Harlan Yu
Principal, Robinson & Yu LLC

W3C Workshop: Do Not Track and Beyond
November 26-27, 2012

I participated early on in the W3C's tracking protection process. I served on the Program Committee for the Workshop on Web Tracking and User Privacy in April 2011. Since then, I've watched the ongoing Tracking Protection Working Group process with great interest, and growing concern.

The Working Group has made many significant strides, and has forged meaningful consensus, on how the technical mechanism for Do Not Track should work. The mechanism appropriately describes not only the HTTP header field, but also a consistent JavaScript representation of the user's preference, and a way for websites to establish exceptions for tracking from the user. These achievements open the possibility for more transparency in corporate data collection practices, and for consumers to better understand how their information is used and to express preferences about such use. Consensus on the technical protocol is critical to paving the path for policies that can improve consumer privacy.

One guiding goal for the Working Group has been to devise a single, global substantive meaning for the tracking preference, that all sites on the Internet would abide by. If that were feasible, it would carry many obvious advantages. However, after a sustained and robust effort, the group has been unable to define tracking, unlinkable data, first parties, or third parties, and has been unable to agree on how servers might address "rogue" user agents, among a host of other hotly disputed substantive questions—including the scope and goals of compliance.

Given the participants' positions, if the Working Group does publish a Compliance and Scope document, either or both of two outcomes seem likely: (1) key stakeholders in industry will vehemently object to its contents, making clear that they aren't part of a consensus on what DNT should mean, and will decline to implement the standard, and/or (2) privacy advocates will press policymakers to institute regulations that are more stringent and exacting than the W3C document specifies.

In either case, it seems unlikely that the resulting Compliance and Scope document will be able to stand on its own to protect user privacy, absent formal regulatory action. While the Working Group's charter specifies that one success criterion is the "[a]doption of deliverables by user agents and compliance by industry," a standard permissive enough to achieve widespread industry adoption will likely, by the same token, be too permissive to address the underlying concerns that triggered this process in the first place.

What this means in practice is that the non-consensus substance of the draft Compliance and Scope document will serve as a starting point for governmental regulators, who can, when and where they so decide, incorporate elements from the W3C draft into their respective regulatory requirements. This may indeed fracture Do Not Track into various regulatory regimes, and while this may not be an optimal outcome, it may in fact be the final result of this process. The end goal of one single global standard is the right one. But whether the W3C can itself establish a meaningful Do Not Track compliance policy, based on the discussions that have taken place, seems unlikely.

I believe the Working Group needs to be realistic about what it is able to accomplish on its own, and seriously contemplate how formal regulatory action can complement the technical consensus that the Working Group has already achieved.