

Who Runs the Internet: Internet Consolidation and Control

Chatham House, December 10 2019

(Corrected for coherence)

Emily Taylor

Ladies and gentlemen, welcome to Chatham House, and welcome to this session on Who Runs the Internet, which is about Internet consolidation and control. We've got a great panel here for you this evening. But first of all, I'll do the housekeeping. So this event is on the record. It's not under the Chatham House Rule despite being in Chatham House. It will be live streamed. It is being live streamed. And there's also, very important, a drinks reception afterwards, when this finishes, at 7:15 sharp. This is all part of a collaboration between the Internet Society and Chatham House on these issues of consolidation and control. And, as part of that, we have a special issue of the Journal of Cyber Policy. We've got some copies of the Journal, not the special issue because it's to be done in the Spring. But, what we have got this evening is a sort of preview of that special issue, and some of the topics that will come through. We've got two of our authors here, Jennifer Cobbe and Jesse Sowell, and also Andrew Sullivan, who's the President and CEO of the Internet Society. Before we get started with the panelists, I was asked to just set the scene a little bit, not just on the special issue, but what we're talking about. We're really talking about the bits of the Internet you can't see. I think we're probably all familiar now with the concept that, at the application layer, the platforms, the services, whether it's search or social media, that there are a few winners who've come through, whether it's Facebook, Google, or Alibaba, Tencent, and so on. Here, we're actually digging deeper into the so called Internet layers, what really makes the Internet the Internet. When I was getting started in the Internet 20 years ago, the exciting elements of it were that it was based on these open, interoperable, very simple protocols which were open, in that they were non-proprietary, as well as available to everybody. Nobody really owned the Internet. Nobody really controlled it. Nobody could be said to be running it. It was an environment where you had a survivable network, which was so useful during times of conflict, or times of crisis, but it was also an environment where there would be permissionless innovation. When we look back over the last 20 years, whether we agree or disagree about whether we're going in the right direction, I think we could all reflect that there's been incredible innovation, that our lives, we were talking about, just before, all of the skills we don't have to have any more, like map reading, memorizing things, memorizing phone numbers, things that I was actually quite proud that I could do, like I used to remember everyone's phone number, and I could read maps really well. Those are skills now that we don't really need as people, we have the Internet instead. So, when we look beneath that application layer, whatever we choose to look at, whether it's Internet exchange points, which you'll hear about from Jesse, or whether it is the those complex supply chains of online services, the same

names keep coming up again and again. They're often the same names that we're familiar with from the application layer. But what does this mean? And, what does this mean for the Internet? We've got a wonderful panel here. I'm not going to do big bios for everybody, but just a quick introduction, Andrew Sullivan, who's the President and CEO of the Internet Society. Jennifer Cobbe, who's a research associate at the Department of Computer Science and Technology at the University of Cambridge, and Jessie Soule, who's an assistant professor at the Department of International Affairs at the Bush School of Government and Public Service, Texas A&M University. It's quite something for a business card isn't it, quite the title, but welcome to Chatham House, all of you, and thank you for joining this conversation. We'll also throw it open to the floor, when we've gone through the initial remarks.

Emily Taylor

So, Andrew, the Internet Society's been a leader in defining, really, what the characteristics, what you call the invariables of the Internet, are. We are talking here about consolidation and control. When you start to chip away at those invariables, when you start to have big companies really owning a lot of the infrastructure, or running a lot of it, is it really still the Internet? Or is it just another network?

Andrew Sullivan

It's a good question. There are a couple of dimensions to this that, I think, are worth thinking about. One is that, in some cases, we've seen this before. We've seen periods of consolidation around particular technologies, and then the people who consolidated found themselves in trouble, and they haven't held on. I remember many years ago, for instance, everybody was panicked that America Online, AOL, was going to take over things, and was going to ruin the Internet. Instead, of course, they faded away, and we got the contemporary Internet. I remember, when I was first coming online, everybody thought that Microsoft had basically sewed up the world, and that there was no way that you were ever going to break that monopoly. And yet, they don't look so mighty today. I mean, they're still strong, but nothing like everybody was concerned about in 1998, say. So, I think that there's a dynamic there that it's very easy to fall into the prey, the mistake, of thinking that your current conditions are exactly what's always been true, or what's always going to be true. Nevertheless, it's the case that certain kinds of consolidation have the effect of stanching the ability of other people to enter. I think that's one of the real questions that we have to ask ourselves, are we creating an environment in which it's impossible for somebody to come along and make a better mousetrap? I think the answer to that is, is wide open. It's quite interesting, because there are two different kinds of pieces to this. One is a development, that we've seen over and over again, which really replaces open protocols with proprietary protocol interfaces. If you think of open protocols are the ones that are developed by the Internet Engineering Task Force, anybody can build on them, anybody can use them, and so on. A lot of online services today are, instead, developed over so called API's, application programming interfaces, that are proprietary things that are basically web services, and they're exposed that way. What that means is, if you're the company who owns that API, you can change it at will. If somebody comes along and starts building something that's better than your own stuff on top of your API, you just change the rules. And now, it's all over for them. We've seen this actually happen more than once. It's, by no means, a stretch of the imagination, because it's actually happened. That's one kind of development that I worry about, and I don't know what to think about it. The second issue has to do with the use of formerly open protocols like SMTP, or DNS, or things like that, and yet, to see a lot of

consolidation around that. When I first started using the Internet, essentially, everybody ran their own mail server.

Emily Taylor

Usually badly.

Andrew Sullivan

Generally badly. It turns out that running my mail server well is complicated, running it badly is really, really easy, and running it so that it doesn't turn in to this fountain of spam for everyone else on the Internet, is quite difficult. We saw people gradually offering that service over and over again. What's interesting is, if you look now at the mail systems, there's really between 5 and 10 companies that run them all. Everything else is completely in the noise. So, that's a formerly open protocol, anybody could run their own. And yet, there are these giants, who operate in that environment, who really can make or break the market. That is a case that seems really paradoxical, because you would think that there'd be a natural point at which the consolidation would stop. Instead, what we see, actually, is that there's an acceleration. I think the economics of that are quite interesting. But again, it's hard to know whether that is just that we haven't reached the equilibrium yet, or whether it's something that betrays a real fundamental problem.

Emily Taylor

In the Internet Society's Global Internet Report, which was published earlier this year, you talk about this sort of, we've got to a stage where it's rock star economics, where only the largest and most famous can profit. When we think back to the early days of the Internet, where it is this environment where anyone could make their mark, their fortune, that really is a change, isn't it? It's not really going to be possible for two kids to start up the new big thing in their garage anymore, is it?

Andrew Sullivan

Well, I think we feel that way, anyway. Whether that's true, I think remains an open question. I would point out that some of the things that look like rock star economics, are also opportunities for the people in the garages. From one point of view, for instance, AWS, Amazon Web Services, or any of the other large cloud ones, but I'm going to pick on AWS because they're the biggest concentration here in the UK, they look like an overwhelming dominator of the field. Yet they now deliver, and you can rent by the hour from them, facilities that, 10 years ago, nobody in the world could afford to build. That means that if you're starting up, instead of having to build your own data cloud, and your own database systems, and have all of the storage and everything, which you couldn't have afforded to do, you can get started very, very cheaply, and then build up. Now, there's a problem in this market, because it turns out that, somewhere in the middle of that, it's very easy to go bankrupt, because it costs you more to expand, and there are problems, and I think we'll hear a little more about understanding that supply chain. I'm not suggesting that it's all puppy dogs and rainbows. This is one of those things, what is the Tom Waits expression, the large print giveth and the small print taketh away? There's this idea that there's actually some benefit and cost here. We need not to concentrate only on the cost. That isn't to say we should ignore the costs, I think we should pay attention to them. But, we also should pay some attention to the benefits that certain kinds of consolidation can bring us.

Emily Taylor

If we have an Internet that is now really in the private hands of a few companies, what's the role of the Internet Society in that environment? I think of the Internet Society, I might be wrong, as very much a part and parcel of those original visions, of that very open network, that distributed permissionless Innovation network. As we look forward, and particularly to new technologies, the ones that you've been talking about, what will be the role for the Internet Society?

Andrew Sullivan

If we head down a path where it really does consolidate down to a few companies, then you would have to ask whether you had really an Internet or whether it's really a, I don't know what you would call it, an oligopolnet, or something like that. I am not too sure that we're there yet. It seems to me that we're, we're at the top of the hill, and we're not exactly sure which way we're going to fall down it. The other possibility is that the open protocols, on which all of these large companies also depend, remember, they can't turn them off because they need them. So, those protocols remain there, and they remain there to be exploited. The question is, will there be enough of a critical mass to pull away from this? Some of this is a social policy issue, some of it is economic policy, some of it is industrial policy, some of it is decisions that we each of us have to make. We maybe have the potential that we could had in a different direction, which involves using some of the innovations that some of these large players have provided, and yet using them to the advantage of the people who are trying to build things on the Internet.

Emily Taylor

And so is the Internet Society's role just to observe this, or should you be raging against the dying of the light?

Andrew Sullivan

I'm not sure that I am too comfortable with the dichotomy, because I think maybe the answer is both. There is a really serious responsibility to have some clear-headed understanding of what these issues really mean, some measurement of what's actually going on, because that's very, very difficult. It's important to understand how difficult it is, and why. At the same time, it's necessary to point out when certain kinds of concentrations are unhealthy for the network, what that does to the ability of people to use it. In a lot of cases, this is not merely a question of people's politics or values, they are also technical consequences. There are certain kinds of innovations that tend to only to happen when you've got a pretty healthy disagreement about who's going to win. When you get strong monopoly or even just a few players, or when you get regulated monopoly, which is a tradition that we had, you actually see that the the innovation tends not to be quite as great. That's something that we need to pay a lot of attention to. On the other hand, one thing that I've noticed is that some of the consolidating companies have started to act defensively against the fear that they're going to be regulated, or that they're going to be broken up. We've seen this in other industries previously. Those who have listened to me before will know that one of my favorite analogies is with the automotive industry. One of the things that's quite interesting is to watch car companies as the automotive industry consolidated. They started to get into a defensive crouch against regulation against them. That meant that a number of innovations that were potential, or possibilities, in the automotive industry much earlier on, only

happened later due to regulation. What's weird is, if we'd somehow managed to convince them not to worry about regulation, they might have shifted earlier, and yet we ended up in a paradoxical situation.

Emily Taylor

Thank you. So before moving on to Jennifer, I need to ask you about the sale of dot org. For those of you who don't know, the the Internet Society owns the domain name registry, for dot org, or had a financial interest in, and has recently sold it. My question to you, Andrew, is were you surprised, personally, by the ferocity of the reaction from some aspects of the community?

Andrew Sullivan

First of all, I should say that this is a pending transaction. It hasn't actually completed yet. I was a little surprised at some of the reaction. I expected that some people were going to be unhappy. I expected that there were some people who would think that this was the wrong thing to do. I expected some people to not to care about it. And I expected some people to be supporters. The people who are opposed to it have been very, very vocal, and I think there's a significant number of them, and I understand that. I do think that we did not perhaps explain ourselves as well as we could have, and that's on me. I was a little bit taken aback by some of the remarks, but I did expect that some people would oppose it.

Emily Taylor

Thank you. Jennifer Cobbe from the University of Cambridge. You have been looking at the complexity of the supply chains in the online services. You think you're working with an application, but in fact, there's a whole load of stuff going on underneath the surface. Can you just talk to us a little bit about what you discovered, and particularly whether you think that the existing regulations we have are sufficient to cast light upon that?

Jennifer Cobbe

Sure. We, as you say, looked at the supply chains of services that really connect together. If you're using an application layer service, what's actually happening, in many cases, is it's relying on a large number of infrastructure level services to provide whatever functionality it is that you're using. And, of course, at the infrastructure level, some infrastructural services can themselves be relying on other infrastructural services. So, you end up with a supply chain. It's really a complex system of systems, and the problem is that it is exceptionally complex. That complexity leads to a couple of different problems. The first and foremost being that it's exceptionally difficult to understand what's happening within that. There's real lack of transparency around this kind of stuff. The second major problem is that you run the risk of things like emergent property, and emergent behaviors, arising in that kind of complex environment, which means that you have a supply chain that is unpredictable, and potentially unreliable, with significant downstream effects. Unless you can understand how that's functioning, which is difficult to do because of the opacity, it's really very difficult for people to understand what those effects might be. We tried a few different legal and technical methods for investigating this complexity, to try to figure out what's actually going on. We are very interested in the data flows between these different services. If you can see where the data is going, you can see what services are being used. We tried some technical mechanisms,. What we found is that if you, as an end user, have a direct relationship in some way with the infrastructure service, whether it's taking data from you, or

returning data to you, then you can, to some degree, figure out where the data is going to. It's very much at the very top layer, you can't really see into second level and third level. There's almost a barrier at that point. But, if you don't have that kind of indirect relationship, it's exceptionally difficult really to figure out what's going on. There are very few mechanisms at a technical level to investigate that. From a legal point of view, there are very few frameworks that would require the kind of transparency that would allow you to look into this. The one that we tried using, partly because we're interested in data flows, was GDPR. GDPR requires that all processing of personal data would be fair and transparent which, in theory, would allow you to figure out where your data's going to. Companies, in complying with their obligations under GDPR, will often put information in their privacy policies. If you submit a Subject Access Request, you will sometimes get information back from them. What we found was that, primarily, what they tell you is not actually useful information in any meaningful sense, They tend to tell you that they use categories of third parties, rather than which actual third party. They might say, we transfer data to a range of (unintelligible) for analytics, or something like that, which makes it really very difficult to understand where it's actually going to. There isn't the granularity you might need. The flip side of that is if you have too much granularity, you might kind of be overwhelmed, for the end user trying to figure this stuff out. We found it was a real limitation in what we wanted to figure out. What we were able to find, from our investigation, was some evidence of consolidation. Particularly Google and Facebook, as services, were absolutely dominant in many areas. Amazon was a big player as well. We didn't have enough evidence to really figure out exactly what's going on, and we don't have the legal and technical mechanisms to really get into this. If we want to be talking about consolidation, and perhaps, what we should do about consolidation, or what kind of responses we need to consolidation, we need, I think, a much better understanding of that consolidation, not just at the top layer of infrastructure, but deeper into those supply chains, to figure out what's really going on there, and we probably will need some degree of legal reform for people to find that. GDPR isn't sufficient. The other legal frameworks that they have, themselves also aren't sufficient. We might need to try to align the incentives for these service providers to provide more information. A lot of this stuff may be commercially sensitive to some extent. A lot of the commercial pressures would be to not reveal where data flow is going. We need to try to shift those incentives in some way, to try to open that up a bit, whether that's through legal reform, or policy reform, or whatever else. Those were the things we were looking at.

Emily Taylor

Rolling back a bit, why is this important? Why does it matter to an Internet user? If they're just using a service, what would they care, really, about all of the different twists and turns in the supply chain? Why does this matter? It's obviously an interesting research project to see if you can find it out. But why does it matter?

For end users, not everybody is going to care where their data is going. We know that from the way that people engage with the application layer services, a lot of people don't really care what happens with their data. But I think there are a core of people who do genuinely care, and actually have a reason to understand where their data is going on, or perhaps are privacy conscious, or perhaps just interested in this kind of stuff. The real benefit isn't necessarily for end users themselves, it's for researchers, or policymakers, or regulators, to really understand what's going on here. If we're talking about systems

that are heavily consolidated, these systems can be quite fragile. If the infrastructure services provided by one company go down, a whole range of application layer services then could themselves go down. They could be quite critical things that we might rely on in society. More generally they could be quite trivial things. But that's a problem. Unless we, as researchers, or academics, or policymakers, or regulators, have a better way of understanding what's actually going on below that (unintelligible) surface, we can't possibly begin to come up with responses to address those kinds of things

Emily Taylor

It goes back to Andrew's point about the resilience of the network. An example, that you gave in your paper, is about how one expired certificate managed to take down O2's global network, and also a mobile provider in Japan. When you have that sort of very consolidated supply chain, you can have surprisingly big results. Of course, Andrew lived through the Dyn attacks, and all of the follow-on, the implications of that. As a policy maker, what questions should policymakers be asking about these supply chains at the moment?

Jennifer Cobbe

The big question really to be asking of service providers, whether at infrastructure level, or whether at application level, where is your data going to? What are you getting back from people? What services are you engaging with? Why are you engaging with those services? When we can ask those kind of questions, we can begin to map out data flows, begin to map out the consolidation, we can begin to map out what kind of services are perhaps more consolidated than others, because it's not necessarily the case that it's a uniform level of installation. Until we start asking those questions, and until we start really investigating this a bit more, we're kind of in the dark to large extent over this.

Emily Taylor

So, we need to respond to what the evidence is, what the data is showing us, in a way. Thank you very much. I'm going to move now to Jesse Sowell. You have been looking at Internet exchange points and consolidation from that perspective. There are people in the room, I can see, who know exactly what an Internet exchange point is, but there's probably one or two who don't. Could you just explain what they are? But also, talk to us about the impact of consolidation in that area of the infrastructure, whether it's a good or bad thing, speaking simplistically.

Jesse Sowell

Sure. Building a little bit on Andrew's framing, and this notion of the power of private networks, the Internet is a network of largely private networks, and somewhere, geographically, in many points around the world, you actually have to interconnect those networks. A lot of these folks are nominal competitors. They don't necessarily want to build their infrastructure into the facilities of their competitor, and vice versa. So, in many areas around the world, around the late 1990s, folks said, you know, we don't want to pay exorbitant prices to AT&T, and Cogent, and all the big telecom providers, to move our bits across the ocean, and move our bits around the world. We're network engineers, we can build our own interconnection facilities. We can build our own facilities to connect our networks, to keep our local traffic local. And, in many ways, that's the essence of what an Internet exchange is. It is there to ensure that the networks that are geographically close to each other, their traffic stays geographically proximate. They have evolved from just moving traffic back and forth amongst local actors to becoming

these platforms, supporting markets for a lot of the infrastructure services that animate the Internet. You have transit providers, you have transport, you have a lot of the mechanics of moving bits back and forth. You also have DDoS mitigation services, that are increasingly seeing these platforms as places to deploy, and places to be able to absorb a lot of DDoS attack traffic. In many ways these Internet exchanges are the baseboards for connecting the networks that make up the Internet together. And, geographically, in terms of where they're distributed, they are, when we think about this consolidation, this notion of concentration, it's often the case that you want to have one, or maybe two, of these exchanges in a given metro area. The whole point is to have economies of scale. The point is to have one-stop shopping for all of your interconnection needs. If you go to this market, you want to be able to find all the other networks in that region. You want to be able to find all the other folks, that you want to interconnect with, in one place. When we think about the larger discourse around consolidation, building on some of the discussions we had earlier -- again, I'm going to want to steal some of Andrew's points -- when we ask about the question of whether concentration is good or bad? Building on Jennifer's points, yes, we see a lot of negative consequences of what we call consolidation. One of the major problems in some of the discourse that we see, especially in the media, is that we're confounding this notion of consolidation, and the negative anti-competitive consequences of consolidation, with concentration and economies of scale. I would argue that in many ways, the concentration and economies of scale are necessary to get a lot of the benefits that Andrew pointed out. These small actors that by virtue of being able to get on AWS, playing on again, the AWS example, yes, we see some degree of consolidation, or concentration. Yes, we see that there is some degree of power by being in one facility, but it also gives these small and medium sized actors the global reach, that they wouldn't necessarily have, and they don't necessarily have the expertise, or funding, to invest in. When I look at this, I look back at the Internet exchanges and I see that their very governance mechanisms, the means by which they manage their concentrated platforms, they have intentionally developed specific sets of norms. These sets of norms are developed by nominal competitors, by folks like Google, Akamai., BT, AT&T. These folks have gotten together and said, we need this common platform to interconnect. It is beneficial for all of us to have this common platform and, as a consequence, this common market. They have identified effective norms, like non-compete and neutrality, that keeps the platform in check, that ensures that platform does not go off the rails, and that they have a voice in how that's managed. In essence, a diverse set of sub industries all have a voice in how that platform is run. When we think about that, as in going back to the story of whether concentration is good or bad, we see there is the possibility, for these very same actors that we are critiquing at the application layer, folks from those firms have been able to develop governance mechanisms for a concentrated platform that are effective, that have the kind of transparency and accountability that I hear in Jennifer's comments. And so, for me, a lot of this discussion, the consolidation debate, is an opportunity to say, Can we innovate in the regulatory space, not go to one end of the spectrum and just fine the heck out of folks for behaving badly, not go to the other end of the structural divestiture spectrum, where we say just break them up, break them up, break them up, but how can we take the insights from neutrality and non compete? How can we, not necessarily, cookie cutter, copy and paste them up into the application layer, but how can we use those insights and see, yes, these firms are willing to engage in these cooperative norm building efforts to manage a platform? How can we transfer those insights to the application layer?

Emily Taylor

It's a very interesting point that you make, that the governance of these Internet exchange points was really, quite thoughtfully, set out at an early stage. The risk was recognized that these are powerful nodes on the Internet, you've got competitors, all sharing traffic, you've got all sorts of things that, if you were doing a risk analysis, that would be raising red flags, but that they set up those norms that you speak about. Have you got any examples of the norms working in practice? You said, if it goes off the rails, if something bad happens, can you give us any examples of that? We all can recall things that looked like a fig leaf. We have a multistakeholder panel that will stop anything bad happening, but actually, it doesn't really mean anything. Do these norms really have teeth in practice?

Jesse Sowell

I would say, I will say, that yes, they do. There's two good examples of that. One is right here at home in London, the London Internet Exchange in their early days, a slight modification. Whenever you say that these norms that were really well thought out, they were really well thought out, but it was a process of experimentation. These folks in the late 1990s, they didn't just come along and say, neutrality, non compete, and mutuality, are canonical norms, and magic will happen after that. They had to do a lot of experimentation. In London folks got together and said, Why are we paying these exorbitant transit costs to send our traffic across the pond and back, just to go up the street in London? They came together, somebody in the audience could probably correct me, they had a little SPARC box that was essentially their core router that they plugged into. They said, we're all network engineers, we can put this together, we can actually save ourselves some money. We are still nominal competitors, we can save ourselves some money exchanging traffic amongst ourselves. That was initially a volunteer effort. That was something that they put together with just a bunch of engineers, who got together and said, Hey, we're going to build this, and we're going to save ourselves some money. As people started using it more, they saw that there was a demand for more professionally managed services They wanted stronger guarantees that this little box, and the larger set of switches the grew out of that, and then this multi-node platform that became the LINX platform, that now has platforms around the UK, they wanted stronger guarantees about that. But, in those early days, when they were transitioning from the volunteer model to a more firm based model, they said, it's not just a bunch of volunteers, we're going to hire staff engineers to manage, and monitor, these switches, and make sure of things, and make sure we have service level agreements in place that say that, yes, we want to have this uptime. There was a period of time where they thought about going commercial, where they thought about moving from this nonprofit, communal based, mutually-based model. That was a big inflection point for them. The community, the industry actors of play, it was a contentious debate, but they did say we want to maintain management of this. We do not want to see this taken over by large actors, that just have their single profit motive at play. We like to have this management role. So, there was this element of, what we would call in policy world, policy experimentation, norm experimentation. That's really that big inflection point for them.

Jesse Sowell

The other story that I really like to tell is the one of France-IX. If I remember my dates correctly, back in 2008, a lot of folks, some of whom are in the room, pointed me to look at France IX as an example. Paris was a really bad place to interconnect back in 2008. The market for interconnection facilities was

very fragmented. There were 8 to 10 very small places where you could interconnect, and most of them were managed rather poorly. A few folks from some of our consolidation objects in this discussion, Google and Akamai, two international firms, and then Jaguar and Neo Telecoms, two French firms, the four of them got together and essentially to the rest of the community, the rest of the European industry, Would you like to see a common platform in Paris? Would you like to see a mutually managed, a commonly managed, platform? Would you be willing to pay for it? They actually put together a survey, and sent this out to industry, and industry actors came back and said, Yes, we would, that would be really great. That'd be really nice. They went ahead, and they built that platform. They built France-IX, and France-IX is now one of the most successful Internet exchanges in Europe and, I would argue, also in the world. This is an example of where they took that deep experience from those experimental days. One of the co-founders, that I interviewed for my research, literally said, they took the articles of association of two of the biggest Internet exchanges in the world. They compared them, and said, I'm going to take the best pieces of both, I'm going to tweak it a little bit for the modern Internet, and we're going to create this model. I wouldn't say that these norms were made purely by design, they were made by experiment, but, in effect, these actors, these nominal competitors, came along and corrected a dysfunctional interconnection market in a particular region. I think that's a fantastic illustration of how you can build a platform that, under certain forms, under certain antitrust literature, would be considered a concentrated platform. They created a concentrated market that, yet, is still managed by the industry itself, by the broader set of industry actors, not just a single firm. When we look at this, the big difference between these platforms at the baseboards of the Internet, and the application layer, is, here, we have an illustration of really good governance norms for a concentrated platform, whereas, at the application layer, we say, okay, we have a single actor that's often a participant in the markets their platform is directly supporting. How can we how can we transfer those norms?

Emily Taylor

How confident are you, that you could actually apply similar norms to companies that are now very well established, extremely powerful, and, as you said yourself, they operate for their own profit motive? Do you think that those sort of touchy-feely IX type of norms could really be applied to those platforms?

Jesse Sowell

They're not necessarily touchy-feely. They're absolutely there because these actors, at that level, have to collaborate. Whenever I look at applying these kinds of norms, neutrality, non-compete, as a way to encourage these platforms to share additional information, this is where I'm going to drag Jennifer into this answer, because she's looked at more of the supply chain of the application layer than I have, what they have, by virtue of these individuals, with participants being experts, they know precisely what to look for. The challenge at the application layer is taking an even more diverse set of participants on these platforms, encouraging them to say, What kinds of questions do they need to ask? What kinds of supply chain questions do they need to ask to avoid the emergent behaviors that Jennifer is concerned with, and I'm, quite honestly, concerned with as well? From the regulatory side, what is the appropriate regulatory incentive to make these powerful actors introduce these accountability mechanisms, these transparency mechanisms, so, in my case, the participants can perceive precisely how their data is used. Looking at Amazon Marketplace, those third party marketplace participants, how is their information about the risks they're taking Introducing new products and services? How is Amazon using that? Is Amazon using that, just to say, Oh, look at them, they took a risk, they deployed an interesting

service, or interesting product, now we're just going to go and deploy it ourselves. Oh, if they can't keep up, too bad, so said. Being able to have that kind of information, that kind of transparency about precisely how that data is used, that, for me, is one of the million dollar questions. What comes out of the IX work is that you do see that these competitors can actually come to some degree of common accord. I don't think it's impossible, but I think it's an extraordinary challenge that lives between, again, this fine-based behavioral structure, and complete divestiture.

Emily Taylor

Andrew, you had a quick point on this?

Andrew Sullivan

Yeah. To double down a little bit on that.,there's an important point about this. It's not only a regulatory, or social, problem that we're talking about. This isn't merely a policy issue. The policy issues that Jennifer is talking about, are also technical problems. If you're designing that service, and you simply don't know that actually this thing is consolidated upstream, and you're dependent on a single point of failure, you are really in trouble. You're designing what you think is a reliable service, and instead, it turns out to be, three steps down, critically dependent on a single point of failure, that you thought you designed around. So actually, there are just good, solid, technical reasons to be informing people of these things as well.

Emily Taylor

Thank you very much. Any, anybody like to ask a question? Maria, at the back.

Maria Farrell

Just, on the single point of failure thing.

Emily Taylor

Can you introduce yourself?

Maria Farrell

Oh, sorry. I'm Maria Farrell. I'm a writer and speaker about tech, and the future. I met Tom Leighton of Akamai last year. He was saying that they, they are a big CDN, that they host 19 of the top 20 American banks. I was asking him, afterwards, I was saying, Are you not worried about that? His answer was no, because we're really good at what we do. I think they are, but I'm still a little worried. So, I was just wondering about hearing a bit more on that.

Emily Taylor

Okay, thank you. I'll take these two from Joyce and Stacie down the front, and then come back to the panel.

Stacie Hoffman

Hi, I'm Stacie Hoffman from Oxford information Labs. I work a lot in technical standards. So I know quite a bit about this area. What I see happening right now is an Internet that will be fundamentally different in five years. We've got 5G that's moving a lot of services and tools to the operator's level. The

infrastructure, we've got DNS over HTTPS, which is moving resolution to an application layer in many respects, or could do. Then we've got actors like China who are looking to centralize, or consolidate, things like data collection. Looking at all the different incentives for this, state actors and an industry, do we really need to just rethink our Internet layers in terms of what consolidation will mean for us in the very near future?

Joyce Hakmeh

Thank you very much for this very interesting, engaging panel. My name is Joyce Hakmeh, I work here at Chatham House with International Security department. I'm also with Emily, co-editor of the Journal of Cyber Policy. Also, thank you for talking about the challenges and opportunities about Internet consolidation in terms of the newcomers, and what that means, and also how they are supported by these big platforms. We've been hearing, increasingly, some politicians calling for the break up of these tech firms. My question is, is this the right approach to the issue of consolidation? If not, what are the risks that are associated with it? Thank you.

Emily Taylor

Thank you very much. Great questions. Are you worried about that example of Akamai looking after 18 of the 19 biggest banks, or whatever it was? Even if Akamai are not worried, should we be? When we think about the Internet in the next five years, whether it's 5G, DNS over HTTPS, or what China's doing with centralized data connection, do we need to rethink the Internet layers? Do we need to rethink those invariables, and describe them differently? And, is breakup, you know, hard to do? Is it the right response to the sort of consolidation that we're seeing? Who'd like to go first, you don't have to do all of them. Just pick one. But do pick one. Jennifer?

Jennifer Cobbe

In response to Maria's question, the question that these people seem to think they're completely infallible, in some way. We've seen this happen with Facebook, for example, over the years. People have said to them, you have a problem. In a couple years time, this is going to be a real issue. And they said, Well, we couldn't possibly have a problem, you being daft? No, we don't have a problem. A couple of years later, they go, Do you know, actually, we think we might have a problem here. We might actually have missed something. And then, at some point, it turns out, it's actually a really serious thing that nobody really knows how to deal with, because they were completely unwilling to address it, when people were telling them about this. I think it's probably the case, that a lot of these companies think they're really good at what they do. They're really good at security. They're really resilient. They're really reliable. But actually, in practice, that's maybe not necessarily the case. We have to not really take them at their word for this. We have to say, look, it's great that you think that, but I don't want to rely on your delusions of infallibility, and we going to have to try to find a different way to figure this out.

Emily Taylor

In fairness, a lot of them are very good at what they do, and probably better. This is part of what Jesse was talking about, and what you were talking about as well, Andrew. Whether it's on resilience, cyber security, or just a professionally run network, a lot of these companies that have become so successful, do so because they're good, and we like their services. But yes, point taken.

Jesse Sowell

There's one thing that one of my interview subjects, very early on in my research, highlighted to me is this saying that might scare some people, but it's a fantastic quote. In the early days of the Internet, as a network engineer, or as a peering manager, if you haven't fat fingered something, which means you make the equivalent of a typo on a routing announcement that says who you should talk to, and how, if you haven't knocked out 30 to 40% of the Internet, at some point or time in your career, you're not doing it right. But the corollary to that is, doing it once is learning, doing it twice as incompetence. The mentality, amongst a lot of the network operators that I've worked with, is that there is some degree of some of the conversation we had earlier, moving fast, breaking things. There's something to be said for that experimental model, that permissionless innovation, that ability to do those experiments, because there's just as much you can learn when you break something as there is from saying, oh, I've done it well. I agree that there there are plenty of people out there that are extraordinarily overconfident. There's also a counterbalance. There's lots and lots of really great operators, and engineers, out there who have said, we will inevitably break something. That's going to happen. But, you have to have the discipline to actually learn from that mistake. That's something that, I think, is invaluable. It's something that's very characteristic of an industry that has to manage this extraordinarily complex system.

Andrew Sullivan

This is an interesting point that I want to pick up on, in order to address this question of whether there's something fundamental, that we need to rethink. One of the things we see in some of those efforts, and not just in the efforts that you mentioned, but, in fact, in a number of other regulatory efforts, in the idea that, Oh well, we will magically break up these firms, and then everything will be fine, is a kind of magical thinking that the managerial approach of centralization is going to solve this problem for us. The solution that many people seem to be proposing to the problems that we're seeing on the Internet, that arise from consolidation and concentration, is, Oh, I know, we'll just replace the concentration. Instead of leaving it up to, say, Facebook's engineers, or AWS's, or TenCent, or whoever, we'll put the government in charge of this thing, and then it'll be fine. That's a bizarre design. Right? The clever thing, that we did with the Internet, was to build a reliable system out of unreliable parts. The reason we did that was because we depended on the idea that, yes, somebody was going to fat finger things, and they were going to take down 30% of the connections. The solution to that was just to expect it to happen. The previous architecture depended on perfection, and it was extraordinarily good. But when it failed, Boy Howdy, did it go bad. The solution to that was not to have that problem. The point that Jennifer has been making, in her research, is precisely that we are wandering into that problem by accident, because we're obscuring, from ourselves, the fact that we've actually got the single points of failure. Another way we could do it, for instance, is with a new IP system, that centralizes management so tightly that, in fact, you don't actually have the intelligence at the edge, you put all the intelligence back into the network. We tried that thing, and it fails hard. So, don't do it again. You would think that we would learn that. I almost feel like we're going to demonstrate, empirically, that old saw that those who do not learn from history are condemned to repeat it. I'm quite nervous about that, for that reason.

Emily Taylor

Thank you. Anyone else? At the front.

(Participant)

Hi, my name is (unintelligible) from (unintelligible) from Turkey, which page is still blocked. I would to ask about Tor. There's another network that you can reach with Tor. It's called the dark web, and also the blockchain, these things are trying to decentralize the Internet. But we are talking about more centralization, consolidation. What do you think about this subject? Because the thing is, we are not aware of the whole Tor network, which is very decentralized. Thanks.

Emily Taylor

Thank you very much.

(Participant)

Hi, my name is [Ian Grosso]. I work for Fujitsu. it's along the lines of your question, but in a slightly different way. Where you talked about collaboration rather than competition, that was only done for economic or technical reasons, and was done at quite a low level in the Internet. Trying to extrapolate that further up the stack is going to be a problem. We talked a lot about centralization, government. My question is, where centralized? Which government?

Emily Taylor

Okay, so, where centralized, which government? And also, is Tor, or the dark web, the blockchain, are they going to bring us back into a decentralized network? Who'd like to have a go at those?

Jesse Sowell

I can answer the question about the cooperation. It's an old term, but it wasn't necessarily pure cooperation, or pure competition. it's that notion of coopeitition. which I think you implied, I don't put words in your mouth, with regards to the economic and technical rationale and incentives for this. I do want to clarify that I'm not suggesting that we cookie cutter, just push these exact norms, and rules structures up the stack. Those particular norms, this notion of non-compete, that the platform should not directly compete with the participants on that platform, these notions of neutrality, that the platform should be neutral with respect to all of its participants, should not necessarily privilege any particular sub industry, or any particular subset of those, as principles for governing a concentrated platform, are useful learnings that we can take, and insights that we can apply. Do I think that it's easy to do? Do I think that there's a single hegemon that can do that, that can come in and enforce that? No. It's a transnational Internet. My students often say, Well, can't the government come in? Which one? It's just not feasible. I don't necessarily mean to come across as I have the answer. The lesson here is that we have two instances of concentrated platforms. At a conceptual level, they are similar. We have one that's governed effectively, and one that is suffering from bad governance. It's not the fault of concentration. It's not a consequence of economies of scale. It is a consequence of bad governance. So, can we take those insights and apply them? Can we say, How can we incent the application layer firms? How can we incent Google to give us more information about how our information is used for ads? Facebook, how our information is used for ads? Following Jennifer's argument, you are providing services that are actually quite critical. What incentives can we give them to say, you need to give us a little bit more information about your supply chain, and give us stronger guarantees? We saw that at the IX level too. We saw the demand from industry, and, in some cases, governments, say, we need stronger guarantees that these platforms will continue to function as we expect. We can no longer

follow this idea that, Oh, it's just going to magically create this collateral public good. No, we have to actually have those stronger guarantees. It's quite an open question about the precise mechanics, but I think we shouldn't ignore these lessons we can learn. Does that give a relatively decent answer to your question?

(Participant)

Yeah, to a point. There is a financial element that is missing here. (unintelligible)

Emily Taylor

Thank you. For those watching on the live stream, the point was that, as you go up further in the stack, the financial rewards are much higher. So, the incentives to collaborate and cooperate, and throw yourself open to that mutuality that we see down at the IX level, perhaps, is less applying. Jennifer, which government?

Jennifer Cobbe

It's a good question. The first thing I want to talk about is really the Tor blockchain question. We can look for technical solutions to this, if we want, but actually a large part of the problem is the business models of the companies here doing this. We rely on getting more data, growing in scale, operating at transactional level, on a (unintelligible) driven purely by profit. You can get them to introduce things like Tor, or using blockchain, if you want, but, as long as the drive towards more data, more bigger scale, more money is a fundamental component of this, that's not going to solve that problem. I don't think. Also, breaking up these companies isn't necessarily the best idea, because it kind of assumes that there is some degree of a market solution, to a market problem. I'm not necessarily sure that that's true. I think, actually, we need much more creative and inventive ways of dealing with this. I think you're right, there has been bad governance, to some degree. We haven't really stepped in. We've taken a very hands off approach to this kind of thing. We need to be much more interventionist, less reactive. That brings us back to the question of, which government? Which is an exceptionally difficult question to answer. There's no straightforward answer to that. What we have seen is that certain regulatory blocks can influence very significantly the way the technology is developed. Particularly, in Europe with GDPR, we've seen the emergence of privacy enhancing technologies, and this kind of thing. Because of the size of the European bloc, because of the size of the single market, the evolution of European data protection laws has had a significant impact on the evolution of data protection law in other jurisdictions, as well. It's not necessarily the case that we need to create a harmonized set of global rules across this, but if you have a bloc the size of, say, Europe, that is willing to take a lead on this, then that can have to knock on effects, and move us on. One other thing to come up is this idea of neutrality, that we should be having these infrastructure services be neutral. I'm not sure that's necessarily a realistic option. I don't think we can ever be neutral. I don't think it's a neutral act, to say (unintelligible) online, because (unintelligible) i don't want to give it DDoS protection. That's not (unintelligible) whenever people are learning online how to attack mosques in Christchurch. These aren't neutral things. These aren't neutral services. We know from research, that's been done by Langdon Winner, by Lawrence Lessig, people like that, there's no such thing as neutral infrastructure. It always constrains, and shapes, and permits, and allows behavior. We have to be aware of that. We can't just go, this is a neutral thing. It should be neutral. We have to be aware of those things, I think,

Emily Taylor

Andrew, I'm going to go out to another question. If I may. We'll take these three questions together.

(Participant)

If the Internet's going dark, will there be a role for the Chatham House rule on the Internet of the future? And in reference to Andrew's defense of the dot org, Why should I give a dot fig for the US National domain?

Andrew Sullivan

I'm sorry I didn't hear the last bit

(Participant)

(Unintelligible) like dot com, the US national domain, so I don't really care.

Emily Taylor

Thank you. And then we've got a gentleman there.

Michael Cooper

Hi, Michael Cooper, a member of Chatham House. I was intrigued by a couple points that Jesse made, particularly. If I look at Internet, generally, you could argue that it's been slowing down, and struggling with more and more things, over a long period. We've had an addressing problem for a couple of decades. TCP is known to be problematic, and fixes, and corrections to that, just can't get out. We can recognize problems, but we can't seem to fix them. Do we not need concentration, either in the form of government, or large companies, who've got very smart people, and the resources, to actually drive forward the changes that are probably needed for the Internet?

Emily Taylor

Thank you for that. And, just the gentleman behind you, if you'd pass the microphone back.

Participant

I was intrigued by the comment that Akamai host so many of the American banks. I would have assumed, however, that the American banks, like the British banks, have four or five cloud suppliers and they multiple source. We've got in one direction this drive for consolidation, economies of scale and so on, but any large corporate, if they are critically dependent on online services, are going for multiple independent routings, via multiple independent technologies, using multiple suppliers. The only glue is the fact that these things can communicate with each other via IP. Now, is there not, because of the sheer need of industry, and government, for critical national infrastructure and resilience, a force that's driving in the other direction, very healthily?

Emily Taylor

A range of questions there, one about dot org, really a sort of I don't care about.org, but also a very interesting point here about, where you have a single company that can do something, they can implement changes very quickly, that we haven't been able to do with say, ipv6 adoption, for example. And then, the fact that if you've got four or five suppliers, are you genuinely resilient? I suppose a lot

depends on whether those four or five suppliers don't just all terminate at one point, and it's, in fact, one provider. This is why we need to know about supply chains. Can I come to you first, Andrew, on those points, and then, we'll come back?

Andrew Sullivan

Sure. I didn't fully understand the question about.org, but since it's "I don't care", I guess that's okay.

Andrew Sullivan

The point about v6, and this point about multiple providers, is actually two poles, two ends, of a discussion, in a market that is consolidating in ways that might be dangerous. That's the central point. There's not a lot of evidence that single operators are any better at v6 adoption than anybody else is. To pick another example, DNSSEC, which is a similar sort of thing, there have been a number of governments who have mandated certain kinds of adoption across their entire operation. This is literally the government, on government-owned infrastructure, and they often can't get it done anyway. So, the idea that, magically, if we just scaled that up to the whole Internet, it would work, I'm skeptical at the very least. The real question is, Is there a real problem here? In the case of v6 adoption, for instance, we've actually seen market mechanisms work there. It's turned out that's been much more catastrophic change than we expected. People seem to be willing to hang on to their lousy old approach to things, and add more and more epicycles on the outside of their model, until, finally, it just collapses, and then they go all the way over. What's interesting about that, is the question of whether somebody has actually done this diligence? In a lot of the providers, what you see is incrementalism tends to be preferred. This is true of engineering, generally, right? Engineers like incremental solutions. It's better to just build a little thing on top of the stuff you've got, than to think about it more deeply. Under those circumstances, then, the question is, if this supply chain is really long, this is an absolutely critical point for providers. Yes, banks, and other people, are trying to put themselves in multiple clouds. If it all turns out that they're completely dependent on a single provider, some way, way down the stream, and this is true at every layer. I want to point this out. We used to see this all the time in data centers, people would be, like, completely redundant connectivity, and everything like that, and you go, like, four kilometers down the road, and you'd find the trench in which all of the cables came together. Naturally, that was the place that the backhoe showed up. That's how we know, right? Because the backhoe showed up, and so we got a hole in the ground, and we got a hole in the fiber. It was actually interesting, near my hotel, I was watching them do this one day. I was, like, I'm pretty sure that's some fiber, and sure enough. This is still a problem today. That's the critical point. In terms of this question about neutrality, and whether there's a policy thing, I think it would be very, very interesting to say, look, as an urgent public matter, there are certain cases where you absolutely need to be able to trace this thing all the way back. I'm pretty sure that the Department of Defense people really do know where their cables go. I'm pretty sure they really do know who owns them. I'm pretty sure that, under those circumstances, they can force redundancy, if they need to. Why should it be only the Department of Defense that has that ? That's a critical policy question that I think would be good to to address, under these circumstances?

Emily Taylor

Jennifer, anything to add to that?

Jennifer Cobbe

I think the question, should we be waiting for one company to try to change these things, or one government to try to change these things, the fact that we even considered the possibility that one company could change these things, is in and of itself a bad sign, because it means that company has a significant amount of power over how the infrastructure works. That's probably not a good thing, really. The way that this is consolidated, the way that it operates, it means that a very small number of companies are operating at an absolutely critical level in society, completely, without any degree of accountability, without any sort of legal framework that is treating them in any particular way. They can exercise their power, essentially, freely, however they want. The question is a really revealing one, because it brings up all those issues. I would much rather wait for a government to change things, than for one company to decide, because it's commercially advantageous to it, to try to fundamentally change how these things work.

Jesse Sowell

I heard that question slightly differently. Part of your question was, should we encourage companies to make some of these investments, in pushing forward, because they can. For instance, if we go back to, again, late 90s, early 2000s, whenever we look at some of these intermediary services, that we now know and love, and enjoy, cloud compute, CDNs, etc. In those days, that was one or two companies that were sufficiently well heeled, either through VC money, or some other means, that were able to go and innovate, and develop, and it was a risky proposition. There was a lot of learning, and risk taking, involved in deploying some of these large scale CDNs. Now, we see that CDNs are, and I know a few people in the audience might throw something at me for saying this, CDNs are...

Emily Taylor

Content Delivery Networks, a way of getting the content very close to the user, often a make or break thing in an Internet exchange point. Sorry, just to jump in.

Jesse Sowell

Sorry, I was too comfortable with my language. These content delivery networks essentially make sure that your Netflix doesn't jitter, because the content is the content caches are close enough to you that it has a much shorter path to get interfered with, or to experience congestion. A lot of these folks, in the early days of developing the necessary technologies were making that work, for a lot of things that we really enjoy, they had to make those investments, they had to make those innovations. They were only a few at the time. In those early days, those guys were the concentrated market makers in that space, but, as technology does, it was commodified. Now we see, and this is where some people in the room might throw stuff at me, CDNs are a dime a dozen. A lot of the big actors that relied on CDNs in the early days, Apple for one, Netflix for another, they have now gone off and made their own special purpose CDNs. So, in many ways, there is some virtue to elements of what of what was baked into your question, which is that sometimes you need these big actors to use their capital resources to invest, and innovate, and push these things forward, not only just the technological innovation of, how do we make an efficient cloud cache, how do we do effective geolocation, but how do we develop that deep operational knowledge about how to deploy it in this complex network? That is not something you can do in a lab, you have to do it in a live Internet, and that's actually a risky proposition.

Emily Taylor

With five minutes to go, just before drinks, I think we've got time for one more round of questions. I think we had...

(Participant)

Is there a role for Chatham House rule?

Emily Taylor

Okay, we'll ask the panel to respond to that. You were waiting at the front, and then the chap at the back, and then we'll come back for some closing remarks.

(Participant)

Emmanuel. Member of Chatham House, I want to find out {unintelligible}, in integration of, for example, DNS in the Internet integration, how can difference scaling techniques. Can that be a use factor? Is that related only to policy, or to the technical aspect of it?

Jesse Sowell

If I understand it, you are asking about different scaling techniques.

(Participant)

Yes, different scaling techniques. Which of this technique applicable to policy, to the technical side of the integration?

Emily Taylor

Thank you very much. Before you answer that, I've actually got two questions. So go to Patricia, and then the gentleman behind you.

Patricia Lewis

Thanks. Patricia Lewis from Chatham House. Just a quick question, what difference is quantum computing going to make to the Internet. Thank you.

Stephen Castell

Should I give you my question? While you're thinking about the answer to that one?

Emily Taylor

Yes. Please give your question.

Jesse Sowell

Stephen Castell, independent consultant. Kind of related. We all interface with the Internet through applications, essentially. This is what this seminar's about, what lies beneath the application level? What does lie beneath it, is a whole lot of other software that we as users don't directly interface with. Largely that software is kind of dumb. It work to standards. It works to the standards implicit in the IP protocol, and the communication layers. I guess there may be some intelligent software to do with routing, and rerouting, of the transport layer, for load sharing, or failure of a node, but that's deep down

somewhere in the servers, of the communication software. To what extent do any of you see more intelligent software coming further up, near the application layer, so, it's not in the domain of any one particular application supplier, like an Apple, or a website owner, or a Google, or a Facebook. Could it be very intelligent. We could all get at, and use the intelligence, but also its intelligence will be monitoring our use of it, much more at the gross Internet level, not through one particular person. And then, who does who does the results of all that intelligence then belong to

Emily Taylor

Thank you very much. So a lot there. I'm going to ask all of our panelists to keep their remarks really brief, so that we get to close on time. Starting with the last question, it's about more intelligence in the network. I suppose, 5G, we have a more intelligent software-driven network. Who owns that, who owns the results of that? Will quantum change things and how? There's the question about integration, scaling, and also the role of Chatham House in the Internet of the future, in this consolidated network. Jesse, should we just start with you? A couple of quick remarks? And then we'll come back to Jennifer, and close with you, Andrew.

Jesse Sowell

There is quite a lot of opportunity for integration and scaling. That's fundamentally what has made the Internet work from the beginning. What's really interesting is, when we look at some of these non-state institutions, some of these groups like the IX associations, like Euro-IX, or IX-F, or pick your favorite IX association, they are about sharing information that is necessary for scaling of, at least, Internet exchanges, in this case. I know that there are other operator groups that also are concerned with this. There are opportunities, and I'm happy to chat more during drinks, just to close this off and move to the next one.

Emily Taylor

Thank you. Jennifer,?

Jennifer Cobbe

I guess the first question that I'll answer is the role of Chatham House.

(Participant)

Specifically, the Chatham House Rule. Attribution and anonymity.

Jennifer Cobbe

That's not necessarily challenged by the Internet, There's a degree to which, once you put things on the Internet, everyone can see it. But also things tend to disappear from the Internet, as well. Anonymity's many times a good thing, there's not any necessary reason to take that away. That's probably fair enough. The question about more intelligent software. It sounds like you're asking whether we should roll out a really extensive global surveillance apparatus.

Emily Taylor

We have one already, don't we?

Jennifer Cobbe

We do. We don't need another one. I think that's really not a great idea. And then the quantum, what difference does it make? I guess, in the short, medium, and perhaps even long term, probably none. because we're plenty of years away from it.

Emily Taylor

Last year, or a few months ago, there was an announcement, was it from Google, about quantum?

Patricia Lewis

We've got quantum networks being set up now.

Emily Taylor

Okay, Andrew?

Andrew Sullivan

The fact that people have managed to get something to work, doesn't mean that they can get it to deploy. I have personally been responsible for several things that I've managed to get to work, but couldn't get to deploy. I want to take these three bits about, the Chatham House rule, the quantum computing, and the question of this complete surveillance system, as one question. I think that there's an important point here. The Internet is maybe the most human technology we have ever invented. Its entire point is to get communications from one end of a thing, to another end of a thing, without putting a lot of intermediaries in the middle. That's an ingenious mechanism that we've come up with. We recognize that there are all these intermediaries, and what we've done is, we've gradually given people a toolkit by which they can build themselves a network, such as Tor, that was mentioned earlier, that allows you to do things in a way that preserves your anonymity, perhaps it preserves your freedom, perhaps it preserves the things that you want to do, whatever it is. There are a lot of interests aligned against that. One of the dangers from quantum computing is just the sheer horsepower that it gives, because it's going to roll out in an uneven way. Some people are going to get it first. That means that we've got a potential threat to the way that this technology has been deployed. Right now, we have ubiquitous technology, right? Every one of us is carrying way more power than the original IMPs, way back on the ARPANET. That means that we need, actually, public discussion about those kinds of issues. Many people accuse me, from time to time, of being too interested in the technology in my job at the Internet Society. I disagree with them, because the real questions here are mostly not technological. Most of them are social ones. What that means is that we need to have that conversation about how we're going to address the various threats. The approaches that we've been taking, so far, to these kinds of threats are, to get the large, various, platforms together, in a closed room most of the time, to make decisions about how that surveillance is going to work. The question, then, Will the Chatham House rule be possible? is one that we have to keep alive for ourselves. Is this something that we can trust? Will we be able to trust our devices? Or will they be tattling on us all the time? The current designs that people are proposing, sort of require that the devices be tattling all the time. If we're getting end to end encrypted, the only place that you're going to be able to do that surveillance is on the end. It seems that the current push is towards that kind of surveillance. The real question, that we have to ask ourselves, is, Are we going to have that debate, honestly? Maybe we want to build a better surveillance system for the globe. I personally don't, but, I mean, maybe that's the social choice

we make. If we make that social choice, it seems to me we ought to do it with our eyes wide open. That's something that maybe we're not having an honest discussion about.

Emily Taylor

Thank you very much. That brings this part of our session to a close, there is a reception upstairs. I would like to thank the audience, for your questions, and your participation. And of course, I'm sure we would all like to thank the panelists, for their insights, and sharing their knowledge with us. Thank you.