

Mark Nottingham
mnot@mnot.net <http://www.mnot.net/>
2010-09-23

THE REAL-TIME WEB

Position: the issues raised in this workshop are not unique to real-time interactive voice and video applications. Instead, they should be considered in the context of the overall evolution of the Web.

To wit:

- **NAT Traversal** is an issue not only for video and audio "chat" applications, but also textual ones, games, file sharing and any other communication pattern where you don't want servers to intermediate.
- **Bandwidth Estimation and Rate Control** is likewise useful not only with video and audio that's real-time, but also for more traditional video streaming.
- **Notification** is a problem for a wide variety of applications, from stock quotes to games and so on, as exemplified by WebSockets use cases.

The important point here is that any new capabilities, APIs and protocols should be designed not just with real-time video and audio in mind, but the real-time Web as a whole.

For example, NAT traversal is a well-worn topic. To make it useful on the Web, however, you need to be able to tie the server exposed behind the NAT as a meaningful URL - because URLs are the bedrock of the Web.

Real-time, peer-to-peer applications also elevates application state to not only live on the server in a client/server architecture, but also to be local on the peers. As such, what have previously been "anonymous" resources on the browser need to be identifiable.

A successful outcome would be one where it's possible to build peer-to-peer applications that leverage the existing Web infrastructure seamlessly; where HTTP can be used to stream video and audio in two directions, alongside images, HTML markup and JSON or XML data. Where caching and intermediation can be used -- and extended -- to add value.

Following this path would be a deliberate decision to invest in the current Web architecture, rather than partially duplicating it. While SIP and XMPP were reasonable directions to take given the state of HTTP at the time, it's worth revisiting these decisions in light of the evolutionary path of the Web, where developments like SPDY will make it possible to layer peer-to-peer communication in much more seamlessly.

Therefore, my recommendations are to:

- Define real-time video and audio in terms of HTML5 media streaming over HTTP
- Develop NAT traversal and rendezvous techniques which are "RESTful"
- Explore bandwidth estimation and rate control techniques for HTTP as a medium-term solution
- Accelerate evolution of HTTP with proposals like SPDY and WAKA as a longer-term approach.