

## **The Policy Implications of End to End**

December 1, 2000

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

Panel 2: Hans Kruse, Bill Yurcik, Gary Larson, Molly Van Houweling  
and Michael Kleeman.

LARRY: These issues are going to survive the rest of the day. We're going to shift to a different set of issues raised by the next panel. So, thank you this panel. We have the next panel — is Hans Kruse, Bill Yurcik, Gary Larson, Molly Van Houweling [phonetic] and Michael Kleeman.

[UNRELATED DISCUSSION]

LARRY: So we're going to go through this panel. The idea was to keep these panels from between thirty-five and forty minutes. And then we're going to take a break. But we've got to get through the four panels this morning or else we're in real trouble.

So here's the second panel. It's raising a different set of issues.

[UNRELATED DISCUSSION]

LARRY: We can't quite put the fires out. We want light, cool the heat in the room. We're going to raise a different set of issues in order to raise that. This is the network on panel security.

A number of technological functions which are being deployed at different places in the network. And the question raised by this panel is where it makes sense to be raising them. Let me just make sure that everybody knows who's on the panel here. So, why don't you just go down and say who you are and where you're from. Starting with Hans here.

HANS: I'm Hans Kruse from Ohio University.

BILL: Bill Yurcik from Illinois State University.

DOUG: I'm Doug Van Howling from Internet 2.

MICHAEL: Michael Kleeman from Merry Networks [phonetic] in UC Berkley.

LARRY: Okay. So the particular topics we want to talk about here are firewall technologies, IP security technologies and we can also talk about NAP technologies which are not directly security technologies, but they will raise some of the same questions.

But let me start with Hans here. Take one of those or two of those and give us an example of an implementation of these technologies that's both e2e friendly and then one that's not e2e friendly.

HANS: Well, I'm going to be hard pressed to find one that's e2e friendly. And there's some reasons for that. Let me maybe just take about thirty seconds and just lay...

[UNRELATED DISCUSSION]

HANS: That better? There are some security concepts. And security gets — depending on where you're coming from and what your idea serves, people tend to focus on one aspect of security. And it turns out that different aspects of security actually have different requirements. And most of them conflict. Which is what makes part of this so hard.

Most people, I think, we think about security on the internet, think about the security concept of confidentiality. You send something out. It doesn't get intercepted. It doesn't get murdered along the way. Nobody sniffs [phonetic] your passport or your credit card number, what have you. That's one aspect of security. And consumers particularly tend to think about.

If you're a corporate network manager, you don't care so much about that. You have some interest in that as well. But your corporate so-called security manager is going to worry about availability. You know, they're worried about denial service attacks. They're worried about unauthorized access to the network.

You get very different ideas about what you want to do in a security perspective when you look at those various requirements.

LARRY: Give us one of these examples, then.

HANS: So look at availability which is where the firewalls come from. A corporate network is, among other things, vulnerable to the side effect of the e2e — which is I consent [phonetic] IP packets with any content I want and an address if I know the address.

Now, I can send a lot of them. I can send them with very malicious, sick content. And certain operating systems react in a very bad, unpredictable and blue screen kind of ways when you send them certain types of packets.

So, as a corporate security manager, one of the things I'd get very worried about is I don't want every packet to get to its intended destination. So as a corporate security manager, I want to break the e2e model because in e2e, not only again the

good packets get there, the bad packets get there, too.

And so the security (??) to the firewall. And if you're a good security manager, your approach to e2e is it doesn't exist unless I say so. A package stops at the firewall unless I let it through explicitly. So good security implementation from that perspective breaks e2e except for those applications that you specifically said I want to let through my boundary. Yes?

QUESTION: Except that it's based on this naive idea that you can look at couple of (??) fields and say this is a good (??), this is a bad (??)

HANS: Nobody says it's perfect. What we're saying — what are some of the examples? You break e2e to some degree to at least filter out those things you can recognize as quote unquote bad. And in fact, what I just said about how these things are usually implemented is based just on that. The idea is if I don't recognize it as good, it ain't going through. So, of course, from the perspective of this conversation, that means that deploying a new application in that environment requires that someone with the — someone at the choke point has to be involved in allowing the application.

QUESTION: Alternatively, everybody now builds their applications to use the magic numbers that get through (??)

LARRY: If that choke point is at the end, then you haven't infringed on the e2e principle. I mean, the simple fact of the matter is that the idea was if it gets to me and I am the end user device and I don't want to open it, it's not a failure of intent at all. As long as no other end users...

HANS: Define end. I think it's an important point. There's a concept of administrative boundaries. And if the firewall — this filtering device — happens to be at the edge of some cooperating group of users as at the edge of a (??) of domain, you might argue that that's perfectly okay. You know, that means that this group of users has to have a common set of ideas as to whether or not something is good or bad.

Now, in a university context, I will tell you that our network administrators says Napster bad. Our students say Napster good. So even though it's one administrative domain, our students would argue we've broken e2e where our administrators said no, no, no. We're all one perfectly happy group of users here and that's what we do. And that's e2e.

LARRY: So the important idea here is depending on where we define the end, we either create problems or don't create problems.

HANS: What is yellow? Is that...

LARRY: Yellow means we ran out of green cards.

[UNRELATED DISCUSSION]

QUESTION: It seems on the e2e argument, the paper that discusses where it comes from and where it should go, should a particular function go into the communications system or should it go in the end system. And one of the principles it seems to me is that if you have to do it in the end system anyway, then doing it in the communications system is at best redundant.

So, if say the best example — I was going to use a corporate example. The best example is a university. But if I have a university firewall between the university and the internet, then if I'm sitting behind the firewall on the university campus and I'm depending on that firewall to provide security for my system, then I would consider myself an idiot. Right?

Because I'm not protecting myself from people behind the firewall who are going to be attacking my system.

HANS: You mean students. Other students.

[UNRELATED DISCUSSION]

LARRY: There's a very good point here. And I think...

QUESTION: I've got a firewall and a personal firewall running on my laptop. [no mic] firewall connected to the internet. And everybody else is behind the firewall.

HANS: Okay. You clarified — you brought another principle out of the paper which was if it needs to be done at the ends, it shouldn't be done in the center. And then David Clark raised his hand with a red card that's disappeared now.

CLARK: Well, I couldn't tell whether I was honestly confused.

HANS: Pretend to be confused.

CLARK: But using firewalls as an example to illustrate something else, there are interpretations of what a firewall is. And one of them is that it's a second best substitute to a poorly implemented operating system. And some security people would say that your operating system said they worked [phonetic] so that they could deflect that (??) Then we wouldn't need firewalls.

HANS: Is that Napster packet a bad packet?

CLARK: [no mic]...mandatory versus the discretionary access control. And especially access control is one that you as a user chose to put in place to protect something that you had. And a mandatory access control is one that a third party put in place, even if you didn't like it because he thought he had the right to.

A classic example being the system of classified information. Just because we wanted [no mic]. And another interpretation of a firewall is distributed implementation of actually undesirable policy by a third party who says he had the right to do so.

And he [no mic] Napster, not because you agree with him, but because he says I trump you. [no mic], I say Napster bad. So a firewall can either be the second best [no mic] manifestation of the fact that there are parties with interests that don't necessarily align who are puzzling over who's in charge.

LARRY: Okay. There's a very helpful distinction. What's the cost, Michael of a third party veto or a person who believes he has the right implementing the firewall in the context of the internet. What's the problem to this?

MICHAEL: It raises sort of the question of what are you really securing against? For the administrative university who blocks MP3 packets, it's probably purely an economic question of flooding the network with traffic that is quote, unquote in one person's mind undesirable.

That was not the intended consequence of the construction of the firewall when the vendors first delivered it. Neither were proxy servers originally or proxy re-direction originally set up to allow the government in Singapore to censor and edit virtually all the traffic going in and out of the country.

So I think one has to distinguish between security and functional control over traffic flow, if you will.

PANELIST: Again, I think Napster may be a bad example in that context. But look at your corporate or university administrator might argue that Napster's essentially a very sophisticated form of download service attack.

Most of the universities I talk to — Napster had a two-hour outage. And there were several postings on various administrative mailing lists. Most university's access pipes during that outage dropped from 90 percent utilization to 60 percent utilization during the two hours. And then it jumped right back up to 95 percent when it came back.

QUESTION: This is an economic question. But if someone — but if your ISP were to turn around and say here's a pipe that's an order of magnitude larger, and I'll maintain that kind of head room for you, would you care? Is it really a matter of economics or is it really a matter of security? I mean, to the question that was asked, what is lost through these forms of security?

If one just wants to eliminate rogue behavior that is technically undesirable because it causes damage as opposed to economic

pressures, one could argue that the real damage of these firewalls is very, very — a great complication in the deployment and the diffusion of innovation.

LARRY: How does that work? Describe it.

QUESTION: Well, you know, I'll take a very crude example. You know, AOL's got on its instant messaging a talk capability because it's basically voice over IP. And I've got a relatively crude firewall at one location and he has a more sophisticated one on the other side.

He can do it and mine, which does simple math, basically can't pass those packets very neatly. So there's a function that, because I desire to grieve [phonetic] security, I can't implement. And there's a thousand examples like that. Simply maintaining all the address clearances for new innovation — quick time. I couldn't get quick time through my firewall for four months. And so you see this lack of diffusion of innovation because of an overly tight control.

PANELIST: Okay. Lack of diffusion, but the other side of lack of diffusion is that innovators realize that the existence of these technologies out there will raise the barriers to them being able to deploy their new innovation.

PANELIST: No, they typically don't. What they do is they gristle out a complaint about the firewall. People constraining them. But it doesn't slow down their...

QUESTION: [no mic]

PANELIST: Right. And then really contaminate everything.

LARRY: Okay. So then another issue — let me just to clarify one issue...

PANELIST: [no mic]

LARRY: He actually owns the mics.

PANELIST: This plays into another economic thing. Real networks is one of the big corporations that hit this problem in a big way. And there were people who came to Real Networks who made firewalls and said if you pay us enough, we'll fix it.

LARRY: Nice. Okay. Now, I want to bring out a point. I'm sorry, you're...?

STEVE: Steve.

LARRY: That Steve was making. And that's exactly what Michael, I think you were just suggesting. So Steve, if the firewalls exist out there like this and you're an innovator and you realize they exist out there like that, what's your strategy?

STEVE: Either you ignore it [no mic] or you [no mic] implement it over

HTTP or something else that is known to get rid of firewalls. And the result is you get an implementation that's very artificial. It's more complex than what it would otherwise have to be.

LARRY: Okay. So you implement it over HTTP means...

STEVE: HTTP is the web coder [phonetic].

LARRY: No. We know that, but I mean — so there are ports that are standard — that used to be standardly assigned by the people who are just neutral assignors of ports. Which means particular uses...

STEVE: [no mic] this is FTP — file transfer. One that says its [no mic], one that says its Telenet, one that says its RTP.

LARRY: So the web port is eighty, right?

PANELIST: Right.

LARRY: So if you phase firewalls out there and you're an innovator, you know that nobody's going to be blocking — or people are not generally going to be blocking eighty because that's the killer application. So you channel your content through eighty. That's what you're saying.

So another cost of this deployment of technology is that applications begin to hide what they are. They're beginning to pretend they're something else so as...

PANELIST: [no mic]

LARRY: Masquerading.

PANELIST: In some sense, they're hiding what they are because the network [no mic] can tell what things are by looking at that. But it's a guess [phonetic].

BILL: Larry, I think I have something to contribute...

LARRY: Please, Bill.

BILL: One of the things that Steve Deering is bringing up is that it isn't really a static situation. There's counter-measures. As complexity rises, as things grow on the internet and innovation, to configure a firewall — now I'm going to [no mic] conferences where people are talking about automatic scripts and logic language that — people write firewalls, configurations. One person cannot understand — you can't go to an administrator and say well, what goes through my network? What's allowed through?

They can't understand their own configuration. So that's a little bit, I think — what's happened is to make it from the policy side, you've taken like a managerial network management function of

how to keep your individual host secure. And say well, it's easier that instead of doing a thousand hosts and looking at their operating systems, due to this one point — there's a choke point or maybe a performance bottle neck, I can implement my security policy at this one point and that'll make it easier to do it at that one point.

But now, you don't eliminate the complexity of a near [phonetic] because now people are doing things, innovating and just as Steve says, they're reacting to what you're doing.

LARRY: So there's an arms race.

PANELIST: There's a cost to e2e deathing [phonetic], that there's — I think there was an indication in the last panel. I think there's a cost to e2e because if you create these types of innovations, this type of rapid change, ultimately — somebody talked about you don't want to secure as a show point only and not secure the end points — that there is a cost here because we ultimately are going to have to deploy a lot of managerial capability at every end point.

And it hit me in the past panel. Eventually at every damn telephone and that is expensive. If you're a corporation with ten thousand telephones and you have to treat each one like you treat the Windows 200 box today, that's expensive.

QUESTION: [no mic] that argument that AT&T back in the late '80's...

LARRY: David, how do you know that?

DAVID: I can't tell you because I'm still under [no mic] with them. People at AT&T — they said we have to make it backwards compatible to the dumb telephone. It has to be as easy to use as the dumb telephones.

And if you told people that there was a significant sector of the population that would spend their evenings and weekends hacking their communications terminals to make it do things that the telephone company never conceived of, they would say what are you smoking? Go home and come back sober in the morning, you know.

PANELIST: It is true with all monopolies, including AT&T, which I'm on the advisory board of — or Microsoft which [no mic]. All power corrupts. And obsolete power corrupts [no mic].

LARRY: Okay. Firewall technology is one of these technologies that, depending on your perspective, might be interfering with e2e. There's a second technology that's particularly illustrative here. NAT technology — network address translator technology.

Now, some would say that the reason we developed this

technology which someone's going to explain — not me, in a second — was because we had a restricted number of addresses out there. The man who's going to end the restriction of addresses here. Doug, who's from Internet 2. Could you start by just giving us a sense of the purpose of the NAT technology and the problems that it creates from this?

DOUG: Well, there are other people here that who could do a better job than I. But I'll say it real quick and then if anybody wants to help, they can.

There are a number of circumstances in the internet today where it is either infeasible or difficult to assign a permanent address to an end device. Sometimes, it's because there's a shortage of addresses. Sometimes it's because of the way networks are designed for security purposes. Sometimes it's simple economic motivation.

And so what's done is you take the address that I thought it was going to a particular end point and change it to another address before it gets there. And then, of course, I have to do the same thing on reverse — on the way out.

We see that a lot now — especially in these home wireless networks — just to give one very particular example. Where you get one address which is what the provider provides coming in. And there's a sub-network with a bunch of other addresses used inside the home to connect all the devices together.

LARRY: Okay, so in that context, it's the providers providing a single IP address and many different devices want to connect to that. What's the cause — what's the problem that's created by this?

DOUG: The problem is that any application including security applications interestingly for this panel — that requires there to be transparency of addresses from e2e to operate stops working.

LARRY: Breaks.

DOUG: And among those things — one of the things that typically stops working are any applications where your end is trying to host a service of any type.

PANELIST: And equally importantly, the thing that breaks is [Unintelligible]. And that's a real nasty problem.

PANELIST: It also turns out that simple address translation of the IP addresses for many applications insufficient because IP addresses are carried around as data. And so real commercial maps know about specific applications in order to make them work. And so if you have an application they don't know about

that requires transmitting address data, they don't work through the maps [phonetic].

Again, you're in a place where it's very difficult to deploy a new application because you don't have to just change the ends, but you also have to change these devices.

LARRY: Okay, but then we're back to what David was describing before where to develop an application, you now have to go out and negotiate with a NAT technologist who are providing NAT technology.

LARRY: [no mic]

PANELIST: Part of the problem, Larry, is that people don't necessarily know they're [no mic] there.

PANELIST: Exactly.

PANELIST: Often with firewalls — a thing that is said about firewalls is people generally tend to know that the firewalls are put in. But NAT's can be anywhere.

Windows 2000 comes with a NAT filled in — that software built in. The airport controller on Macintosh's have NAT's built in. You just — when you're sitting there using — you don't know whether the NAT between you and the rest of the world is [no mic]. Possibly, you will know there is a firewall.

PANELIST: There's actually a NAT built into a number of cable back end systems. And it's — Long Island Cable Company implemented a NAT at its head end so that the addresses were translated on every packet going outside of [no mic] Cable Network. And that's a common means argued again for this economic reason that for some reason, the ARIN [phonetic] or somebody says it's too expensive.

We want you to — there's a policy decision that's been made at sort of the ICAM [phonetic] level. It's too expensive to hand out addresses. So they make you justify to them that you need the addresses. You have in some overly conservative interpretations, my customers have said well, if we can deploy in that, we can just get one.

LARRY: Right. So from a boring lawyer policy perspective, this is the structure of the problem, right? There are organizations who are trying to minimize the cost of their doing business. So a university wants to minimize the cost of copyright lawsuits by their recording [phonetic] industry which seems to be very high these days.

So they want to deploy a technology so that they can control copyright violations by their students and that technology is a

certain kind of firewall technology, let's say.

That behavior could impose costs on others outside of the university. And in economic speak, in externality is imposed on others. And we've been describing the externality. The externality is on other application developers who now must take into account these various technologies out there for facilitating firewalls and NAT technologies and IP...

PANELIST: But I would argue that there's a larger externality that we haven't discussed yet. Because these technologies tend to vulcanize the net in very important ways.

We all have for a long time subscribed to the notion that the effectiveness of the net — the usefulness of the net was up as the square or something like that. Some function of the number of organization/people who are attached effectively to the net.

The instant you vulcanize the net, you reduce the net's overall effectiveness. This doesn't have to do with sort of how you deal with a new application and so on, although that's part of the vulcanization you suffer. It has to do with even more difficult issues which is why most universities don't have firewalls.

Or if they do, they're extraordinarily minimal because the whole notion is we don't want to have the network between people on the campus and people other places somehow interfered with. That's a part of the value that a university tries to maximize.

But let me give you an example of how this vulcanization and particular instant hurt [phonetic]. University of Michigan developed a virtual automotive college. And we had a bunch of online resources. And the people at Ford Motor Company wanted to consume those online resources.

In fact, they wanted people at the desktops in the Ford Motor Company to be able to consume those resources as part of their training. But, of course, it wouldn't work because the server we had on the campus to provide those services was not within the Ford Motor Company firewall.

And since it was not within the Ford Motor Company firewall, those educational services were simply not available to the employees of Ford Motor Company anywhere they were in the world.

What we ultimately had to do to get around this is Ford Motor Company had to hire one of its suppliers which was within the firewall to mirror this set of content. But, of course, once that happened, the faculty members on the campus would have been the people who would consult with these students about what they were learning had to go through the same kind of

mirroring process.

Which of course, required a hole to be punched through the firewall. Which ultimately a little while later, allowed somebody who had nefarious notions about what could be done inside Ford from the university to use that hole to cause a problem for Ford.

Now, the reason I'm using this particular example is to say that — is to point out that get in the way of e2e that give you some false sense of security every time you try to open the network to more effectively use the broad network — you discover that you're compromising the very solution you were originally intending to use.

And as a result of that compromise, you wind up having to push the security need all the way out to the end where it should have been in the first place in my opinion. So this issue of social value is not just about companies making money. It's about the value the society receives from the network and the investments we make in trying to use the network effectively.

LARRY: Okay. Now the trade off then that's implicitary is between — clearly there's going to be legitimate reasons for companies or universities to control what goes on in their network. And clearly there are harms that are caused by some of these technologies being developed.

The trade off that needs to be made then is a trade off between these benefits and these harms. Who could possibly be making this trade off for us? I mean, Carl Auerbach's here, but we'll hold him off for (??) later.

PANELIST: I have another example with ISP's. So ISP's at some points, they have to have technical control of their network. So you look at how their network's being managed or traffic flow and things. They say well, listen, we need to filter out things within the network.

They might filter out routing, advertisements and things like that. Or they might actually filter out services. And it just so happens that these services might be competing services. They might actually be services that say, okay, you're going to be behind a network (??) translator in our network and you can't run a server.

Okay. But by the way, we have this server. We'll sell it to you as a service and we'll make money. So it's sort of — you say who's making the trade off? You're giving a company a power to make a decision based on (??), but also enable their economic...

LARRY: It's not just a reduced cost. It's also increased profits that you

might be deploying these technologies...

PANELIST: It's just that they have a technical basis that say these things. But you sort of are giving them the keys to...

LARRY: Okay. This has been great.

PANELIST: Arguably, there's someone here who's ultimately — other than — you know, the providers are in there to make money.

LARRY: Yes.

PANELIST: But ultimately, there's somebody here who's trying to draw the benefit of the network which is the end consumer. Now, it may be a corporation through — providing service to their employees or maybe a home user. But ultimately, that end user ought to have the ability — and I don't think they always do right now — but ultimately, that's the person — corporate entity or otherwise, who ought to make that determination.

If I'm a corporate user, it ought to be my determination. Do I deploy a firewall because of cost of ownership and cost of management? Do I push things out e2e. Ultimately, the people that gain the benefit from the network ought to be the ones making that determination.

PANELIST: And let's remember that there is one other user of the network that has considerable standing here, and that's the nation state [phonetic] in which the network is being implemented.

Because they are trusted in some fashion or another by their citizenry to somehow maintain that nation state. And this is where all the wire-tapping and all the rest of the national security stuff enters the picture. And so, you somehow have to figure out — you were talking about balances, Larry. You somehow have to figure out how you achieve these conflicting objectives.

Now, the interesting thing about all of this is there are a whole set of more distributed, more e2e technologies that one can use to accomplish security objectives which require substantial additional thinking about how the core network is actually operated and implemented.

And what we have now — and we've talked about this in a lot of the literature for this meeting — is we have a bunch of people putting solutions into various parts of the network to compensate for the fact that we haven't really thought through and implemented the right kind of stuff at the core network level.

LARRY: That's right.

PANELIST: That doesn't mean making the network smarter at the core. It

means thinking about what services need to be deployed in what fashion around the...

LARRY: I want to end with one clarification, though. I mean, so Doug, you've described this policy decision that needs to be made, isn't being made. Hans, you suggested that it could be made by the individual users each on their own.

Now, if it is an externality in the sense that we're describing — externality that's (??) above the network as a whole as the way Doug describes it. Or particular application developers as Michael was talking a little bit about it.

Michael, is this the sort of thing that individuals on their own could solve? Is that a problem that's through the invisible hand self-solving or is there still a policy judgment that has to be made.

MICHAEL: Yeah, I think there's clearly a policy judgment that needs to be made. And it's also a technical complicated — there's also a policy judgment in that you have to protect the society against the rogue operators. All right? And I think that that's a real issue that gets certainly much more amplified in the case of a network like the internet.

I mean, we could have had a denial of service attack when people could flood a switchboard with telephone calls which is comparable. But there was an economic cost to it, it was highly traceable, didn't happen that often. And it was also a lot more difficult to do.

We have far more damaging things that can occur. And so there is a reason for some of this. The other complication I'd suggest and there are people here that are a little closer to this than I am — is we almost talk about this environment as if it's a quasi-stable environment as opposed to one of rampant experimentation. And significantly well designed chaos on top of a relatively stable core base. All right?

And until there some of the experimentation sorts itself out a little bit, that's why we can't have VOIP [phonetic] inter operating yet. We're still experimenting. And why we still need the core telephone network to be able to make phone calls. I'm serious.

It's like the old joke. I used to be a park and you used to call up people to get their e-mail addresses. Right? Because directory services still don't work. We're still experimenting.

And until some of the core intra-structures get stabilized, we can't hand over the tools to the end points to really manage this I think in a reasonable way. So I think it's both a technical and a policy question.

QUESTION: One of the unspoken fallacies in this discussion has been that the end points are all kind of equi-potential. And in fact there's all kinds of different end points coming along that are going to sit potentially behind at NAT, but you won't be able to differentiate.

LARRY: Okay. So we're going to pretend to continue these questions to the next panel because we're going to take a break right now. And the break is going to go for exactly 7 ½ minutes.

We're going to start with quality of service which raises these same questions again. Yes.

QUESTION: [no mic]

LARRY: I've been intentionally ignoring you. No. I apologize. Is there something that you'd like to say right now?

QUESTION: I think it's very important for an end policy to realize NAT did not come about as a security measure.

LARRY: No, of course.

QUESTION: Because globally (??) interests could not be assigned or the operator charged you for them. Which is in a way specious. Okay? Policy to get e2e working with regard to NAT is to not charge for globally route-like [phonetic] interests and get more out there.

PANELIST: And drive B-6 out [phonetic] ultimately. We have.

QUESTION: [no mic]

LARRY: Okay. So, we'll take now a seven-minute break.