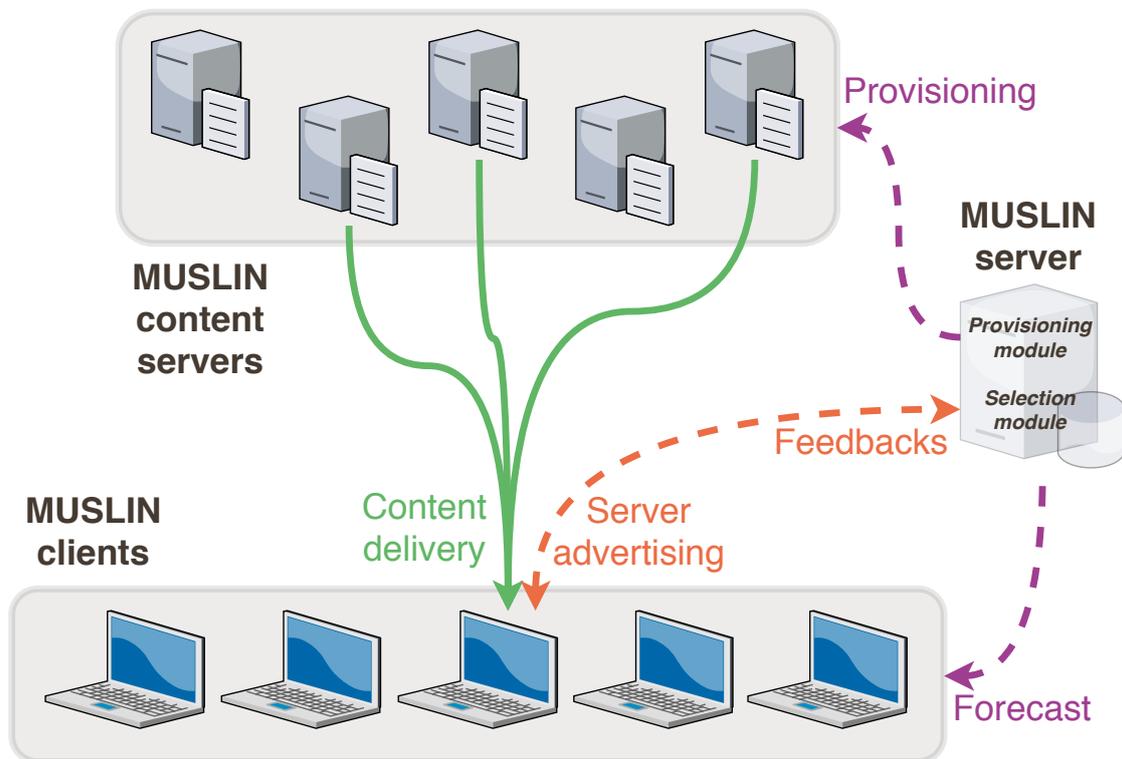
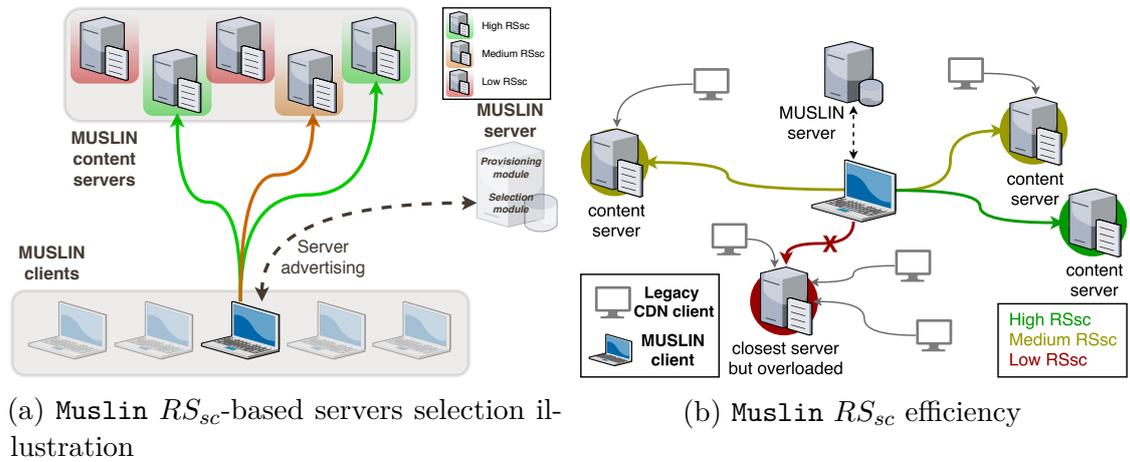


Figure 1: Muslin overview

Muslin is a solution supporting a high, fairly shared end-users *quality of experience* for live streaming, while minimizing the required content delivery platform scale. **Muslin** leverages on MS-STREAM, a content delivery solution which aggregates video content from multiple servers to offer a high *quality of experience* for its users. **Muslin** dynamically provisions servers and replicates content into servers, and advertises servers to clients based on real-time delivery conditions.

As shown on Figure 1, the **Muslin** server provisioning module periodically estimates the required throughput to dynamically adjust the infrastructure scale according to real-world needs.

The **Muslin** server selection module then advertises relevant content servers to clients depending on multiple criteria such as distance, bandwidth and server load. aggregated into a RS_{sc} ranking score (see Figure 2a). As illustrated in Figure 2b, if nearby content servers are already overloaded, the **Muslin** server selects and advertises other content servers with a higher RS_{sc} to the client.



We have used **Muslin** to replay a one-day video-games event, with hundreds of clients and several test beds. We showed in our experiments that thanks to the coupling of MS-STREAM with the proposed **Muslin** system, end-users experienced almost no rebufferings, a higher video bitrate, and more evenly shared *quality of experience*, compared to existing state-of-the-art streaming systems setups, with a smaller infrastructure scale (about -18%).