

The Policy Implications of End to End

December 1, 2000

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

Panel 3: Barbara van Schewick, Karl Auerbach, Jamie Love, and
Stuart Benjamin

LARRY: Okay, the subject of this panel is quality of service. An issue raised by the fact that the initial architecture was a best efforts system. And for some applications, best efforts is just not good enough.

And so this creates the inducement for people to develop other technologies for getting beyond the best efforts paradigm. So we're going to start again as we've done before with an attempt to describe at least two — and actually this time it'll be three architectures for addressing this particular issue.

I'm sorry. What we should also do is we should start by having everybody just say who they are and where they're from. So Barbara, why don't you start?

BARBARA: Okay. My name is Barbara van Schewick. And I'm a visiting scholar at the [Sounds like] law school.

CARL: I'm Carl Auerbach from Cisco and at Cisco, I research into things like autopilots for networks.

JAMIE: I'm Jamie Love. I work the Consumer Project and Technology which is a group Ralph Nader started.

STUART: I'm Stuart Benjamin. I'm a law professor currently visiting at the University of Pennsylvania Law School.

LARRY: Okay. So Barbara's going to start off with a description of these three potential responses to this quote problem.

BARBARA: Yeah. Maybe I'll start very briefly for all law people in here to explain what the best efforts service model really means. It means that in the internet, all data packets are treated the same way. So if there are a lot of data packets, everybody slows down. If there are only a few, then everything goes very quick.

And now with real time applications, you have some applications that are very sensitive to delays. For example, if you want to talk to somebody on the phone, people have found out that if it takes more than 300 milliseconds between the time

you talk and the receiver receives your words — and it doesn't really feel like a conversation at all.

If you want to have industrial control, then you would like the command to reach the RAM of the Robot [Sounds like] before it crashes into something. So these are the main motivations to say well, best effort isn't enough. We need certain guarantees from the network for certain types of service. And these guarantees are called architectures that provide these guarantees are called quality of service architectures because they provide different classes of service to different types of data.

And there are basically two architectures which have been discussed by the engineering task force. One is the integrated services architecture that provides quality of services to individual applications or individual flows of data. And the other one is a more coarse grained approach. The difference (??) architecture that provides quality of service to different classes of service.

And I'm going to explain very briefly what that means. Both architectures basically have defined certain specifications. How they should be treated if it belongs to a certain class. And in the integrated services architecture, what is done is that certain reservations are made within the network for a specific data flow.

And before a data flow enters the network, the network has to decide whether it will be able to provide the type of service that has been asked for. And if it can, then all the reservations will be made within the network and the data flow will be admitted. If it cannot provide that kind of service, then it'll say well, sorry. Well, nothing will happen.

And to do that, you need to move a lot of functionality into the network because as you can imagine, in order to have reservations in a particular network, you need to store this information for every flow that's going to a certain network node. And at every computer in the network, you have to deal with the flow and decide what type of service it should get. And maybe add some other functions.

So this means a lot of functionality in the core of the network and that constitutes a huge deviation from the insurance [Sounds like] principal. And the differentiated services architecture — what you do is you classify data at the edges of the network and then in the network, they just treat it according to this classification. But the more complex functionality is left to the edge of the network which may be the boundary of an internet

service provider's network. But may also be the edge of a bigger network.

[UNRELATED DISCUSSION]

QUESTION: How is the classification then done? Does everybody get individually de-classified? Is my package really important or is there some sort of centralized classification system even though it's implemented [Off mic].

BARBARA: That's what you need to decide when you implement that system. You could imagine different ways to implement that. David, you want to say something to that?

DAVID: For example, if you look at DifServe [Sounds like], there's a clear idea in DifServe that there's a new kind of element added to the network which is...

LARRY: I'm sorry. You said what? DifServe?

DAVID: Differentiated Services.

LARRY: Okay. Yeah.

DAVID: The assumption is that you cannot [Off mic] in your head arbitrarily. The assumption is that there is some element which is inside the network boundary. That is to say the trust is associated with the network which is policing or verifying or constraining how you consent [Sounds like] with this.

Over an internet, that might be done on the basis of economics that you pay for the privilege. It's [Off mic] and so forth. But (??) is the idea that you didn't give them the power to be arbitrary. You didn't let everybody ask for a Cadillac. So there has to be something in the network that enforces some policy.

BARBARA: Okay. So what we've seen from this is that both of these architectures need functionality in the network's core. But to a different degree. If you want to implement integrated services, you need to move a lot of functionality into the network.

If you want to implement differentiated services, then in a lot of end nodes, you just need to implement the functions that treat packets according to the classification as the more complex classification of packets and the adjustment of the flow to the traffic policy is done at the edge of the network.

And then there's another alternative that isn't often discussed with respect to quality of service. But which could solve the same problems which is just adding more capacity to the network. Because if you have a huge capacity, then all applications would be very quick and maybe you wouldn't need to have different classes of service.

And that's obviously the solution that wouldn't change the architecture of the internet as it is at the moment. So we have three architectures which are compliant in different ways. Over provisioning, which is perfectly compliant; integrated services which is not at all compliant; and then differentiated services, which may be some of a middle way [Sounds like].

QUESTION: I just want an explanation of why integrated service is not compliant e2e. It's [Off mic] definition of e2e, I think, that's important in [Off mic].

BARBARA: Okay. The definition of e2e that I'm using at the moment is that — the effect of e2e is that intelligence migrates to the edge of the network. If you want to do application specific functions as high in the layered network as possible, then that also means that the application specific function migrates to the edge of the network.

So you could say that, in an e2e network, you have simple core of the network like...

QUESTION: That's a key, though. It's a simple network, e2e meaning that the decision on how to implement applications at the edge is what I would argue that in a network that provides multiple — best effort and multiple classes of service — that doesn't have to violate e2e. You see still [Sounds like] do all of the application design innovation at the edge. That's important I think in this context.

LARRY: Okay. Scott and then David.

SCOTT: I'd like to actually phrase the difference differently. In integrated services, the network elements in the middle of the net are aware of individual flows on the net. So it is your individual phone call, your individual instance of application. This is putting per flow awareness into the network.

And that is something that in the e2e model doesn't occur. The network element should not be aware of individual flows in the network. Differentiated services does not put individual flow [Off mic].

So all of the phone calls are treated as if they were one phone call. Whereas in integrated services, [Off mic].

LARRY: David?

DAVID: Well, it's actually a fairly subtle design area. Just because you build a simple mechanism and put it in the center of the net doesn't mean you win.

LARRY: Win what?

DAVID:

[Off mic] but what you put in the center of the network has to be right. And in some respects, you can say that either architects are lucky or [Off mic], you could have screwed it up.

Integrated services is slightly less impish [Sounds like] in one way which is you set up a reservation that [Off mic], and therefore there is some pre-standardization that characterizes what those reservations are. If you want a reservation [Off mic] before, you can't have it because the founders never heard of it.

In exchange for which the founders can enforce the characteristics of that reservation very precisely. So the sacrifice you make in (??) is again — I think it's a tail of a curve [Sounds like] sacrifice. It could use a little bit of innovation because you might someday want a reservation clasp [Off mic].

DifServe is a different technological approach which is it takes traffic of potentially different kinds. And it blows them [Sounds like] together into one glob and treats them as aggregate. And a bad engineer would love to gather inharmonious [Sounds like] classes and everybody would be unhappy. A good engineer would create multiple DifServe classes and promote harmonious traffic.

You can't control it as well because what you're doing is controlling an aggregate. But in exchange for that, at the edge of the net provided you can characterize the impact of what kind of traffic on the aggregate, you don't have to tell the center what that sort of profile of the traffic looks like.

So that the feeling behind DifServe is you've got more flexibility by pulling the characterization of that traffic out to the edge in exchange for which it moves the ability to absolutely police it. And it's a trade-off. We've had [Off mic]. We're killing each other (??) because they said I can't ship my product unless you give me [Off mic].

So nothing stops me from shipping the product unless there's standardized mechanism [Off mic] going to kill you. And it was an interesting thing to try to propose a general technical compromise to see if we could get it all the way through the standards process.

LARRY:

Okay. But the trade off you've described here is perfectly consistent with the general point about e2e. And so the extent you're compromising the optimization of the network to preserve the opportunity for innovation — that's the e2e — that's the more compliant e2e strategy. And the only costs we have identified so far with the integrated — with the IS structure is the slowness or the cost to innovation of having legacy systems that

don't really yet understand some new particular use.

PANELIST: There's also stealing [Off mic] to the amount of statements kept [Sounds like].

LARRY: Okay. Well, this is the question. What other consequences might we have from this? Jerry?

JERRY: There's another element that I think also has to be factored in here to kind of understand exactly what you're getting in return for what you're giving out. The world's a dangerous place. This is for packets and [Off mic] everything else.

And there are no guarantees for the real world. And therefore, [Off mic] actually is not true. There will still be some things that get lost. And because that's the case, all you can actually do is push the probabilities around to different positions.

[Off mic] because some of the arguments that are made in favor or in some of these (??) are making the assumption that if you only implement this feature, it will be guaranteed that people [Sounds like] will get this kind of thing.

LARRY: But help me in showing where this argument cuts on the difference between these more compliant and less compliant technologies.

JERRY: It seems that a more compliant technology may [Off mic]

LARRY: A slight amount.

JERRY: [Off mic]

LARRY: Okay.

PANELIST: I think another way to sort of characterize which I didn't quite get out of Dave's distinction between the two — InServe and DifServe [Sounds like] — is that the InServe model is extremely application specific. The idea is that you propose a particular application that exists today and you want to solve this problem. Like telephony [Off mic].

Telephony is fixed rate isochronist [Sounds like] strategy. Isochronist meaning that the time delay from one end of the network to the other is always identical for every packet that goes through.

In order to — you might another description of something that's very like isochronist, but has different, for example, error bounds on isochrony [Sounds like]. And claim that it is a different category because different application of different things whereas in DifServe, you would lump those two together into a class.

And require more adaptation at the edges. Basically, you would force the edges to adapt a little bit more. So, for example, you might have telephony that only has some [Off mic] buffering at the end. I'm getting really technical, but all this talk about buffering at the end. And any more buffering and that would be too costly to make [Off mic].

And then another type of thing which like voice phone where the amount of buffering is much larger and perhaps the errors would be different. And even though those are isochronist categories, they're not in a structured space. And the trade off that you build into the network becomes application [Off mic] a variety of integrated services and a variety of providers. And the edges force trade offs into a network that might not have otherwise been [Off mic].

I believe that it's sort of a management issue also around this — this sort of forces you to think in advance of what are the really fundamental different categories of things you might want to put through the network that you can't anticipate? Build some categories and then offer them for general use.

As opposed to having say telephony service providers [Off mic] a category and be able to define it in a closed net cost on the network.

PANELIST: [Off mic] integrated services architecture is fairly generic in terms of how you characterize the circuits. You don't say I want telephony service. You specify it in a (??) way.

PANELIST: But there's a small number of...

PANELIST: The battle has been that people want lots and lots of parameters [Off mic].

LARRY: Okay. There's 46 parameters in Net Serve.

PANELIST: I can't help but notice that you've all very quickly dismissed the third architecture that was put up. [Off mic] The original telephone network was actually the only true implementation of end twin. There were no switches. You had a private line.

You still have private lines, which are true e2e at least at the physical level. All the rest of the discussion from there on out is just not intellectually honest unless you say we're actually making [Off mic].

The only pure e2e is hard wire and dedicated circuits everywhere. So there's no technical debate to this. It's economic...

LARRY: There's actually a twist to this. Who gets access to the cleaner circuits if we introduce noise because it isn't all time...

PANELIST: The SS7 Network nails up a (??) and it gives you much better e2e service per voice than anything you can do on the internet. Everything from there on is deprivation for now [Sounds like].

LARRY: We're identifying what issues — what things are being traded off. We've had a very helpful I think technical description about relative flexibility innovation between these two kinds of designs. So Peter, you've come on and said there's that whole other design that you ignored here which Barbara originally described which is...

PANELIST: [Off mic]

LARRY: Right. Right. But how do we think about that trade off? What is that stake to think about that trade off?

PANELIST: I'd just like to solicit some clarification about the trade offs that DifServe imposes that say, an over capacity architecture would settle or would [Off mic]

LARRY: Don't think of DifServe and InServe as mutually exclusive. Integrated services signaling is often used at the edges. And as the traffic is aggregated — which is the core of the nets — we often shift to a DifServe type mode.

PANELIST: But let's assume we have a pure DifServe architecture. There would still be trade offs.

PANELIST: If they're actually the same in the sense of what DifServe does is it lets you build [Off mic]. If you over subscribe by giving more capacity to the goal of service [Sounds like], that's equivalent to over building the underlying service.

Putting in more fiber is equivalent to pushing up the percentage utilization for individual class of (??) — the same way.

PANELIST: In order to get people to pay money for the premium service, at some point you have to make the bad service crummy. And you have sort of a dis-incentive to...

PANELIST: You could give it a more positive...

PANELIST: Selective fairness.

PANELIST: Well, yeah. Except that if it's a premium service, you don't need the premium service. You're not going to sell many units. And I think that one of the problems you've got is who's providing the service?

Within Ford it's one thing. If it's a Time Warner to the customer and ISP, the incentives change quite a bit. So I think that we have to be realistic about what's driving this process.

PANELIST: [Off mic]

LARRY: All right. But let's bring out a little bit more about why it's important — why you might need to have different quality of service here. And therefore need to find ways to discriminate.

PANELIST: Yeah. When I first saw this debate develop several years ago, it was about do you want to have a go fast switch that you'd pay for? And to make packets and a lot of talk about congestion.

And it's sort of a classic problem where you have high fixed cost in the network — low marginal cost and you've got congestion cost in the architecture. And so economists like Varian [Sounds like] and other people immediately said okay, congestion costs. Let's charge money for congestion. That's how you address the congestion problem.

That is sort of the initial take. And then there was a bigger reaction against this because people felt that they liked not paying enough [Sounds like] to do things on the internet. They liked the fact that all the incentives were for the ISP's to continue to build out capacity because that's how they competed against each other was by the performance which was based on big capacity.

And they saw themselves heading in the wrong direction. And they saw this as a back door way of introducing the kind of payments for content and services and transaction which had been not only been missing from the internet, but actually people enjoyed the fact that that wasn't working very well on the internet.

And then there was this other discussion about this sort of virtual circuits on the telephony thing where people thought you could have sort of real time bidding for virtual circuits on the internet. And it turned out I think the cost of negotiating the prices and billing and things were astronomically higher than basically the cost of over capacity.

So again, there were a lot of people who didn't want that to work either. And some people did and some people didn't. But there's obviously huge issues, not the least of which, is what it will do to content markets in this sort of possibilities for monopolization and rent seeking and Bi-SP's [Sounds like] and sort of changing the entire character of information flows on the internet.

LARRY: There's also the aspect of doing, I'm going to watch Steve cringe — of doing MPLS virtual circuits on top of the internet as another way of doing traffic engineering. That's for longer live flows. And you'll get a lot of objection to that approach as well. Stuart?

STUART: The concern that I have, maybe on Peter Huber's point, is we're

talking about all the possible advantages of these different kinds of modalities. One, we could say let's just pick one and we'll go with it. It seems that a lot of people would think that might be premature and we'd like to actually have some experimentation.

If we had multiple networks, there's an easy way around this, right? We could have a whole set of different networks and each of them could have their own systems. The thing we think there are certain positive externalities that we get by having a single network.

And then the question becomes how do we want to structure it and who do we want the we to be? And I guess as the designated lawyer here, I'm not being clear?

QUESTION: [Off mic]

STUART: Which sort of gets to my point. I'm sorry. I leaped ahead. My point being when we try to move, when we have as the law — the law should actually have some significant role to play in the e2e network in making sure the e2e network will remain e2e. Then we have a we that is actually imposing a structure upon us rather than having a series of competing alternatives.

Or even if we ended up with one alternative, we have one the idea of being good code like from various people in this room deciding it rather than from external organization which might have less confidence.

And I suppose the concern that I have in listening to this is given all of these difficulties, what makes us think that a non-technical regulatory body — mainly Congress would be leading the way — would actually do a better job at this such that we think it's actually appropriate to cut off the debate? I see a green card, but...

QUESTION: [Off mic] the question is not well, we either have sort of all of these different networks and they compete with each other or we've got one network that's regulated in a particular way. They're all our means and the internet it seems to me has demonstrated precisely that means by which you can have a single network and the benefits of interconnection between them. And yet have experimentation within it.

And it seems to me that the e2e architecture is precisely that means. And so if you talk about well, are the regulators insist upon a particular structure, I think you're asking the wrong question. The question really is are we going to have a technical structure in which we can continue to have experimentation or are we going to be forced to a choice?

The only way we get to keep the options open is to have

regulators protect the e2e architecture. I'm not sure that's such an intervention...

LARRY: ...insisting on over building. That's their regulatory strategy.

QUESTION: [Off mic]

PANELIST: I'm questioning what you mean by e2e? There's selective deployment of technology. What is — it's multi-cast e2e — that part of it we're seeing selectively deployed. QOS is ripe for deployment by different ISP's.

BARBARA: But I think the issue we have to think about from a policy perspective is how much control do you want to give to network owners, internet service providers because what quality of service architecture means is that someone classifies traffic. And someone asked before, who does the classification? It's not just that you say oh, this is simple.

And traffic (??) going to go through very fast. But someone is going to say and this is important traffic. It will go very fast. And this opens up as you said a large potential for discrimination, for example.

I think Cisco in some of their marketing documents once said this gives you the potential to slow down offerings by your competitors and to speed up your own offerings. And this is — maybe in not a very realistic scenario, but they have advertised that. So this is something...

LARRY: All you have to do is...

BARBARA: ...some of the policy issues that need to be thought about. There may be a lot of technical arguments that say that quality of service may be implemented in a way that's very e2e or not very e2e.

But the value of that control is more at the edge and not within the network. It's only really saved by just adding capacity to the network and then everybody can have quick transfer through the internet. And then you need to find out whether that's an option or not. Some people say it is and some people say it's not.

PANELIST: I'm objecting to some of the premises here. When you talk about e2e, how long do you get it for? Is it the life of the application? Is it only for (??) Because we're allocating resources — valuable assets within some series of carriers between hither and yon.

And the routing typology could change. Something could break. And are we going to suddenly demand that a third party carrier that wasn't involved in initial allocation of resources suddenly

has to dump new resources into this? Provide an application with guarantees of services?

A lot of this falls back on the application writer to be able to say uh oh, the quality of service I asked for is no longer available. I'd better find Plan B. Application writers have to change their attitude. And also realize that the network can and should push back on them.

LARRY: David and then Bill.

DAVID: [Off mic] I want to follow up on this. I did want to make a comment about the concept of what were provisional. I think that if [Off mic]...

PANELIST: I object [Off mic]

DAVID: It's important to remember that we're not provisioning a static network here. You talk about the rate of growth of the internet in terms of how many users you have.

But it's important to remember that there's also what I might call pants us [Sounds like] application behavior. Why did MP3's happen? Well, it didn't happen until there was enough capacity that a significant population could take advantage of it.

And then all of a sudden, certain users, capacity went up by a factor [Sounds like] of ten. And you say well, I'll just over provision and then we'll get internet television and we'll get HTTP television.

And so all the provisioning is a then daunting objective (??) willing to chase the applications [Off mic]. I think that DifServe is less about giving — making the tail of the curve better for the good user. But is more about regrettably creating values for edification and allowing people to pay less in exchange for getting more quality of service.

And you wouldn't think of that as [Off mic]. You think of it as — if you block television service, that's fine. If you block MP3 service, that's fine. If you block [Off mic]. And I think that kind of stratification makes economic sense [Off mic].

LARRY: Okay. Let's make sure we understand the economic sense it might make. So, we can use words like discrimination and it sounds as if it's always or necessarily a bad thing here. And I think you're saying that it's not necessarily a bad thing to facilitate this kind of discrimination. Tell me the story why. Or maybe there's a policy person here.

PANELIST: I think it's just value pricing. It's value stratification. Let's go outlaw Saturday night stays [Sounds like].

PANELIST: I think you have to look at the market context. It occurs. And if you look at one of the big facts of e2e is the market penetration of Microsoft Windows and Microsoft Office. I mean, it's not like you don't have a monopoly problem on that end. And what is one of the tools that Microsoft has done to monopolize the application? Well, it's different levels of service.

Microsoft applications get thumbs up — great levels of inter-optability [Sounds like] with Windows. Other applications — lousy inter-optability with Windows. Now, cable companies — they make their whole bread and butter is basically keeping unaffiliated programs off the platform, giving them poor places and locations on the channel dial — even when it's in the first ten channels, everything is sort of discriminatory and things.

And so their whole thing is sort of shaking people down to get preferred — that's what they do. Now, you have to ask yourself, where we're ahead in terms of an intent structure on this last mile thing. And you have to be realistic about it. You can't ignore whether or not you're going to have problems there.

Because if you have a system that gives people the opportunity to degrade their rival services, and to promote their things, they will — and timing's [Sounds like] a scale, a huge part in the content market too because if there's enough consolidation even in certain parts of the country or just on the cable platform, but not the DSL — to disenvalue [Sounds like] certain kinds of content providers, that can shift the whole market to the ones that have that home market and trade policy.

They talk about countries — they used to sort of protect — like the Japanese — like the home market, so they could actually build up their economy as a scale so they could be competitive in the export market.

Well, you can do that if you have kind of a captive content market too. It sort of gives you sort of bigger units, sort of bigger home thing. And you can sort of do well. And in the cable industry, that's what happened in CNN and things like — not just CNN, but a lot of cable channels.

If they were on the Time Warner and the platform for TCI, then their numbers were pretty good. And it helped them even on the platforms that they weren't affiliated because they were good players. So, it's as big a communist [Sounds like] a scale in video programming, too.

LARRY: I have a cue here, Bill.

PANELIST: Just a quick comment. Since we have the original authors of the (??) Paper and they have a subtle argument that I think applies

here is I thought (??) was going to mention right there when we talked about dynamics and complexity. This is just a (??) situation where you're going to have clues coming on and off all the time. And some of the things with the NM [Sounds like] model allows you to understand what's happening at the IP layer like the technologies they talk about MM1Q's [Sounds like] and (??), you have a fell for it.

But in layman's terms, it gets more crowded. You expect more — longer line. You might have like an express lane for let's say under twelve items. What happens when you walk in a store and there's all sorts of priorities. What needs to be done intellectually is there's not a real understanding of that.

What the expected interactions will be — if you're going to have a stable network, if you have all these different types of flows, then you could have possible instabilities that arise. I mean, it happens in routing.

As a technologist, there hasn't been much work to stand on this. I know a little (??) and that's something we have a network that has production on it. This issue is sort of important. I mean, we might actually not look at it until we have a problem.

LARRY: Let's get back [Sounds like] to the provision of QOS Over [Sounds like] — network transience?

PANELIST: No. What I'm saying is you're going to be dynamically provisioning. You're going to have like an impacting problem. All these different people are going to want different priorities.

So, at a certain point, you have a certain number of small parties — even with the smaller parties, the possible commonatorus [Sounds like] is very large to model. So you're not going to really know what that dynamics or network [Sounds like] are going to do until like maybe you start having these instabilities start occurring.

So if you lock in a certain quality of service regime, you might find out later you have unexpected interactions that you didn't expect.

LARRY: On the QS models, one of the things — there's policing, marking, shaping — things like that. And one of the things that policing and mission control — well, they're different. You're just simply saying no to new clients. Say nope, you can't place your phone call with this quality of service right now. Maybe a lower quality...

PANELIST: Co-admit and you police correctly all the piece are there, but it's not really — there isn't...

LARRY: Okay. But I'm hearing two classes of problems here. One is that a class of problems about innovation, about new applications given structures we deploy for dealing with QOS. And the other was a class of problems which Jamie was bringing up about ways in which this might interact with monopoly problems we imagine in content markets or other problems about facilitating discrimination.

Not technically — not really — can't say technically here — really about this innovation problem. So there's a layering on another set of issues onto the set of questions we've been talking about — which is the set of issues about the monopoly...

PANELIST: ...a third part of that which is the notion of lumping e2e. Which is — we're growing the network with new technologies and we're going to — we just talked about firewalls. And that's sort of a negative way of looking at it, but deployment of new technology is an enabler, but it's not going to be deployed all at once everywhere.

So we're going to have certain kinds of new e2e services which will avail only the small populations.

LARRY: Yeah. But certain of these services interfere with other people's ability to deploy new services. That's the set of innovation questions we're talking about. But it's distinct. That's all I'm trying to make clear.

It's distinct from the issues that I think Jamie was raising here about what other problems this might create. Michael, you...?

MICHAEL: Two good points. One is the whole QOS [Sounds like] thing came up as a response to trying to provide quality of transport services on a network that's increasingly congested. So it (??) the trap [Sounds like] to go through.

We have taken a relatively simple network and we have layered on the simplest order within one network a level of complexity that we've never seen in this environment that's totally unintentional.

Okay. B - we're now talking about networks even in today's world, if we look at [Off mic] behavior on hearing, take QOS mechanisms, layer it on top of that and then have five different carriers involved, I have no idea how that will work. It has largely been turned from an engineering solution to a problem that I think is a transient problem that we need to view [Sounds like] the network properly to an economic leader for certain characters.

This is your point that really concerns me. I don't believe this is over [Off mic]. This is just decent traffic engineering that

everyone says can't be done. And I would challenge us to say let's just do piece of traffic engineering and build the networks the right way and be mindful of where the bottle necks move over time. And try to keep ahead of it.

But we're layering the level of complexity and cost and equipment and nodes and software. And I will guarantee it will blow up in our face. We have a much simpler set of options in front of us that is simpler, provides e2e and people call it over provisioning. I don't understand the content.

[UNRELATED DISCUSSION]

LARRY: To you, a centric view.

PANELIST: Why?

LARRY: Why?

PANELIST: I don't have that capacity to make use of most places in the world.

PANELIST: But that's merely an investment question.

PANELIST: No, it's not an investment question. It's a pricing question on the basis of the underlying asset holder [Sounds like].

LARRY: But then you say blow up in your face. That sounds like a technical term. Right? But on the other hand, you said the thing that really concerns you is the points that Jamie was raising about the ways in which certain providers can leverage power here.

PANELIST: What I'm suggesting is that the near term effects of this will be market control. But if the technologies get implemented, it is inevitable that they will lead to a set of complexities that will be unanticipated. And will really create problems that we'll have to go back and fix. And the question is is there a way to avoid going into that (??)

PANELIST: Who's going to be the we?

PANELIST: Who is we?

PANELIST: Yeah.

PANELIST: Engineering people, applications people. All the different consumers. One of the problems is we're now saying we're going to dump on the application developers that need to be mindful of not this problem, but this problem. Well, there goes innovation. All right?

And then we're going to say when the complexities arise and the unintended consequences arise, then everyone from Microsoft to the guy who works in development to all the

management people — that everyone — the ITF [Sounds like] are going to have to go back and figure out how do we simplify the problem we've created for ourselves?

LARRY: Okay. I want to bring another side on the table here. We know people in this room who are skeptical of Jamie's claim about this being an opportunity to leverage power in a way we should be concerned about. Doug?

DOUG: The problem — an active quality of service effort going on in the Internet 2 project. The reason for that is that we do not believe within the higher education environment that (??) that we can't prevail on everybody in a synchronous way doing [Off mic] to provide adequate capacity everywhere in the network. So we need to have a way of coping with the fact that some portions of the network at any given moment are not going to have the capacity to run all of the applications that we would like to run.

And therefore, we have to have some way of allocating for that. we also have a concern, and this goes back to our discussion with our corporate partners, that unless there is a way to provide e2e across multiple economical network systems, that it isn't possible for a network provider to offer high quality service to its customers at the end because they have no way to carry that service on through the network to the other end.

If that can't be done, then it's very difficult for them to argue from the capital to build the network that is required. So we view that, in the larger commercial context, that this (??) of thought and services and efforts is necessary in order to create the conditions which stimulate network providers with adequate capital into the development of the networks.

In today's network environment, it's very much like a slum landlord game where fixing up your house because the one around you aren't getting fixed up doesn't add any value to your house. And so you can't get more money to it — for it.

So we see these issues of quality of service — while they certainly introduce complexities about market power and so forth — as is always the case when you're looking for large investments and so on — we see them as very important issues for the naturation [Sounds like] of new applications of the network space.

LARRY: All right. Harold.

HAROLD: There is, I think, an underlying (??) about regulation at its being discussed here. And [Off mic] interact with technology that I'd like to bring out. Which is number one — the subtlety or grossness of regulation. You can have a regulatory process for

regulation due to the (??) and actions of large (??) players either through the view of merger processes [Off mic].

Or the regulations that take into account the size or take account of functionality. And to give it a precision and a minimalist approach that can handle in terms on one hand of market power, but at the same time, not getting in the way of [Off mic] particular technology.

So I would really urge people not think of regulation which is a very strict on/off kind of thing. The other thing is that (??) and Carl's point which I think is very right about lumping e2e and its development is again a probable regulatory problem, if I will, to just encourage this sort of expansion and this sort of [Off mic]. That it allows a certain amount of innovation to take place out here and also prevents any kind of manipulation that's going on in here.

LARRY: Okay. Now, I want to close this down, but I know I can't close it down without asking Mark to say something.

MARK: I want to make it clear — and this is probably like a moral and ethical thing — when you QOS controls and differentiated services into the network — it's a couple of processes. [Off mic] then you decide which class you want to put it in. Basically, it's three outcomes if you've got differential services.

You could say you could put packet into a more favorable class, you could not do anything with that packet and [Off mic], or you could look at that packet and duplicitously in an unfavorable class.

One of the things that we're seeing — and I hate to mention cable operator networks on the network and stuff — is that the way some of these operators are allocating their capital, it's for their own content, their own services.

Where they make packet classification the main — say, this came from not one of my places, so I'll give it this favorable access. The rules are important, but there's also the ethics that go along with how you use the rules.

LARRY: Okay. But there's an agreement about the description of what's happening. But what's interesting is radical disagreement about what flows from that. So, Debra, you...?

DEBRA: I had a question for the gentleman down here. You talked about how you can regulate [Off mic]. How does a regulator do that without being accused of taking the first step of regulating the internet? Which is an issue that we certainly face. I'm with the FCC. [Off mic]

PANELIST: Well, we've already started. You won't be taking the first step. We've already started doing that. I mean I have.

LARRY: Okay. The second issue that Jamie brings up is going to be at the core of the next panel. So I just want to switch to the next panel. And we're going to race to the end of the morning so we can get to lunch.