

Mozilla Position Paper

Interactive / Real Time Communications on the Web

September 20th, 2010

Christopher Blizzard
Tim Terriberry
Jonas Sicking

Re: RTC-Web Summit (Oct 6th, 2010)

Background

Tim, Jonas and I work at the Mozilla Project. Mozilla is best known as the public benefit organization that develops the Firefox web browser. Firefox is used by over 400 million people around the world and is now the majority browser in several countries. Firefox is open source software, developed by thousands of people from around the world.

Our interests and positions are our own, but they closely reflect the positions of the organization that we happen to work for. In terms of outcomes, Tim, Jonas, myself and Mozilla-as-an-organization share an interest in the same sets of outcomes.

Some background on us as individuals: Tim is the maintainer of the Theora codec and is also one of the leads helping to develop the CELT codec. He's also been contributing heavily to the new VP8 codec that was recently opened up as a result of Google's purchase of On2. He's got a strong background in coding theory and does a large amount of the actual implementation and testing as well.

Jonas is one of the lead developers of Gecko, the engine that powers the Firefox web browser. He's deeply involved in the web-facing APIs that developers use to build the web. He's also got a long history with security issues that affect those APIs, including cross-domain attacks and privacy problems that are created by improperly designed APIs.

I have a very different background, but am also quite technical. At Mozilla my role is to do product management for the web-facing side of our platform. That includes most of the APIs that are used by web developers - HTML, CSS, JavaScript and everything that binds them together.

Tim and I have spent a lot of time working together on the questions around open video and I've been a strong public advocate for HTML5 video.

Outcomes

We hope for a few things through participating in this summit. Our hope is that we can achieve rough consensus on a set of questions, outlined below, with the rest of the participants in the conference. In the absence of consensus, our hope is that we can at least have a process outlined to achieve closure on the open questions.

Questions

1. What are the assumptions for a base set of IP licensing terms for development of technologies to bind to the web? This is important for those of us shipping open source implementations, but also for members who participate in Internet-standard working groups like the IETF or the W3C.

2. How do we want to bind Interactive & Real Time Communications to the web from a technical perspective? In particular, how do we want to make sure the power of the web, expressed through HTML, CSS and JavaScript can also be expressed with RTC as an integral part?
3. Assuming a particular architecture, we're likely to allow web clients to connect to each other in a direct, point-to-point manner. In the security models of the web, this will be new. In general, when communicating with other individuals there's some expectation of pseudonymity. How can we expose that to users in a way that makes sense?
4. In the area of code collaboration, what are the areas that need work, and who will do it?
5. What are the other standards bodies that will be involved, and how do we engage with them?

Positions

1. Our belief is that any base technology for the web, real time communications or not, should be royalty-free. We also believe that this should apply to anything else required to implement a baseline stack as required or implied by the RTC standards. This allows for the kinds of experimentation that we've seen on the web itself and allows implementations to co-exist without a permission model around it to inhibit innovation and interoperability.
2. Any successful solution is likely to be a multi-vendor effort that includes open source implementations as part of the ecosystem. It would be good for all parties if there were at least one production-quality reference open source implementation for servers and clients so that people would have a transparent, inspectable implementation to test against.
3. As with #1 and #2 above, Royalty-Free is seen as important for interoperability. That is, anyone can implement or re-use an existing implementation without worrying about getting permission from others.
4. Any solution in this space needs to have strong commercial support. Not only so that people can build products and services around it, but also to maintain the kind of investment required to make the stack one that's generally useful to the entire Internet - that is, users and commercial and non-commercial entities.
5. History has shown us that a successful API design model for the web is to start simple and expand out from there. Large and complex APIs built on large standards are often not interoperable and are also rarely fully implemented. We prefer a process of specify, implement, learn and grow. This has worked well for the rest of the web and has resulted in a relatively lean platform. We suggest that we stick to that with RTC as well.

Conclusion

We want to ship something that will allow people to communicate on a real time basis, with audio and video, and be able to combine it with the technology of the web. We know that will require the participation from a large number of organizations and individuals and our hope is that the set of people involved in this group will be part of that.

Chris Blizzard
September 20th, 2010