

The Policy Implications of End to End

December 1, 2000

Stanford Law School Center for Internet and Society,
Stanford, CA

Panel 4: Dave Clark Christine Hemerick, Ann Brick, Doug Chad and
Tim Denton

LARRY: All right. Here's the final exercise panel which is not supposed to be raising all these deep fundamental policy questions — just as a much more superficial level seeing trade offs. And the trade offs that we're talking about now in this context is the context of caching.

And so I'd like to as Christine — first of all, people introduce yourself and where you're from. And then Christine to introduce that trade off question of what various architectures or caching might implicate e2e.

DAVE: Dave Clark from MIT.

CHRISTINE: Christine Hemerick from Cisco systems.

ANNE: Ann Brick from the ACLU.

DOUG: Doug Chad from the University of Minnesota, Law School.

TIM: I'm Tim Denton. I'm a lawyer and policy type in Ottawa, Canada.

LARRY: Okay. So Christine.

CHRISTINE: Thanks. I want to take just a minute to make just a couple of general comments because I've been sitting very quietly over in my corner over there. But on this general subject of e2e — ness and whether or not by breaking the e2e model, we create a fatal, ethical, moral sin. I think it's important to remember that a lot of the growth in this whole industry, data networking, has really been driven our ability inside the core network to break the e2e model in ways that would favor economically overall growth scale and development of other applications.

Networks that are successful and consequently that grow by definition create a scenario in which the overall benefit and priority of the user community at large differ from the benefits and priorities of any one individual or any one sub community of users and network.

And so you will always have — I think a lot of times we talk about the e2e-ness. There's a lot of focus of what's best for a particular end user or a particular sub-class or of user when the trade offs are much more complex than that.

My company certainly made a relatively nicely successful business out of our ability to insert features that (??) network that made up for the fact that for example — it's very difficult to take thousands and thousands of end systems and simultaneously upgrade them at one time to put all of the new functionality that you'd like to put into the end system.

Had we not been able to implement band to a management [Sounds like] capabilities within the network, nine, ten years ago, a lot of applications that we see running on networks today would never even have been implemented because there were other applications that were not properly behaving and therefore did not give the opportunity for new applications to grow.

So I think that — I just reject the notion that somehow e2e is always good and breaking e2e is always bad. The situation is much more complicated than that. And a lot of it is about the relative economics of what's going on...

LARRY: There's a clarification question here.

QUESTION: [Off mic] because I actually think that Cisco has almost never broken e2e and it's a really interesting for (??) approach to argue that since we've broken e2e and the world has turned out well, therefore e2e is not a good thing if you've never actually broken it.

CHRISTINE: Perhaps that goes to the ambiguity of what people mean by e2e. But certainly we were amongst the first to implement NATS and firewalls and caches and load balancers and manner of systems that intervene between what we've been describing here this morning as "e2e".

QUESTION: [Off mic]

LARRY: It turns out I don't either.

[UNRELATED DISCUSSION]

QUESTION: Some of those things that you're claiming have been implemented with no harm, there's actually no proof that they've been implemented with no harm. And some of them, like (??) management, are not actually violations of [Off mic].

LARRY: But let's focus on one which is the caching one.

CHRISTINE: And I think to be specific about caching — in my view, the future of caching technology and what's referred to now as a whole set

of technologies that are about content, delivery networks — that whole technology has at its heart this kind of perpetual, yet another arms race between the declining costs of band width and declining costs of CP and memory.

And over time, while there have been a couple of exceptions, and it's my personal belief that history has shown that there are going to be a glut of band width at any one period of time. Or as a long term solution to the problem have lost.

And again, bringing in the issue of thinking about this not just from a parochial view [Sounds like] point of view, but from an international view, I think when we look across the development of the internet on the international basis, the disparity between the availability of band width on a global perspective has never been greater.

And when we talk about issues like digital divide, digital divide's not just about access. Digital divide is about applications. And I think that the future of caching and content delivery networks is a critical trajectory to our being able to make a wide variety of applications available on a universal scale because we're simply not going to be able to uniformly have lots of bands width available in all places where we'd like to deliver many of these applications.

PANELIST: Can you say what a cache is for people in the room?

CHRISTINE: I guess in the broadest sense, a cache would be an intermediate, somewhat temporary storage point for information in the process of delivering information from one end point to another. I think there are — again, one could debate about where do you draw the line between what is a cache as opposed to what is a store and forward server.

LARRY: Okay, so we've got....

QUESTION: What is a buffer?

CHRISTINE: What is a buffer? Exactly.

LARRY: So we've got a claim — and I don't think there's a lot of contestation about this — there are advantages provided by caching that would be more important in some areas of the world than in others perhaps and have not been eliminated by endless band width. But let's now focus on are there now ways to implement the caching such that it's more e2e compliant or less e2e compliant?

[UNRELATED DISCUSSION]

QUESTION: [Off mic]

PANELIST: That's part of the question Larry just asked.

LARRY: How can we implement caches? So, Christine do you want to say something? Or Dave...?

CHRISTINE: On that particular...?

LARRY: Ways to cache.

PANELIST: Transparent versus directed?

LARRY: Right.

PANELIST: I don't have organized notes. It was a communications glitch. I'm inventing my thoughts as I go because I don't have anything prepared.

At one end of the — caching makes sense in the context of a class of applications that broadly look like content distribution. It's hard to imagine how you do caching in the middle of internet telephony. Caching preemptively is solving a problem — two problems.

One of them is that retrieving a page from a distant web server over across a network of unpredictable band widths may put a lot of load on that band width and cause long delays.

It's also a case that if you're doing a website that turns out to be unexpectedly popular, if everybody in the world knocks on your door, your server melts. And so at its simplest, the idea of a cache is a semi-benign idea that if seven people pulled out a copy to a distant part of the network, why don't we save a copy the first time and the next six can get it from there.

In that respect, it tries to be a somewhat transparent mechanism. It isn't, of course, completely transparent in the sense that many people who produce web pages — especially when you look at web pages that have interesting complexity like advertisements and per consumer tailoring — don't necessarily look at caching as a benign thing because they don't get the hit count.

They don't get to know where you're coming from. They can't build demographic models so it begins to corrupt the view that you have of all these people knocking on your door.

There's another variant of this game which is often not called caching. And I don't know whether you ruled it in or out of bounds — it would be called content mirroring. And content mirroring is a slightly different model in the sense that in a cache, the only reason that you store the material is that you've made a probabilistic guess that it's going to be needed again.

In other words, it's sort of the network operator's decision that

it's worth saving a copy because this appears to be popular. In content mirroring, the reason why the content is replicated around the net is because somebody paid you in advance.

And so it becomes a service for fee whereas in the case of a cache, in fact, it may be very well simply the network operator is saying look, I'm doing this to try to make this system work better.

In both cases, you find attention between the people who are getting cached and didn't want to or the people who would have liked to be in the mirror, but can't afford to pay for it. So there's clearly a mechanism for discrimination going on here. Well, I think I see a hand, so maybe...

LARRY: Carl, did you...?

QUESTION: Are you drawing a distinction between mirroring bit for bit or [Off mic].

PANELIST: I did not make that distinction.

QUESTION: [Off mic]

PANELIST: Some people are saying yes. Some people are saying no. I'm looking around the room. What he said is you could imagine building a system who's doing something very simple which is there's a web page. It's a static thing. Ten people retrieve it. The first one pulls it all the way down, the next nine get it from the saved copy. You could imagine something much more complicated in which the web page contains descriptions on how to customize it. It's in some sense an active web page or — I'm trying to think of a word that would capture this.

And what gets sent down is sort of assembly instructions. To me, that's less caching and it's something more complicated because I think it involves more participation by the agents who are scattered around the net in the business objectives or at least the distribution objective of the original source.

LARRY: Okay. But what I'm missing in this conversation so far, we're starting at a technical issue. A technical level. The relationship between various ways of caching and this e2e principle. You're describing two ways of caching — one was transparent and the other was mirroring. Scott had a third.

PANELIST: This is part of my life, so I know how these guys — I work with companies that do this. It's a little more complicated than that. And just technically, it sort of blends into (??) business model, but caching is about [Off mic] cache stringing video because those are the most important [Off mic].

They've made a useful distinction between content mirroring

and caching. They are often called the same thing. For example, most people in the market think that (??) does caching and in fact, it does no caching whatsoever. It does content mirroring on a fee for this wire [Sounds like]. And the stuff that ends up in the cache is paid for by the supplier and nothing that's not paid for ends up in the cache. It's a very...

LARRY: It's a mirror box.

QUESTION: Is the only difference between mirroring and caching economic in terms of who chooses to do it and pay for it?

PANELIST: There's an architectural distinction. And that's what I was going to get to. That relates to the economic distinction. Alchemized [Sounds like] mechanism for doing so does not involve reading traffic and trying to analyze it on the way to the end user of end service.

Alchemy [Sounds like] reads the web sites — stuff off the web site — Alchemy basically runs an analysis tool at the source.

Looks at all the stuff and transforms it into an equivalent mirrorable content and when it actually tells the end user what to fetch, it gets the end user to fetch from the relevant cache other than getting the end user to fetch from the original source.

So it in fact transforms the content — it can personalize the content in the following simple way which not an AI [Sounds like] process or whatever — but just recognizing a web page is a symbol for many components dynamically already. That's the standard way [Off mic] bid for bid.

The assembly process happens on the user. So if you send assembly instructions are different per user. Almost all the components are cacheable. So a CNN web page that contains a bunch of pictures in it might contain different pictures for different users because the very tiny description that describes the overall layout might actually be personalized and might come from the original source not having a need to be cached.

But the 89 or 95 or 97 or 99.9 percent of the data comes from the cache even though it's selective based on the user's...

LARRY: Okay. But what's the issue?

PANELIST: So the issue about this is that the other — the alternative caching model which were the early caching models try to hook the stream. They try to use this Magic Toward 80 [Sounds like] mechanism and actually analyze the content as it flows.

And try to make a dynamic distinction about what makes sense to cache independent of the sources and in fact, independent of the destinations which is at the wish of whoever operates the

cache, which might be an ISP.

It might be in service of a better quality of service that they can sell to their customers which is sort of indirectly the wish of the user. [Off mic]

The key thing here is that the control comes at the result of violating the e2e nature of the protocol. In fact, the world wide web standards, HTTP, HTML and all the other stuff — do not suppose a cache. Whereas the content mirroring thing works within the HTTP, HTML standard on an e2e basis.

It's perfectly feasible to do the dynamic caching not just the economic model where the server pays — it's perfectly feasible to do the caching at the HTTP, HTML layer if you use a proxy. And the proxy stuff is part of the HTTP, HTML stuff.

And it basically makes the proxy server or the cache server an end if you will — and outside the network function that's part of a standard. But the problem Dave alludes to with the caching stuff not getting the hit counts, not getting the advertising revenue or whatever comes from the fact that, instead of a protocol between the ends that are interested, someone inserts themselves and corrupts the whole model.

LARRY: This is getting interesting. Corruption entering into the story.

PANELIST: Violation of the e2e model impinges on the presumed economics or the presumed behavior that was built into the standard HTTP [Sounds like].

LARRY: But working with the ends. Tim, did you...?

TIM: Some caches intercept packets that are not addressed to themselves. So the source, the browser is sending ID packets that it thinks it directing toward a web server. And some caches intercept those and answer them. And so one style of cache.

The other style is where you actually explicitly choose to send your packets to the cache or the proxy. And then if it doesn't have the answer, it forwards it on to the web site. And those are the sort of — and the first one is the one that violates the e2e law.

You're sending a packet somewhere and in fact, something else is going on. And that's the shifting sands that Dave was talking about. I put this thing — I told it never to do this — and it did something else. And I had no way to know it was going to do that.

If it's doing the wrong thing, I have no way of knowing what's [Off mic]

LARRY: Right. But without a lot of food, which we don't have right now, what's the problem? What's the issue? Tim?

TIM: Well, that's been my entire trouble trying to get my head wrapped around the question of caching. It seems to be a fairly sensible procedure for minimizing your bandwidth costs. And...

LARRY: So what's the evil?

TIM: Yeah. I'd like to hear some people talk about the evil of it.

QUESTION: [Off mic] what's interesting is architecturally generally, is the regional versus global or national internet division you might have. Because the news in the industry right now is that, when you look at broadband carriers, as much as possible, put content delivery architecture inside their local proprietary network.

And then this ends up being kind of like a quality of service discussion we had earlier. People who are more in an area where they have a cable provider or DSL provider, you put content delivery architecture inside that proprietary cable or DSL network. [Off mic] video on demand.

So what that means is what internet you experience depends on whether you're in an area that's served by a proprietary broadband network or not. And that's really interesting issue I think, in all of content delivery architecture.

It's not necessarily — I mean, e2e may be interesting to the network designers. But from a policy perspective, I think it's a regional — the original problem with the internet is you go anywhere and you get the same kind of content no matter where you are.

But what everybody in the industry is trying to do is to turn it to where you are depends on what kind of content you get at high speed. [Off mic]

LARRY: Jerry, then Anne.

JERRY: [Off mic] someone will figure out the correct answer. But one of the things I do is I want to buy a book from Amazon.com. I open a secure connection by using a special protocol [Off mic]. And they present me with a certificate.

And if I take the trouble, I can [Off mic] certificate from Amazon.com. But today, I discover I've got a certificate from Achimon [Sounds like]. Who's that? Because the e2e model broke down somewhere. I'm talking to someone in the middle instead of talking to someone at the end.

LARRY: This is disabling functionality that we expect. But I think there's a separate set of issues that I think Tim was also pointing to that gets raised by this.

CHRISTINE: I think it boils down to perhaps the questions of what's getting cached and why. And are there neutral principals applied to answer those questions and to what extent are value judgments coming into this? When I talk about neutral principals, for example, there can be a rule that there are no sound trucks allowed on the street. Or there can be a rule that there are no sound trucks that broadcast above a certain decibel level.

Or there can be a rule that only Republican sound trucks are allowed and no Democrat sound trucks.

LARRY: Sound principle.

CHRISTINE: And sometimes the judgments can be economic and sometimes they may be political. But those are the kinds of questions when you're asking what's getting cached and why. Basically, we want content neutral principle. You have to decide what values you want to further and to what extent are those economic and to what extent are they political?

And then are there neutral principles that do that that don't do violence to basically the availability to content to everybody on an equal basis.

LARRY: Okay, so we might be generally agreeing without political decisions affecting caching, although I'm sure we could get an argument going about it. I think the core issue here would be whether economic non-neutrality is a relevant thing to be upset about. And Jamie, you wanted to jump in on that.

JAMIE: I probably came here with less armor than anybody else. But let's just try to see if I'm understanding this correctly. My turn on the board. It seems to me as if there's two types of caching going on. I like this example of Jerry (??)

Here's Amazon — that's one end, right? And then let's put Jerry over here. And it looks as if they're two potential places where the caching's taking place. One of them — the one that Jerry mentioned is at Achomine [Sounds like].

And if I'm getting this right, Achomine is engaging in the what's called the active content mirroring. Taking a lot of Amazon stuff and saying, hey, instead of overloading Amazon because everyone wants a part of Amazon, why don't you to come to Achomy [Sounds like] which is a much bigger server and then we'll take care of that for you. And then whether the (??) coming from — no, that's no how it works. Maybe I'm wrong.

QUESTION: [Off mic]

JAMIE: That's my point. Achomy is charging Amazon for the service, right? So it's just a mirror image of the Amazon content on Achomy along with any number of other providers of content.

Now, the other kind of content caching which I understood and to be worried about — and this is perhaps the more constitutional types are worried about — is at the ISP level. Or it doesn't even have to be at the ISP. The example I've seen in one completely bizarre context is the government of Singapore. The government of Singapore says there's a lot of stuff we don't like here. Asian values and all that.

Any reference on this world wide web we don't like that refers to our government in a negative way, we'd like to filter that out. So why don't we kind of combine these two functions.

What we'll do is we'll take the content caching ability — I mean, the content mirror ability, except that we'll apply it to the entire web. And we want 99 percent of it, not just the one percent that deals with things about repressing dissidence in Asia. We don't want the one percent that deals with people who are too naked or naked in the wrong kind of way.

LARRY: Let's not go there.

JAMIE: Just as a lawyer who doesn't have much of a technical background at all, sounds to me as if you've got not only the act of passive distinction, but also distinction of whether it's being paid for by someone who's originating content versus somebody who's just trying to get the access. And the policy issues are dramatically different.

LARRY: Okay. That's very helpful.

QUESTION: If I pay Time Warner to cache for me, [Off mic] somebody's paying Time Warner.

JAMIE: I don't mean to say that...

QUESTION: Caching narrowing distinction — you've got an exactly symmetric process except in one case, I'm the receiver implicitly paying, [Off mic].

PANELIST: Your third party's paying the ISP to do something.

QUESTION: Somebody is paying at one end or the other.

[UNRELATED DISCUSSION]

PANELIST: Let's get it from Jerry's point of view. Jerry was worried about (??) I think Jerry might have a different kind of problem here. Let's supposed he's dealing with the ISP and there are two

types of ISP that Jerry could deal with. One of them is simply hey, happens to live within the empire of Time Warner — and that's the only cable modem he can get is Road Runner.

So he just wants Road Runner. He thinks it's transparent — it's neutral. It's just going to bring him his content. That's what he thinks he thinks he's getting. But it turns out that Road Runner because of some convenience issues says well, to speed up band because...

PANELIST: [Off mic] the fastest band with possible solut — get it to you faster with a narrower pipe if we have more aggressive caching.

PANELIST: Hold on a minute here. With whom are you dealing, which seems some kind of issue of authorization and authentication. I don't — and therefore the issue may have some importance. I'm trying to find out what the importance of what the caching issue is. And I'm not sure it's caching as such.

LARRY: Let's bring it back to America a little bit. Let's say that Road Runner says these are our preferred suppliers and we're going to serve these suppliers much faster than everybody else.

They've obviously entered into a non-neutrality into the service provided over the "internet". What the problem with that? That's the obvious question which the lack of food is taking us a long time to get to. Yes?

QUESTION: But the point is that Achomy is discriminating as well because I didn't give them money. So what? The problem here isn't whether you have an ISP [Off mic] doing the caching or whether Singapore's doing it or Road Runner's doing it. The problem is transparency that (??) was bringing up.

In a transparent cache, there's nothing the user types in to make it not be transparent, not transparent, use a proxy, not a proxy. The URL [Sounds like] is the same. In the configuration of the browser, you set it up to be a proxy or a non-proxy.

If Singapore is running a transparent cache, they do not cache or permit certain kinds of sites from being accessible. You cannot get access to them. That's the civil liberties problem.

LARRY: That's one civil liberty's problem. But Mark?

MARK: The reason that (??) reacts differently to these two models is because of the idea of who gets to control the packets, right? And Jim was perfectly happy to say Amazon is free to contract with Achomy because it's Amazon's [Sounds like] happy.

If Amazon wants to give it to Achomy and let them run the server soft site, wonderful. But Amazon didn't contract with the ISP. And so if the ISP sets up to alter the content or even to

locate it in a different way, so that it maybe turns out for example, Amazon engages in price discrimination. Right?

And so if you just blindly cache the last person who looked at this book, they get charged a different price than Amazon.

I'm not personally too heart-broken by that, but I can see why Amazon...

CHRISTINE: Reasonable. I mean, come on, the technology is such that reasonable cache is in the area that you mentioned. It can cache certain objects of a web page which are not user specific while not caching other parts of web pages which are user specific.

And the notion of transparent caches, the concept of transparent was originally an issue of configuration. Transparent caches were designated as transparent because you could put them into the network as opposed to the proxy situation.

And virtually with very minimal configuration. Certainly no minimal configuration required of the end users. You could insert them and they could do what they want to do.

Now, for providers whose business model is one way where for example the vast majority of the revenue is driven off of advertising or something like that, a transparent cache which goes in very cleanly, configures very cleanly, delivers the desired advertising to the end user, doesn't violate the business model, I don't think represents any civil liberties concern.

And in many cases, is actually welcomed by the content provider because it allows more users better response time and better access. I think the issue comes up when you have a situation where the cache either because it isn't adequately — doesn't have the adequate functionality to count hits. And by the way, you can also do hit counts in other things in caches and provide that information back to the content provider.

I think the situation that's of concern is where there is a particular business model and that a particular implementation of the transparent cache which don't match. Where you have the potential that somebody loses revenue or somebody loses information that's important to them.

QUESTION: But Amazon owns this. And Amazon can authorize caching whenever it wants...

CHRISTINE: Amazon can cut off anybody it wants to. It look at where requests are coming from. But this isn't purely a business model. It's part of the problem. A lot of the information on the internet is for information and influencing people and

exchanging ideas. And caching can have a real effect on that in terms of whose stuff gets cached and speeded on its way. And whose doesn't.

LARRY: So commercial is fast. Non-commercial is slow.

CHRISTINE: Exactly.

LARRY: So that's the sense in which this could affect other issues other than just innovation. Now, look, here's the structure...

QUESTION: ...bigger issue that David Clark alluded to, but then we lost it. And that is in any kind of passing, there are some kinds of traffic that are inherently non-cacheable. And those are included. And the best one is internet telephony where you can't cache it, you can't delay it. It has to occur in real time. And so that means that any caching that goes on has to make some kind of determination about what's this content.

That to me is the really big red light about caching.

LARRY: That it's making this determination.

PANELIST: Hold on a minute here. David, when people cache, they make a choice of some kind that web pages are going to be cached.

QUESTION: [Off mic]

PANELIST: Yes, that's the purpose of caching. So far as I understand it. Why are you assuming that there's somebody making a determination about that this — what's your problem with someone making a determination that the cost effective way of delivering this content is through caching.

DAVID: That's if you assume that all web traffic or all internet traffic is content. But a lot of it isn't content. Most of it — and increasingly more and more of it — maybe not most of it — more and more of it is communication.

PANELIST: Granted. And so why do you think that that would be cached?

DAVID: [Off mic]

PANELIST: Of course not. But why do you think that it would be?

DAVID: Well, it would be if you are blindly caching everything.

PANELIST: But you're not.

LARRY: David and then Mark and then I want to wind this down.

DAVID: I think the argument is more subtle. You know, we've had a theme that's come up for air and then gone down again a couple of times which is thinking about over provisioning as a solution. And clearly, the people who cache to some extent are doing something subtle. They're reducing latency and so forth.

But the gross argument that motivates caching is an argument which would go away if we had lots of capacity because aside from the fact your server might melt if you inadvertently underestimated how popular Madonna was, what's wrong with fetching the page ten times if the backbone is really over provisioned.

And so I think the argument which is a subtle one but is worth watching is that people are putting their money behind the applications more and more that have proved their popularity. And what that means is there's less and less money going into upgrades to the infrastructure that could be usable in more general ways.

The dollars spent on caching does not improve in it telephony [Sounds like]. Caching doesn't break internet telephony but if in fact, people are investing in caches instead of more backbone capacity, then application that needs a significant amount of capacity and can't be cached may not take off because the infrastructure's too feeble to launch it.

And part of what you would expect, but you could also be unhappy about it as the internet grows up, people more and more put the money behind the known value and less and less behind the general purpose. I don't know what it's for, but let's build it anyway experience.

LARRY: So this is a great...

QUESTION: Just one amplification on that because it's not just about band width investment, it's also about investment in the web infrastructure itself gets biased toward the thing that caches do well. Which is static content.

Even though you can handle some dynamism [Sounds like], for example, downloading very personalized stuff does get discriminated against. And just like the argument of information as opposed to e-commerce or...

LARRY: So here's another dimension of discrimination that's possible.

QUESTION: [Off mic] discrimination that's going on. But when you help one thing, but not another.

LARRY: Mark, I want to wind this down here. Mark?

MARK: I think [Off mic]. In the e2e model, the two ends are processing the content. And what we're talking about here is something that is invisibly inserted in the network that dabbles [Sounds like] with the content. Either by storing it, by modifying it or something like that. [Off mic] when the Postal Service opens up the envelope for whatever reason, may or may not change something, may

inspect something, may cache something. Puts it in a new envelope and sends it off. Would you trust that service?

LARRY: Okay.

PANELIST: I know you want to wrap this up. I want to do this pivotally and aphoristically. In real architecture, the modernist slogan was form follows function. Louis Salvan [Sounds like] in post-modern virtual architecture, Larry Lessig [Sounds like] says function follows form. And what's happened here is what appears to be a very simple short-term stop gap solution to band with scarcity creates multiple opportunities for people to get in there and start playing around with content.

And those intermediaries being present in and of themselves destroys the entire e2e concept, conceptually at a very high level of extraction. And now we're just arguing about the details.

LARRY: This is very helpful. Here's the way I want to draw together a description of what happened this morning. We've shifted from one to another very different kind of cost that gets imposed by deviations from e2e values, let's say.

And we started the morning talking about the cost to innovation. Cost of these deviations making it hard for innovators to develop new technologies. And what we've ended with, both in QOS and in caching, is a separate kind of cost which is the potential for a different kind of discrimination to be brought into the network that either makes it harder for non-commercial speech to have some play here. Or aggravates monopoly problems that might exist given the network providers.

Now, those are very different problems. And we start in the afternoon, after lunch, with Jerry — who is going to frame the economic question I think in a way which helps us think much more helpfully about that second problem. About whether there really is an issue to worry about when we worry about the discriminations that facilitate one way of doing business over another. Or commercial over non-commercial.

But I think this morning has usefully distinguished those two. And they're not answered by — (??) by answering one of them. Okay, now. That was a pretense at ending the morning conversation in a way that makes it sound coherent. But it felt coherent to me.

It's 12:30. And at 1:30, we will continue and there's box lunches and you should feel free to enjoy the Stanford campus. And we don't have any program during lunch. So please just be back here at 1:30 because we've got a lot to do in the rest of the afternoon. And some of the most important.

