ALLAN FRIEDMAN: Welcome back, everyone. Hopefully, we still have a few folks on the webcast from their offices and some folks on the call, I see; so that's great. Thank you or rejoining.

Folks in the room, hope you had a good lunch, made some friends, built some bridges across professions and sectors and expertise. That is the spirit of multi-stakeholderism that we have come to know, love and occasionally roll our eyes at.

We're going to pick up with the discussion from the working group. This next working group, the final one, is focusing on some of the real meaty policy issues that we're going to be facing. Then after this, we're going to talk about it. Then we're going to zoom out and have a little bit of a big-picture discussion about where this broader initiative is going and how we can sort of get the outputs out to the wide world, and what is the pathway that the work that we've been doing is actually going to lead to better security in the ecosystem.

John, you're going to set us up in that direction.

JOHN BANGHART: Thanks, ALLAN FRIEDMAN.

Thanks, everybody, for sticking around. I know we're right after lunch; it's a little warm in here. Under your seat, there are pillows; so feel free to grab those if you want to take a little nap.

My name is John Banghart. I've been helping out with Working Group 4. I've been participating in some of the other groups as well, but I've been co-chairing this one along with my co-chair, Vic Chung, who is on the line, wasn't able to be here with us today. We've had a really interesting time with this working group dealing with incentives and barriers and adoption.

Just to sort of frame this out – just a little bit of history for those of you that weren't on the group, and I know there's at least a couple of you who have been. We struggled right out of the gate to figure out what the right sort of angle of attack is here because the concepts of incentives and barriers are pretty broad. Even in the scope that we've talked about here – that is, limiting this to primarily consumer devices, limiting it to upgradability – even within that narrowed scope, trying to address what are incentives and what are barriers can get fairly nebulous fairly quickly.

We did talk a lot about that at the beginning. We didn't want to take the obvious approach and say, well, okay, companies are incentivized to make profit; consumers are incentivized to buy new things; there, we're done, we did it. That would have been too simple, and that's sort of obvious. We also knew that we could go down an awful lot of rabbit holes here. You could choose a particular type of device, a particular type of industry; and you could spend a ton of time sort of working your way through into a lot of detail.

So instead, as we thought through these questions or tried to arrive at this, we were thinking, well, rather than focus on any one particular device or one particular piece of this problem, let's think about the decision-making process. What goes into the decision-making process? What influences decisions around do we build upgradability into products? Do we *not* build upgradability into products? How do consumers factor into this?

We heard some stuff about consumers about some of the other working groups. So that was the decision that we took, for better or for worse; and we'll find out.

This is a bit of an eye chart; I apologize. I didn't create handouts, even though ALLAN FRIEDMAN gave me the opportunity to do so. I chose not to, so I now regret that. But in any event, these slides are available; hopefully you've received them already. But we had to do some definitions right up front. We had to at least come to an agreement in the group if we're going to talk about incentives and barriers, we need to understand a few things. We need to understand what are the different vectors where incentives and barriers could come into play? Also, who are the stakeholders when it comes to incentives and barriers?

So particularly I think when you look at the major stakeholder groups, we settled on three; and we sort of made these definitions up. They're loosely grounded in other definitions that you're probably familiar with, and certainly we'll be looking for feedback from the broader group on whether these make sense. But producers, that is anybody that is involved in the production of an IOT device or service; and I'll come back to that.

Users can be humans; we also treated users as being machines. We had a notion of machine-to-machine IOT; that is, one machine using other devices for some extended purpose, which is why we chose the term "user" instead of "consumer." We wanted to recognize the fact that devices will interplay with each other; it's not just about the human device interaction. And then regulators – we all know what regulators are. The environmental, interactive and scale will become a little clearer as I go through some examples here.

So what we did was based off of these definitions, we tried to build out what you could call a taxonomy and tried to build out some sort of a framework in which we could develop some use cases. So as you look at this, you can see on the left side we've got our primary stakeholders. In each case, we tried to break it into some sort of logical category. We recognize this is imperfect. We recognize that there are levels of granularity that go far beneath this. But for the purposes of what we were trying to achieve, this was a useful construct.

So importantly, and we heard about this earlier, it's easy to think about software and hardware in terms of a device; but there is this service element. There is the networking; there are the protocols; there are all these other pieces. So if you think about a light bulb – we've been using that example today – there's somebody who manufacturers the hardware, the physical light bulb. There may very well be somebody else that manufacturers the software that runs the light bulb. There is some sort of a network operator that controls the network between the light bulb and the app that you're running. There's another developer that may actually develop that app or may have an ecosystem of apps, like Alarm.com or

somebody like that. So there really are multiple pieces here, even with one single device; and we wanted to make recognition of that.

Then the environmental interactive and scale pieces will become a little bit clearer when I run through a couple examples, so I won't spend a lot of time there at the moment.

Users, I mentioned, we think about this from both human and machine. Regulators, we used – took a little bit of liberty here; and I used the term "regulator" to be both the normal enforcement regulator but also voluntary. So thinking about trade associations, industry organizations, that come together to agree this is something we want to do, PCI, others like that. So it's a little bit of a loose use of the word "regulator," so I apologize for that; it's a useful construct.

So based off of this, then, we tried to go and actually start doing some use cases. It took us a while to get to that point; and we were struggling quite a bit around, okay, well now we've got this notion of this taxonomy for this decision-making process, but what do we do with it?

So we finally started to go through some use case developments. We developed several of them. We are going to produce a white paper that will have more information in it. It's not ready to go yet, which is why you don't have it; but that will be one of the outcomes. Just some examples we came up with: connected dishwasher; I have that emboldened here because I'm going to use that as our example as we dig into this a little bit deeper. But supermarket devices to track automated inventory, trash cans, things like that; the list goes on and on and on.

We recognized that risk is an important factor here, and it's come up once or twice in our conversations today. We actually went on a bit of a – I won't call it a tangent – but we took a turn at one point and started to think about this in terms of risk and risk management. Perhaps unsurprisingly at this point, what we found that while that is potentially a useful avenue, it's also again another type of a rabbit hole. You could go on and start talking about risk effectively eternally because you have to boil the ocean when you start thinking about the interconnectedness of all of these devices.

I forget who it was, but somebody this morning was talking about the fact – you know, you have a hub that's connected to a thermostat that may be connected to a furnace that's connected to all these different things; and there are different risk elements associated with each of them. So while we did go down that path, we decided that's not going to be an approachable path for us for the purposes of this discussion. So we did do that. I wanted to mention it because it's an important thing to remember – that this is all about ultimately managing risk.

Then we thought about personification and design thinking to try and put ourselves into the shoes of what would these stakeholders be thinking. Our group is fairly diverse; we've had a good core of people that have been there sort of week over week: IT, civil societies, trade associations. We haven't had a lot of manufacturing representation; so that's an area that if that's the space that you work in, we would love to have. Think about that producer category that I had up there; that's an area that I think we could really benefit from some additional insight. But we, again, tried to put ourselves into the shoes in these use cases and move things forward.

Again, a little hard to read – I apologize. But what I did here was think about that taxonomy that I had up on the screen before using the dishwasher example. So here I'm taking a producer of the software that runs the dishwasher; and what we tried to do was insert then some real-life things, some aspects. Some of the things it says up there – even I can't see it, and I'm standing right here – for example, an environmental barrier for a producer of dishwasher software might be – what's a good one? Well, scale, supportive legacy devices, and this has come up a few times.

So you've got somebody who's responsible for producing and updating software over time, but scale is a barrier for them because that's going to introduce complexity into them being able to push it out. There was somebody – it might have been the gentleman from CA or someone from over here, I don't recall – was talking about the fact that if you've got millions and millions of devices that you're responsible for, scale is an issue for you. It is a barrier to actually building upgradability into your products and could influence your decision to say, you know what? It's too big of a problem; it's going to get too expensive; I'm just not going to bother. So that becomes a barrier now for the producer.

I could go on and on through these, but I think you get the idea that what we wanted to do now was take this taxonomy and test it out. Does it work? Can we start inserting real world examples and arrive at something meaningful?

So that was useful up to a point, and so qualitative sorts of approaches make a lot of sense. I don't manufacturer dishwashers, and my dishwasher is not connected. So my input around all of that was largely just I'm going to make this up; we're going to imagine what this might be like. So while we couldn't necessarily come up with fully robust examples, given the group that we had, what we did want to do was now we wanted to take the sort of methodology, how do we think about this problem space, to the next step, to try and move it away from purely qualitative and start to get into a little bit more of a quantitative space.

We thought about it from the perspective of this just one example using a Likert scale to kind of ask questions around the various components that we came up with. Here are some examples, so I'll just read this to you. For example, it says "producer." One of the things that you could put in front of a producer is I expect to support this device for several years; that could be an incentive on scale. And they're going to say, I strongly agree, I strongly disagree, and so on.

For a user, new features are important to me, right? That's an interactive incentive potentially – strongly agree/disagree.

Regulator – this device impacts users' physical safety, right? I strongly agree with that/I strongly disagree with it.

Then trying to turn that into something useful, we went here. Again, I don't know how well you can see it. Where I have the producer up top, it says, "I expect to support this device for several years." I made that "strongly agree."

Producer – providing new features to users is important, "strongly agree."

Producer patching could introduce new vulnerabilities, "strongly agree."

I'm making these up, but you get the point. So what we end up here with now is given the nature of these questions and recognizing too that in any situation like this, you would have far more complexity to try to arrive at the truth; this is just examples. But you end up hopefully with something that says in this case you've got really strong incentives, but you also have really strong barriers. What does that tell us about what needs to happen?

Situations like that could actually be really hard to make progress on because maintaining the status quo is going to be the easiest thing to do because moving off of that space could become really difficult. Pivot that to something else where now we have really strong incentives – producing new features is important to me – but really weak barriers. Somebody says, "I don't think patching *is* going to introduce vulnerabilities," for whatever reason.

So *now*, you've got an interesting case where you've got strong incentives, weak barriers, and you can see that maybe now there's an opportunity to advocate for change internally. So if I'm manufacturing this dishwasher and I know that as a company we've got strong incentives to build up great ability in, but I have weak barriers based off of the analysis that I've done, that might suggest that I can work internally in my company to make the case we need to have upgradability in our products.

The opposite example is of course where you have very poor incentive but you have strong barriers – so again, using the same thing. This could be a device that maybe the manufacturer doesn't expect to have a long life period; they don't expect that the life cycle of this device is going to be very long. Providing to users the features – they don't care about that, whatever; you get what you get, and I don't plan on ever giving you new features. And I'm really worried that if I do patch stuff or if I build upgradability into my product that somebody is going to use that as an attack vector, new code is going to get inserted, whatever the case may be.

Now you've got a situation where you have really weak incentives, the company doesn't really care about it for good or bad reasons, but you have really strong barriers. They're really concerned about the potential risk that might introduce. So in that case, you've got a situation where instituting change on the inside of the company could be really, really hard because you're not going to be able to convince the leadership, you're not going to be able to convince the Board, that building upgradability into your product makes sense because there are not enough incentives; and there are barriers around it that are really going to be difficult to get past.

So that's all looking at it from a single stakeholder perspective; I was using producer there. But you can also extrapolate this to start thinking about this from a multiple stakeholder perspective as well. In this particular example, you've got a strong regulator incentive because maybe the product here is going to be dangerous to consumers. You may have a strong producer barrier with a weak regulator incentive, or you may have a strong regulator incentive and a strong barrier to scale factors and so on.

So we spent all afternoon sort of running through the different scenarios, but the point being that we think that what we have is a way to at least start thinking about these problems as opposed to just saying reduce liability, make it more profitable, convince users that they need this, that or the other thing. Those are all important aspects, but those are also well-known aspects. We think there's a complexity and some nuance around Internet of Things that requires a little of an approach that might allow some of that nuance to come through. This is by no means a complete solution, but it's something that we sort of arrived at and that we think as we start building more examples into it, it seems to be bearing out.

Moving forward on this subject, perceived and actual strength, obviously, are somewhat subjective. So in those previous charts that I had up there, whether somebody strongly agrees or disagrees, there's still subjectivity involved in that; and it's going to be hard to get around it. But that's an area that we're thinking about, and then what influences that. What influences a producer or a consumer or a regulator? What influences their perception of, well, that's a good incentive versus that's a good barrier?

Meaning of change can be very subjective as well. We spent a lot of time talking about the fact, well, is upgradability always a desired outcome? There's sort of an implicitness in this argument that we've been making that, well, upgradability is a better thing; and we've even heard from some people today that maybe it's not always the best thing to do. Maybe there's a risk strategy in certain cases where you might say upgradability is *not* a good thing. So you have to factor in the idea of what does change mean. You could foresee a circumstance where somebody has a product that is upgradable; and they say, "This is a complete disaster; let's take upgradability out of our product moving forward." They may see that as a positive change; and it may, in fact, be one. So we have to maintain that sort of perspective.

I think the key here – and I sort of mentioned this up front – is that we don't necessarily at this point have groundbreaking recommendations around here's how you incentivize this group and here's how you tear down these barriers. I think the point that we're trying to make is that that is so varied and so complex that there is no easy way to just say, "This is an incentive; this is a barrier." So we're spending some time continuing to think about all that and where it could take us.

Similar to the other working groups, what we care about here is first of all, have we identified the right stakeholders? Are there nuances in the identification of the stakeholders, the producers, the users and the regulators that maybe haven't been captured? Is there any utility to this approach, or *could* there be? Even if this approach of Likert scale followed up by mapping that out, developing some sort of combined qualitative/quantitative formula around how do we determine strong incentives, weak incentives, and so on and then ultimately what are we missing? Have we overthought this? Have we made it overly complicated? Should we just have stuck with reduce liability for manufacturers, educate consumers, and we're done?

I'll stop there and take questions/comments, angry lunch throwing. Yes?

MICHAEL AISENBERG: Michael Aisenberg, I'm Chairman of the American Bar Association Information Security Committee. On your "since our last update" slide under the regulator, I would suggest, just as an editorial amendment, that you substitute some other word for the word "standard," which those in the voluntary standards community are very protective of and consider it to be a term of art. I think what you're referring to here generically is something more like a norm or an authority.

JOHN BANGHART: That makes sense. I usually draw the distinction of lower case "standard" versus upper case "Standard," but maybe that's (inaudible).

MICHAEL AISENBERG: And then sort of in general under the concept of user, the way your charts bake out, given the example of a product that could slide along the scale from being a residential consumer product all the way to an industrial product, depending on what context the dishwasher or refrigerator were placed in, for example, sub-zero in a government bio lab versus in a residence, I would suggest the possibility of trying to run that same set of analyses against a static notion of the other two variables, but a dynamic across a spectrum of one of the stakeholders, in this case, consumer. Run the same thing for residential consumer, industrial consumer, government consumer, high-end consumer – around the same product. I think you might find the analysis would vary even within that narrow frame.

JOHN BANGHART: Well, I hope that it would. It would be weird if this whole methodology ended up with the same result every time. I will say that your point is 100% valid. We've specifically stayed away from running analysis outside of the really specific sort of home consumer space because that really is the scope that we were given to try to address this and not get outside of it.

But I think you're right. I think, hopefully, if you start plugging in other things here, what the methodology gets you is some way of getting slightly better informed on given the set of circumstances, given my target market, given the potential risk, how easy or hard, or how motivated or unmotivated am I going to be to actually build upgradability into my product.

MICHAEL AISENBERG: And I think that issue goes to the question of the establishment of norms and what communities can be relied on to produce norms. The pressure for safety and security and efficacy of an FDA type would apply to a product, like a refrigerator, being used in a bio lab or pharmaceutical lab industrial setting at a very different degree than if it's in the residential consumer setting where the only thing it's telling you is whether you need to buy more milk.

JOHN BANGHART: Yeah, certainly, and that was part of the debate we had around do we want to turn this into a risk management exercise or a risk assessment exercise? We decided, for the reasons I laid out, that our group was not going to be able to tackle that problem for all the reasons you're highlighting.

There is, pragmatically speaking, an infinite number of scenarios that you could sit and lay out; and it changes the calculus on this stuff. So it's challenging; we had to scope it somehow.

JAMES SIMISTER: John, can you talk a little bit more about your planned white paper? Like, who is your target audience?

JOHN BANGHART: I don't know, to be honest. I mean, part of what we need to do here is figure out – I took it off the slide, but are we even heading in a direction that's going to be useful? We laid out some of

the other avenues we could have taken but decided on this one; and I think we need to understand from the broader group, and try and get more people engaged on this, is this a path forward that we want to continue to explore? I think that will help us define who is the target audience.

The paper basically has the same content, just with more examples, more narrative; but I don't want to publish something if folks are looking at this and saying, "Well, this maybe, sort of, might be intellectually interesting; but no one's ever going to use it."

JAMES SIMISTER: What do you think?

JOHN BANGHART: I think it's more than just intellectually interesting. I think it needs more investigation. But any methodology or taxonomy is going to be that way. So we have to figure out do we want to define – or further refine, excuse me -- this approach, following even some of the examples that he gave around, all right, let's start plugging in other use cases here and see what comes out of this. Ultimately, what our original target audience we thought was going to be – and maybe still is – the decision-makers around who decides or who influences whether upgradability and patchability are available in a device or not; and that's where we arrived at our stakeholders.

The producers can make that decision. Consumers can influence it by saying, well, I'm only going to buy this product or that product if *Consumer Reports* says that it's really upgradable; and then, of course, regulators can. The goal, I think, is to try and create some sort of a toolset, to use the term more loosely, that decision makers, regardless of where they are – producer, user, regulator – can help to motivate them.

I punted on answer to your question, but that's what I've got.

ALLAN FRIEDMAN: Let's start here, and we'll go around the room to capture the comments.

JAMES SIMISTER: This is James Simister from Panasonic. As you were talking about the stakeholders, one of the things that struck me is that integrators might be another category to think about. They're kind of a unique hybrid between a producer and a user; but I think they have their own sets of challenges, trying to sometimes bridge the gap between a producer and a user. But I was just thinking that might be another interesting stakeholder to consider.

JOHN BANGHART: I like that; that makes sense.

[MEHRAN]: White paper is a good idea; at a higher level, yes, it helps. But like as we are talking already, like use cases and test case studies, those things would be very helpful for the implementation. Use case, you may say application use case, okay, you can do this kind of a thing. Test

case, solution, references if you give like, for example, in one of the industries this has been implemented like this; for example, automotive, this is the kind of a thing.

The same use case may be played in a different way – like in an automotive or in a medical scenario or in a home automation or a lab automation kind of thing. So those kinds of things – if we have to show that as an example, not many, couple of them if we show -- use case, multiple test solutions; and this is the upgradability white paper -- so all these three may help in moving forward.

JOHN BANGHART: Makes sense.

ALLAN FRIEDMAN: Craig and then Ralph.

CRAIG: Thank you, I provided some comments earlier. My recommendation is we focus it on the producer. Users aren't going to become an NTI site. They're not going to be reading this; we know that, I think -- at least, we know that intuitively. But I will tell you that companies that are trying to bring products and services to market are hungry for direction, and we can't solve and serve all three audiences.

The other part – and it gets back to Michael's comments about some other types of devices – is in some of the working groups, we've really tried to keep focus. What types of products or services are we talking about? And I think in some of the groups, we said basically consumer-grade devices that are used in the home, used in the office – including wearable technologies. And by design, since FDA has sort of regulatory oversight in their world, not to swim in that lane; and the National Travel Safety Board has theirs, so not to focus on automotive.

But I think it's very important that we keep focus. It can't be so broad that it covers those other areas. Again, if we're trying to make an impact, when we look at 500 new IoT consumer-grade devices being brought to market every week, they need that prescriptive advice and direction more than anyone else. So that would be my recommendation. I'd be hopefully happy to hear if anyone has comments on those points.

ALLAN FRIEDMAN: Two-fingers, over here.

RALPH: Next time I'll do three; I'm a Boy Scout. From my little world, when we design and build data centers, we try and herd the cats around certain tiering – Tier 1, Tier 2, Tier 3, Tier 4 types of infrastructures. And possibly, under certain consumer groups and industrial, automotive --all these different groups – it might fall out that there are certain reliability aspects to group consumers into to try and define parameters around upgrades in security, if that makes any sense.

JOHN BANGHART: I didn't track it 100%. I'm sorry, you lost me somewhere. Are you suggesting that there is an additional sort of granularity to how to define a user that may impact whether security updates should or shouldn't be available?

RALPH: Yeah, exactly.

JOHN BANGHART: I think that's true; I think all of this is true. What I struggle with – and maybe I'm the only one that struggles with it – but when we start to talk along those lines, the horizon just keeps going further and further out. So you end up with a situation where you say, wow, no matter what we solve, there's always going to be more problems. I think there's a balancing act between recognizing that the horizon is infinite, but that there are immediate concerns. So it's what do we do with the *immediate* concerns?

My background is I'm an operator, right? I'm an operations guy, so I think about how do I solve the risk now that plugging in this connected thermostat is bringing into my home or into my business? I recognize that down the line, somebody is going to create something that's going to connect to my thermostat that's going to connect to something to connect to something. I can't fix that problem; that's just market-driven. So that's going to happen no matter what. But if I can piecemeal the problem out, if I have some way of approaching it and saying for this given device, I can at least somehow gauge the risk and put some sort of mitigation in place that helps the risk to at least be knowable and *partly* controllable, if not entirely. Well, I'm repeating myself; but I think scoping here is incredibly important.

ALLAN FRIEDMAN: Two fingers from Ralph, then we're going to have Jason on the phone and then back to Michael.

RALPH: So with respect to – we've identified the right stakeholders, I suggest that there's one that's missing, which is everybody else. And I say that somewhat facetiously, but it's the fact that both the stakeholders who could be potentially affected, for example, from distributed DDoS attacks, are not the actual user of the device. So there is this impact to others that are affected by what is done with the device itself.

I've heard IoT security being identified as a tragedy of the commons; so in some sense, that might get to incentives. But to be a little bit more practical in the application of the use cases you're identifying, I think there are a actually a fairly limited number you need to look at. IoT Analytics does a great job in terms of segmenting the IoT Marketplace, and they pretty much broke into two different categories at the top level. There's what they call IoT to C, which is to consumers, and IoT to B. I think the incentives are different; but within the IoT to B category, they're fairly similar; and within the IoT to C category, they're fairly similar. So if you look at just those two scenarios, they may drive out the fundamentals. They vary to some degree; but orientationally, I think they're fairly good. So if you think about it in those two broad categories, I think you can simplify the number of use cases you need to think about. Those two generalizations may be sufficient.

JOHN BANGHART: Interesting.

ALLAN FRIEDMAN: We have Jason on the phone.

Angela, do we still have Jason?

OPERATOR: His line is open.

JASON: I was thinking about the target audience question, and it got me thinking about, well, how would this information be useful to *me*? I think getting some information not just about – because we can talk about the voluntary stuff versus the more heavy-handed stuff, whether that's regulatory or if somebody is deploying a solution through a provider and the provider is saying, "You must secure X, Y and Z, or else you're not going to be included in our service offering."

That kind of stuff can be coming down from on high, but I think other worthwhile information would be from either the manufacturers or developers who are building this stuff. What is the lowest threshold that would convince them to do this stuff? What is it that would make it easy enough for them to do it that they would say, well, yeah, once it reaches these criteria, then we would absolutely build upgradability into our devices?

JOHN BANGHART: I think that's right, and I think we try and capture that in the taxonomy that we put together – trying to fit exactly what you're saying as both an incentive or a reduction in barrier. So you could foresee, to your point, a technological advancement. Maybe it's improvements in how we patch products, whatever it may be that lowers the barrier to the point where now the incentives outweigh the barrier; and so the producers are sort of automatically, then, incentivized to do it. I agree with you 100%.

ALLAN FRIEDMAN: Michael and then Justin.

MICHAEL AISENBERG: I want to strongly disagree in part, and agree in part, with the gentleman across from me who commented regarding the market segmentation and applicability issue. I don't know how many other people that participated in this exercise were part of the NIST Public Working Group on Cyber-Physical Systems, but one of the reasons that report – those 200 pages – are, in my view, so inaccessible to the point of being constipated is that there was a failure to do exactly what you *are* doing here, which is develop a suite of use cases as precursors to any advice – and their document was meant to be advice – to standards bodies, developers, a range of premarket communities.

So I think there's great value in doing this, irrespective of whether you're going to apply it to the consumer space, to the industrial space, or some combination of those when it comes to autos and medical devices, while it may be that there is Federal preemption when it comes to the actual use of vehicles or medical devices. I just heard the FTC Commissioner speaking at an IoT panel last week and agreeing with a comment that I made in front of the Attorneys General two weeks ago, which is there *will* be state liability

legislation separate from the FDA's review of connected medical devices based on these issues – on the connectivity and IT issues separate from the medical/clinical efficacy issues.

Those will not be able to be preempted. They're not going to defang the ability of the citizen to get a remedy when they are injured by a device because it incorporated bad software and became hacked. The clinical efficacy, perhaps, will still be preempted. And I think the same thing is still true for vehicles.

So I think there's a huge market for these use cases that define an appropriate suite of preconditions and analytics available to the standards developers, to the early stage product engineers, that helps them define what features need to be incorporated in their products in order to satisfy what will emerge to be consensual demands from the marketplace.

JOHN BANGHART: Thank you.

ALLAN FRIEDMAN: Harley?

HARLEY: So in talking about scope and what will bring the most value on this particular topic, you should feel appropriate; but my personal opinion is that I think that you'll find enough diversity of use cases and diversity of incentives and barriers to adoption among producers alone. I think that gives you the most bang for your buck.

We've talked a lot about diverse and how varied IoT deployments are and all the different products that there are; and I think that you'll find different incentives and barriers among different deployments, depending on the vertical. I think it's also a lot more difficult to nail down things like barriers and incentives for users, for regulators, because I think a lot of them will be a lot more amorphous.

Regulators, for example, have – I've worked in Government; I know that you have. There's politics, right? And that is actually a huge barrier or incentive, depending on which position the person is in, to regulate or not to regulate and how to regulate. Same thing with users – a lot of users come from such a huge array of backgrounds. Anyway, you get the point.

In the end, I feel like what we're talking about is the upgradability of IoT devices; and the only people who are really be able to do that, who are really going to be able to build in that functionality, are the producers. I think that you could have a full white paper – you'll have no shortage of content in just discussing the incentives/barriers to adoption for producers alone. In fact, I'd be a little surprised if it's not been done already; if it hasn't, then this fills a gap.

JOHN BANGHART: I don't disagree with you. I struggle to figure out – and I think our group has struggled to figure out -- how do we separate producers from consumers when you talk about incentives and barriers because they become incentives and barriers to each other.

But what you may be describing is maybe we need to reduce or alter the tact a little bit here to think about, okay, let's recognize the fact that ultimately producers are trying to sell to consumers. Let's recognize that, and let's those types of incentives and barriers off the table and think about scoping it to – maybe it's more of the environmental and the scale issues and less the interactive issues as we've identified because that's where most of the consumer interaction comes in. Focus on what are the technical advancements that might make this easier for people to implement. Is that kind of where you're headed -- maybe?

HARLEY: I was thinking of answering the question: Why don't more producers build upgradability into their IoT? Like the company that made the cameras that were the primary bots for the Marai botnet. Why didn't they include upgradability features in their cameras – why not? I don't know if they're willing to talk about it, but getting some insight directly from that source if you're able to make contact with them – that sort of thing. Hearing from them, what is it? Is it technical? Is it a profit margin issue? Did they just not think about it before, and now they realize that it's serious?

Like I said, I think that you'll find different companies probably have different perspectives on all of those; and you could probably make a full white paper just based on that.

ALLAN FRIEDMAN: It's almost like we have some people who are in the technology division business that have some opinions on this.

JOHN BANGHART: Yeah, it's cost – it's really simple; it's cost.

HARLEY: I was just going to add I bought one of those cameras. It was upgradable; it just didn't follow all the best practices on how to make it upgradable. It was very obscure, and you had to go download a piece of software from their site; download another software manager to download it to your camera. You had to jump through a bunch of hoops in your own network to find your camera.

JOHN BANGHART: So this comes back – I want to respond to this because this is good. I think what you're asking and your answer – this gentleman's answer – I'm sorry, I'm forgetting your name.

RALPH: Ralph.

JOHN BANGHART: It's what's underneath that interchange that I think we're trying to drive to. So you're asking the right question, but your answer of cost is the right answer. But then, what then is the incentive? How do we incentivize *you* by either making it more cost-effective – that is, reducing or lowering the barrier or creating some other way for you to be able to offset the cost of introducing the upgradability, hopefully without having it go to being regulated, which is what we all want to try to avoid. That was the dynamic that we were trying to capture; and you guys just sort of played it out in real time, which was great.

RALPH: That's why I said I think if you look just at the consumer market, the dynamics there are pretty clear. I think it's an example you can use and fairly easily use it as a use case and drive out what those incentives are. So cost is one. Another one is control. so these are conversations I've had with the manufacturers who say, "Hey, if you're going to automate this, I don't want updating my device to be under the control of some other manufacturer. How do *I* control my destiny from that perspective?"

And then is my business model, which relates to cost, which is today my business model is one where I sell the product and I'm done with it, period. The fact that I'm making it smart doesn't really change that business model. So there are fundamental business models that have to change to make it a sustainable product that says, hey, I have to build into my cost ongoing support for this device.

HARLEY: You're getting at something that I would frame as breaking out cost. So if cost is really the big barrier to producers, including upgrade capability into their devices, what cost exactly? Siri's working group came up with a long list of technical capabilities that are necessary in order to form an update. Is it possible to quantify those and to say, okay, these are the most expensive parts. Is it hiring the security engineer who's going to have to come up with the security update, et cetera?

RALPH: So the most common cost that gets raised to begin with is it costs me to put security in. So if I'm going to support upgrade, I have to have twice the memory. If I'm going to encrypt things, I've got to have some ability to encrypt. That means I need to use a more expensive part, which means my product costs more, which means my competitor can undercut my price.

ALLAN FRIEDMAN: Justin and then we'll go to Kent on the phone.

JUSTIN: I have just two points I want to talk about here. One is just some experience that we've had with convincing people to use update stuff that we've provided. So basically, we go; we provide working implementation of it; we show people it works; we do a couple of example integrations; we show them, hey, this is how easy it is to go and get this going. Then usually, once we get our first major adopter, then everybody else wants to play follow the leader. So that's, I think, a viable strategy for getting something into the market -- is actually build good tooling around and have something that people can use out of the box, they don't have to invent.

I think that in this discussion about costs and all this other stuff, I think if I had been asked that question, I actually would have said first to market because first to market because getting that first mover advantage, getting your product out earlier, getting ahead of other people, at least in the very startup, the Internet of Things devices really makes a big difference -- for an established company, maybe not as much. So that was the first point I wanted to make.

Then the second point I wanted to make that was brought up a couple of times is that I'm not sure that IoT – there are different ways to view IoT. We should do updates for IoT, and IoT is everything under the sun that we can think of when it comes to IoT. But I actually from having worked in a lot of detail in the automotive space, and worked in a little bit less detail but having worked quite a bit with some folks in the

medical space who understand that space very well, and worked a lot in the kind of cloud server/desktop space about updates, I think the domains are different enough and the constraints are different enough that if you tried to have a single set of guidelines or a single sort of operating environment or a single kind of description of what an updater is supposed to look like, I don't think it fits because medical isn't the same as automotive isn't the same as IoT isn't the same as an airplane.

I think that are some point, trying to figure out do we want to only collect the high-level ideas that make sense in all those domains and fuzz over all the rest of it; or is the intent to come out with more detailed specifications that describe at a lower level of detail for specific domains our thoughts on how they should be doing it. I think that's also a discussion that should be had because this isn't a one size fits all.

JOHN BANGHART: Let me react to that. I know we're going to get close to time here, but just bear with me. In my mind – and so maybe this is a failing of me in terms of presentation – but everything you just said is sort of like the introduction to our white paper. It's why we took the path in the first place. It's a recognition of the fact that there's no way for any one group of people in a broad context like this to sit down and say, "Here's an incentive that will work for everyone. Here's a barrier that everyone has to knock down," which is why we took this approach. It was can we generate something useful, given what you just described, which is this is a highly-nuanced, highly challenging area. Can we enable decision-makers?

So if I make medical devices, is there an approach I can take to get me started thinking about what is my risk? What are the reasons I can't upgrade my medical devices? They could be technological; they could be costs. What are the barriers?

Now they can start to think about the fact, okay, now I've thought about it in this context. How do I make some changes, either internally or working with external parties, to now get to where to where I want to go?

JUSTIN: So does that mean that the second working group and the first working group and everything should be siloed, but there should be some master worker that does it overall? Do you see what I'm saying?

JOHN BANGHART: This has come up before, and we've talked about – well, maybe you wanted to talk about it later – but we have talked about how do we start to bring the groups together. We've been working largely independently. There's been some cross-pollination just on an individual level, but it's been largely independent. There has been some discussion around, okay, now it's time to start bringing groups together. And I think – I'll be honest – some of the stuff I saw today from the other working groups, it's the first time I've seen it. So it got my sort of juices flowing around, oh, you know what? I see what Harley is talking about; I see what Kent is talking about; I see how that could interplay with what we've been talking about. So I think there's a logical next step there that gets to what you're saying.

ALLAN FRIEDMAN: We're going to go to Kent on the phone and then to Michael. You guys are now having an appreciation of why this working group has had some fantastic discussions, but is still saying

what we're doing is clearly important but how to actually forward. You sort of get the tension of what we've been working with.

Kent on the phone.

KENT: So John does what John normally does; he answers my questions quite often before I even ask them. One of the things that I do want to sort of – and the reason for the question in the first place – was in some respect, it does come down to how we're going to package and deliver this from the standpoint of the intent of the effort is really to try to encourage the manufacturers to do the right thing initially. The consumer areas are fine and useful; but the reality is if we don't convince the manufacturers that there's value in trying to go through this process, this understanding as to an evaluation as to when does that cost become important to them, then I think we're not accomplishing the mission.

But if we can focus this in such a way as there is that kind of decision-making process that can be incorporated into a product development life cycle, does it matter? Is it so cheap that the consumer is going to throw it away if there's a problem? Is this something that is going to have a much longer life? Do we need to make the investments now to avoid the PR nightmares of being used in a DDoS attack and having our name scattered everywhere?

Those types of decision-making processes would be extremely valuable for organizations trying to go through this right now. We don't want them doing this on a gut basis. We want them doing this on some sort of evaluated process that makes sense for them. I think that what I see here is really good as that kind of basis, and we need to focus on the producers and convincing them to use a process similar to this.

ALLAN FRIEDMAN: Thank you, Kent. That's a really good way of framing this.

Michael?

JOHN BANGHART: You said he could talk.

ALLAN FRIEDMAN: Yes.

MICHAEL AISENBERG: I think as we consider the emergence of norms in this space -- whether it's industry best practices or voluntary standards or regulation or, heaven forbid, a statute -- I think it's important to separate those events and risks and liabilities which are preexistent in the industrial product from those that are added by virtue of their connectivity to the network. You may be inducing an opportunity for a privacy violation or an incendiary overload of an electrical panel by the connectivity, but there is still a universe of preexisting liabilities from the device overheating and blowing up its motor or going bad in some other way.

I think it's important to recognize that the norms need to account both for those things that were there in the product before it became a connected device and an IoT device, as well as those that result from the fact of its connectivity.

Again, I will say my read of your work is that it really bolsters the understanding of the additive risk associated with becoming connected when the device has a preexisting identity and a preexisting legacy of known faults and errors and liabilities that might be associated with it.

JOHN BANGHART: That's great, thank you.

I'll wrap it up. So this is exactly what I wanted to hear today, so thank you for that. I mentioned right off the top that our group has been very engaged, but we've been missing a couple of key voices. So I got to hear from a couple of people I would classify as producers today. I want to encourage those folks, whether the ones here in the room or others, to kind of jump into our working group for a little bit and help us to frame out these things in more real terms.

We think we've got an approximation of reality, but what I've heard today is that we need to get more real. We need more use cases – not a ton of use cases – so we need to think about how do we scope what the use case really is; but we need a few more that are more concrete, that can actually be more informative around here's how it has actually played out.

So again, you've heard this from every working group – you'll hear it again from ALLAN FRIEDMAN, I'm sure – but we want to encourage you. If you're interested, let us know. You don't have to come to every meeting; just come to one or two, share your perspective, help us to get that voice into the mix so that we can continue our progress.

ALLAN FRIEDMAN: Thank you, John.

And thanks to everyone who took part in this working group. As I said, I've been able to listen to most of the conversations; they've been fantastic, but they've all sort of been wrestling with this very large, society-wide issue. They've certainly carved off a fun chunk for themselves.

Now we're going to have a broader discussion. We've heard from all four working groups, and we want to have a discussion of where we are and where we're going to be going. There are a couple of different components I see to this, but feel free to jump in as we go on.

One is just the outcomes. We are in different stages. I think everyone has a clear trajectory, but where they are along that vector varies. The Communications group has a draft, and they're clearly looking for

feedback so they can actually get it out. Whereas the Incentives and Barriers group has said, okay, we've got a vague idea; now we just need to sort of make sure we're in the right direction – great, we're going to go forward. So how do we want these to actually go forward and then come out to the world? Do we want one document? Do we want four separate documents? Should there be an overarching document? Do you want cross-referencing?

Then equally important, once we have these documents out into the world, how do we make sure this isn't just a throw it over the wall and assume that everyone knows NTIA and the Federal Register so well they will obviously read it? How do we make sure they actually have the impact that we want, that we've been doing all this work for the last few months on?

The other thing I want to put on the agenda this afternoon is from a stakeholder perspective, we have this odd problem, which how do we, as the folks *in* the room, identify the folks who aren't in the room that really need to be part of this conversation in order to make sure that the outcomes really reflect the entire community and also have the legitimacy that comes from having consulted the entire ecosystem. So who are the folks that we should be engaging? What are the four that we should be engaging? I know a number of you have come here today to sort of be able to have a little bit of mention of the work that you're doing that is quite closely related to the problem and the challenges around software upgradability in IoT. So I want to make sure that we have some time for that.

But first, we'll first chime in and get some thoughts and opinions on the outcomes of what we're going to be doing. How do we want to take the work we're doing and actually aim towards the finish line?

MALE SPEAKER: Allan?

ALLAN FRIEDMAN: Yes?

CRAIG: Actually, it might be helpful to just kind of – what does success look like, and what are the metrics? For example, so we've worked on other NTIA working group efforts in the past. Do you measure unique downloads? What is the measurement of successive in the past that we should be thinking about? What capabilities are there? I think that would be good.

I think the other point is it's not only this working group, but I mentioned earlier in the week about the efforts of other U.S. Government agencies, and I think that's another area to be thinking about. How do we reconcile those things together – not just the efforts of this group, but these other parallel efforts – and what is the role of NTIA in those areas? I just wanted to push those out as we think about going forward.

ALLAN FRIEDMAN: So to sort of address your more immediate questions, I can say that from an NTIA perspective, we are the facilitator. We are not looking to measure or build metrics around success. Our job, as I like to say, is we push; but the steering has to come from the community. It has to be something the community wants. What we want to do is be able to identify what are the catalytic measures we can

take to say the work that we're doing -- which has taken a lot of time and a lot of effort, a lot of skill sweat – how do we actually get it out into the world in such a way that it has a meaningful impact?

And in terms of the broader government discussion, this initiative has never been – in fact, it is illegal to have a group like this and have a talk about other direct government policy. That would break the FACA barrier. But what we want to do is make sure that the work that we're doing here can be useful to our government colleagues and also the NTIA editorializing, sort of demonstrate the value of community- and industry-led initiatives as a path forward for this kind of problem.

Michael then Harley.

MICHAEL AISENBERG: It occurs to me that there are – and this may also go to who's not here -- but there are a couple of bodies that have recurring publications that articulate some of this at a national policy level. There's a triennial national standards policy document that is written in partnership between the Commerce Department and the Voluntary Standards community; ANSI usually holds the pen for it. I would suggest that that's an interesting place to consider going.

The Table of Contents of the Standards Catalog has 15 or 20 voluntary standards bodies who perhaps ought to be convening with this group, or some variant of this group, or at least polled for their collective consensus view of how do they think things might roll out in their individual spaces. And then your colleagues within the White House and the Office of Technology Policy are going to produce a national technology annual that clearly should have some reference to IoT in it, and maybe that should reflect some of the consensus emerging out of these working groups.

ALLAN FRIEDMAN: Harley?

HARLEY: First, on the final form that this work may take, and then second on metrics, I think that the work of the different working groups is each valuable as a standalone document but could also be then integrated into a larger whole. So I think that the target audience for a compendium of the standards, for example, is going to be different than the target audience for communicating the elements of IoT upgradability. I think that it would make more sense for each of these working group outputs to be matured individually so that you can send them out to those target audiences rather than handing them kind of a larger white paper that everybody has to dig through and that may not contain the specific information that they want.

Now, that being said, I think that you can also have an additional document that weaves these different outputs together. Whether or not we are willing to take that on and what that form might look like, I'm not sure; and I'm not sure that we're going to know until each of the working group's work product has matured enough to stand alone. But I do think that each one adds unique value in and of itself, and so I guess I would envision a final work product looking like the four working groups' documents and then perhaps something that synthesizes them all – or perhaps not the thing that synthesizes them all, you just have the four working group documents and that is the output.

In terms of metrics, I also think that the metrics for each of the working group's output is going to be different. For example, to use the same two working groups, the metrics for the body of standards, the compendium of standards, very valuable but I'm not sure how you measure that – perhaps downloads. But if I'm going to measure the output of our working group, the success of our working group, it's going to be adoption. It's going to be whether or not producers actually decide to communicate these things to consumers. So it's very different than, I think, the measures of success for some of the other working groups, which again goes to my point of each one adds unique value alone. It would be difficult to just mash them all together. I think that all four of them should mature and then we decide what to do.

ALLAN FRIEDMAN: Great, thank you. Thoughts – John, then Matt.

JOHN BANGHART: Yeah, Harley said a big part of what I was going to say; so I'll just reinforce it. I do like the idea of the synthesis document, even at a conceptual level, because I think it will force us to understand what we're missing. Right now, because we're operating largely in four silos which, again, is appropriate for this part of the process, I think there may be areas that are not being addressed because of the approach that we've taken. An effort to synthesize the four outcomes – not to necessarily merge them, but to sort of synthesize the concept and see how they fit together – might put this group in a position to say, oh, crap, we forgot about this other thing; or, now we see how this fits into a bigger picture. So just to reinforce I think that's a good idea.

The other thing – and this may be more tactical – is I do think we do need to figure out how to take advantage of the multisector nature of this and find the right ways to do multisector communication. We're doing a panel on this subject, and it's for security guys at a security conference. Okay, that's great and that's fine; but like so what? That's not enough; we need to find events where you can have a security person and a manufacturer and a policy person and somebody else, who are sort of on a panel together talking about this issue, maybe not at a cybersecurity conference. Maybe it's some other type of event, and I don't know what other sort of events cable box manufacturers have; but maybe there's an opportunity to reach out to those types of places and spread the message around to audiences that aren't used to hearing it in the same way the cybersecurity community is used to hearing it.

ALLAN FRIEDMAN: Fantastic, I want to make sure we get time to talk about the different events and venues. But before, I want to focus -- continue this conversation -- about separate documents, single documents, how that (inaudible)?

Matt?

MATT: Allan, thanks.

ALLAN FRIEDMAN: Remind us who you are?

MATT: Mathew with the U.S. Chamber of Commerce.

First, John, thanks for the work on the incentives and barriers. That's a very hard thing to put together well.

I just wanted to offer the thought that I think based on experience with a number of members, I think a single document compiling all the different efforts might be your best way to proceed. Something that combines a little bit of simple messages in terms of what you're trying to communicate so you're engaging both the new audiences as well as the experienced hands, as well as maybe some granularity – maybe not as much as the cyber-physical systems document that NIST put together. But I think some combination of both tends to work well.

Then in terms of how do we gauge use, I think if anything the NIST framework effort at least instructs us that trying to gauge metrics use is a tough road to get down. I think we all like the scientific method, but what policymakers want to do with metrics can often be something that we don't want -- anyway, just a word of caution.

ALLAN FRIEDMAN: I am going to let my NIST colleagues spend their time talking about the NIST framework metrics issue. We'll leave that outside of this particular discussion.

Michael?

MICHAEL AISENBERG: I just wanted to add one more stakeholder community that you might be interested to add, to the extent they're not already here; they may be and I'm just not aware of it. The State Attorneys General have through their Washington Office a technology expertise, and I think they would be very interested in letting folks know where the demands are percolating up from – the sub-Federal level and also hear what the various directions this has taken. At their technology conference a couple of weeks ago in Charlottesville, IoT was probably the most frequently mentioned phrase at the conference.

ALLAN FRIEDMAN: Thank you. (inaudible)?

MEHRAN: Yeah, I agree with Matt on having a single document kind of step, but initially giving some introduction/overview, kind of what we are trying to address in the whole document. And then put together all four workgroups in this one. Also then, as I mentioned earlier, the use case and the test case studies – implementation, that's more key to this, how to implement in different scenarios, like as you said. The same update has to be done (inaudible). So those kind of case studies need to be shown in the document; that becomes more valuable than just putting the words.

ALLAN FRIEDMAN: Craig?

CRAIG: I know it's a little harder to have one document; but from our own experience of having documents that reference others, and I know you may not like metrics but when we see only 10% of people that go actually download the other document, we see it as a disconnect; and they're both really important. So I think it's almost kind of like the life cycle of what we're trying to do or how they all kind of couple together, need to be ideally in a single document; if not, then each document needs to say kind of like, "Chapter 2, you need to go here." It needs to be really prominent – make it very clear that this is only Part 1 of 4 parts or 5 parts of the working group.

My fear is people pick up one document; and the point of reference is we launched our IoT framework that had the must-do principles, and then we had recommended principles. The recommended principles was a separate document, and we had like 18,000 downloads of the main document and less than 600 of the must do's. As people go through, they weren't aware of it; they got the one document, and then they went off on that – so our own experience.

In this area there are a lot of really good parts here; but if it's not a one-part document, I think we're going to really limit the overall goal of having a good solution for everyone. Thank you.

ALLAN FRIEDMAN: Harley and then Claude on the phone.

HARLEY: I want to advocate for both. I don't think that it's either/or – that you have broken up documents versus one document. I think that you can have both documents that stand alone and be individual files, and a single document that weaves them all together. I think that is the way you get the maximum utility out of each of the chapters; and it will also make it easier to target, as I said, which I think is actually hugely important.

Sending a white paper to a body that I know is only interested in part of the document and then pointing them to the particular chapter in the larger white paper which will, by the way, have to wait until all of the four working groups, like all of their work products are completely mature, seems to me to be unnecessary. I think that we can do both – so just putting that out there again. One document would be fine; I would think having the four stand alone as an additional vehicle is also important.

ALLAN FRIEDMAN: All right, thank you, trying to have both aspects seems to make sense. We've got Claude on the phone.

CLAUDE BAUDOIN: Hi there, can you hear me?

ALLAN FRIEDMAN: We can indeed. Thanks for being on the call.

CLAUDE BAUDOIN: Hello, everyone. My name is Claude Baudoin, and I'm an advisor to the Industrial Internet Consortium in particular, not exclusively.

I like what Harley said a lot about the composition of the deliverable. I just want to say that the Industrial Internet Consortium, I see, did face exactly the same problem; and just as an indication – I'm not saying that it is *the* solution – but the way they handled it was to create a sort of super table of contents up front, which is sort of a document architecture, if you wish, for a complete book of knowledge on, in their case, the Industrial IoT, and then to name the volumes of that overall book of knowledge, giving them letters and numbers. You know, there's G-1, G-2, G-3, G-4 for the general volumes and then T-1, T-2, T-3, T-4 for technical things.

Then the deliverables can be developed asynchronously. So the reference architecture was – I don't know – G-2 or G-4, I can't remember -- et cetera. So that gives the ability for each of the contributing bodies, in this case the working groups, to go at their own pace and complete their documents when they can; but it provides the ability up front about what is going to be the overall structure of the documents. But when you see the title and you see that it's Volume G-2, it sort of hints to you as a reader that maybe you should look at the whole table of contents and find what's in G-1 and G3.

So it does require advance thinking about the architecture of the overall compilation or compendium of deliverables, but then it leaves each working group in control of its space and its scope. So that's just an idea.

ALLAN FRIEDMAN: Thank you, Claude; I appreciate that. We've got a two-finger from Harley.

HARLEY: I'm glad that Claude made that point, and I intended to make it and forgot to make it in my last comment, which was that even in standalone documents, I think that it is appropriate for each of those standalone documents to reference that it is part of a larger whole. So that should be explicit. A person reading a standalone document should not get the impression that was the entirety of the output of the multi-stakeholder process, just that they're downloading one chapter of the multi-stakeholder process; and they also have the option of downloading all of the work product.

ALLAN FRIEDMAN: Any strong thoughts – or I seem to hear a vision saying, okay, we like the idea of having the workstreams continue. We want to correlate them together, so having individually addressable documents that can be freestanding but still are clearly part of a whole. Do we like that general vision? Which probably means at some point, I'm going to be asking for folks to do that preface -- just as a heads up. I'm going to be coming and looking for folks who can sit down and actually have that preface that can be so broadly addressable, cleanly written, and describe what's in the overall community. So just letting you know that's something that's going to happen.

Now, one of the things when we talk about these outcomes is, first, NTIA I think is happy to sort of be the host, the pointer, of these documents; but there is also the alternative for other organizations to say, hey, the work you've done – either all or part – this is great. We have a long-term standing mission to do this; we're happy to take what you have and build on it.

So one of the things I always like to make sure is are there folks who sort of think that this might fit in some other organization that is happy to sort take the work that we've done, continue the spirit of multi-stakeholderism, but continue to build on it?

MEHRAN: Continue to build is okay without doing the modification to the existing one.

ALLAN FRIEDMAN: That's a very good preface -- we want to build on it, but we don't want to change and erase certain points. Any thoughts on that?

I see we have a comment from Kent on the phone.

KENT: Actually, first off, I agree with Harley; I think that we can have it both ways. We do need sort of a holistic approach to this document or this effort. The question is from the standpoint of a lot of these going forward, we need sort of an established mechanism to get it out. Personally, I think NIST is a great place to try to get some of the stuff out, if we can get them to sort of accept the work like they've done in other areas in the past.

I've been on other efforts where we have come to NIST with specifics about an effort that they've then turned into an SP. This is one area that it might be beneficial for Commerce to work together to actually get this kind of outcome out; and within there, they can have references to external documents which could be, in some respects, a full-blown kind of each working group's output.

ALLAN FRIEDMAN: That's (inaudible). Certainly for this initiative, I'd be willing to go out to Gaithersburg. It's far. Any thoughts on the sort of long-term relevance? Yes?

AMANDA: Hi, my name is Amanda – formerly OMB, but currently here as an individual. I just had a question in terms of structure. Typically in the past, if you could elaborate on what the process has been before when you reached the end of a working group and what the communication looks like. From what I've heard today, it doesn't sound like you have a communications group, per se, within the working group or a marketing person to go to. So I'm just curious what that looks like and if you could elaborate on that – if you're looking for volunteers for that too.

ALLAN FRIEDMAN: The joy of the NTIA approach is that every topic is different, and they really are led by stakeholders. The role that NTIA has played is to help work with stakeholders to achieve the impact they want. Whether it is purely awareness raising, whether it is adoption, whether it is pushing this into other more formal processes for even greater adoption, that is really up to you as a community; and we're happy to help in any way possible. We are a small agency; we don't have a massive budget. From that point, we have to rely on you, the stakeholders, to say, hey, we can do this; and, for example, we want to submit a talk to this conference where this is an audience that would be really receptive. Let's bring some stakeholders together; or let's work with this trade association, industry group, things like that.

We would love your help. We're very happy that you're here.

Craig?

CRAIG: Allan, I'm going to challenge that -- just because that's the way it's been done before; and I suggest that NTIA rethink that because, again, whether it was official recognition, whether it was the mobile, other ones, after the last working group it really didn't see the light of day. So I think to depend upon – everyone has put their time – I think that is the role of NTIA, is to then amplify that effort. And so to rethink that, I really challenge that. I know you're waiting for a new head, and a lot of things are up in the air. But that would be a role; and especially in this Administration that wants to see self-regulation succeed, I think that would be an important role for NTIA to put resource behind that, get that out there. That would be my firm recommendation to NTIA.

ALLAN FRIEDMAN: Thank you, we'll take that under very strong advisement. And I would argue that even when there are organizations that can devote lots of resources to that – for example, the NIST framework – a lot of the heavy lifting, apart from having folks from the NIST half of the world, was having really strong buy-in from private sector leaders, such as the Chamber of Commerce, who have been saying, "We can take this out and really work with this."

One of the reasons why, by the way, that we started this discussion by having an event co-located with the Consumer Technology Association is because we wanted to make sure that we had some engagement from industry from the start.

John?

JOHN BANGHART: Well, all right, I'll say what I was going to say; I thought maybe I wouldn't. I want to second what Craig said, first of all. I think, as we have all learned, and continue to learn, this is a complex issue that intersects a lot of different things. I understand that from NTIA's perspective, it can be complicated to be the champion; and maybe it's not just NTIA, maybe it is some combination of NTIA and NIST or something else. But my personal opinion is we do need some solid government leadership around this issue for some time yet.

Part of the NIST model has always been – and the reason it's been successful is that NIST has gotten pretty good at knowing when it's time to stop shepherding something and give it up. This issue is not ready to be given up; and so the extent that NTIA or NIST or whoever is going to continue to play a role, I think that role needs to be proactive. I think it *does* need to be a leadership role until it becomes clear that industry is capable of taking this and moving it on, on its own.

ALLAN FRIEDMAN: Thank you, I think that's a great segue to talk about some of the related initiatives that some of you are interested in talking a little bit about.

Michael, perhaps we can talk about the event that you have coming up that might be a good opportunity to discuss this. And I know there are a few others who are interested in mentioning this.

MICHAEL AISENBERG: I think, as I mentioned earlier, I'm the Chairman of the American Bar Association's Information Security Committee, which is part of its section on science and technology law. Coming up on May 10th and 11th here in Washington – and I've got enough of these little brochures to pass out to everybody – we have the second annual National Institute on IoT coming up. It will be held at the offices of Jones Day on Capitol Hill, a two-day event.

Among the speakers are going to be Ann Marie Buerkle, the Acting Chairwoman of the CPSC; Rajesh De, the Former General Counsel of NSA; Sean Reyes, who is the Attorney General of Utah. There's going to be a mock trial on exposing some of the issues in IoT liability, as well as a tabletop exercise looking at the various potential directions for norms and regulation development. So we expect it to be interesting. It is not free; the ABA website has all the information; and the link to that is printed on this document, which I have enough to share with everybody.

ALLAN FRIEDMAN: Thanks.

HARLEY: Michael, how do you envision this work integrating with the work you just described? It sounds like you have an agenda and speakers already all lined up.

MICHAEL AISENBERG: Yes, but certainly any attendee would be free; and there are several panels addressing issues, like norms and standards, where questions from the floor will be entertained. Those that reference this work will clearly be interesting and appropriate. I'm going to be moderating a panel on the development of norms coming from government sources with Bob Metzger, Bob Martin from the Industrial Internet Consortium, and Andras Szakal from IBM Federal looking to the future of—

ALLAN FRIEDMAN: You should think about putting someone from the Government on the panel, you know?

HARLEY: Well, I may also suggest though that if we're going to really make a push for adoption and popularization of this work, then we are going to need something beyond bringing it up in a question at the end of a panel. I mean, we've all been on panels where those things get brought up; and then no one actually downloads the documents or maybe one person does. This is sort of a panel in and of itself or a presentation in and of itself.

MICHAEL AISENBERG: I can certainly make the commitment that in the panel that I'm going to be moderating, I'll certainly make reference to this initiative. But I'm a newcomer here; many of you have sunk investment in the working groups. Plus, I would suggest that there will be great benefit of engagement and exposure with the variety of individuals who will be attending from a range of government and professional bar and practitioner sources.

Eric Hibbard, one of the U.S. Officers of Hitachi Data Systems, even though he's not an attorney, he's an officer of the ABA Internet of Things Committee; and he's going to be one of the conveners. So there will be a lot of interesting folks there to engage with.

ALLAN FRIEDMAN: Perhaps a bigger picture, which maybe we can talk about offline, is if the ABA thinks that this is a good direction to head in, how can this community as a whole engage with the ABA, which clearly has been driving in-house counsels, risk officers, things like that.

MICHAEL AISENBERG: One of the things I can make an immediate commitment to right now is that we'd be happy to do an online webinar with a panel from this community to expose the work as it currently exists to the ABA community.

HARLEY: As well as building it into the agenda of future events – I think that would be really awesome.

ALLAN FRIEDMAN: I think that would also be a great way of getting some input from lawyers both who are part of these organizations and who represent parties at risk to sort of say, hey, this is going on; let's contribute. So thank you; I appreciate that.

Jason, on the phone, we see you; we're going to get to you in a second. I just want to make sure is there anyone else who wants to talk about a sort of related forum or engagement that this body should be aware of?

Yes, Ralph?

RALPH: This is Ralph. I'd just raise the Open Connectivity Foundation, or OCF, which is very active – one of the larger consortiums around IoT currently focused around smart home or consumer focused IoT. CableLabs has been an active participant there. I sit on the Board. We chair the Security Working Group there. So we think there's merit in trying to establish standards in this area that are interoperable, broadly adopted, multiple suppliers participating.

One of the things that I would like to do is take some of the output of this work here and promote it through that organization because I think awareness of that organization of the interest from NTIA would be beneficial.

ALLAN FRIEDMAN: That would be fantastic. I think also, similarly, if there are folks in that initiative who would want to weigh in, I think one of the things we've learned is that serious security people have taken

a look at the work that's gone on and said, "Here's some input that you guys should have as part of these discussions." That would be very useful.

Craig?

CRAIG: It might be premature; obviously, the documents are still pretty rough and we're moving things around there. But I think at some future state, I could see if we really want this group to be ambassadors effectively, would also perhaps be a deliverable of a core deck or so that could be used so people could go. We all have our worlds that we go into, that we have an opportunity to present it. So having that as a resource that NTIA could create for that, that would be, I think, helpful as we look forward.

So we first have to get the core documents done, but if we think about that.

ALLAN FRIEDMAN: That would be useful, especially if one of the working groups is starting to wrap up their work, they'll have some spare time – right, Harley?

On the phone, I have Jason and then Claude.

JASON: Thanks, I wanted to sort of bring up what you had said before and then segue into the (inaudible) forum stuff really quickly. One of the things when I started getting involved in this, I was always trying to put the hat on of like how is this going to be useful to me and to us. So I've been actually trying to take the work in kind of a more agile way rather than waiting for everything to be finished.

I've been regularly feeding back what's been happening in Working Group 2 specifically to the design of the things that we're trying to do to enable IoT upgradability with really any connected device upgradability because what happened in the forum is about 12 years ago, we had to solve very, very similar use cases for home gateways and other CPEs, like set top boxes and [warp] devices and such. So the mechanisms in that protocol were kind of vetted for doing remote upgrades and keeping up with it. Scalability was a huge, huge issue that we ended up solving.

So we've taken those concepts, and we're moving it into kind of a new revolution of that protocol that is specifically targeted towards consumer electronics and revamping some of the old things. But we're working on that now, and it's called the User Services Platform; and we're pressing forward pretty quickly.

ALLAN FRIEDMAN: Thank you. Is this an open platform?

JASON: It will be; the forum is in kind of a state where it's changing the way that it does things. It traditionally develops completed specifications after it's been vetted for a long time, just by the membership and people who have to be members in order to be involved and all that. But all that is kind of changing. So there will be pieces of it that I expect that the forum will help produce that are pieces of open-source software that can be used probably for the more complicated pieces of it. Then the rest will be designed by vendors according to the specifications.

ALLAN FRIEDMAN: Excellent, thank you.

Claude, we've got you on the phone.

CLAUDE BAUDOIN: Yes, sir, so I recognize that the Industrial IoT is in some ways a different beast; but in terms of how the work that you are spearheading could flow into what we're doing and receive some promotion or some usage, the Security Framework that was published by the IIC in January or February has a short section of writing and patching. So it wasn't ignored, but it's very short; it's about half a page, which we did use in Working Group 1, and it's mentioned in the current draft catalog. But it's a great hook to leverage the work of this multi-stakeholder process and adapt it for the use of industrial stakeholders.

I can't commit to the result of that process, but I can commit to put this in front of our Security Working Group. I may be able to do that as soon as our next meeting, which is in Berlin in mid-June – so in less than two months. And then see whether the Security Working Group thinks that it could either revise its paper by adding more stuff on upgrading and patching, leveraging this work, or whether it actually wants to issue a separate report which might be a subset of the NTIA report adapted to industrial players.

So there's definitely an avenue there to connect this work with the IIC work on security and make it relevant to the Industrial Internet world.

ALLAN FRIEDMAN: Thank you, I appreciate that. You said the magic word about next meeting; but before we get there -- (inaudible) ?

MEHRAN: I just want to add to what he said. In the Security Frameworking Group, we are extending that framework document by adding different use cases and test case studies and other staff. In one of the use cases, secure update is one of the topics; so it's coming up there too. So, yes, what we are developing here can be taken there and harmonize it.

ALLAN FRIEDMAN: All right, further thoughts on community building – either people we need to bring in or places we need to go out to?

JOHN BANGHART: Well, maybe this is more generically, but there's been a group of us that have been doing some planning around carrying this message to cyber security events. If you work in the cyber

security space or you like to attend any of the major cyber security events, or even some of the less major ones, and you'd be interested in being on a panel about this topic, please let me know. There are four of us right now; we can't do all of them, but I'm happy to help share the materials that we've put together already.

Again, if you're going to be at a cyber security event in the coming several months, whatever it is, and you'd be willing to be on a panel about this subject, please let me know because I think the more people we can have sharing some of this messaging, to your point, I think the better off we'll be.

ALLAN FRIEDMAN: Thank you. I'll also add to that I've spoken at events for many of the people in this room, and I've already consumed way, way, way too much of my office's travel budget. But if there are events that you're engaged in from your community, where I can help be an advocate or at least be pathetic enough that people will say, "Oh, we'll help them out," let me know. I've *very* good at being pathetic in public, and we can actually make some project.

Mike?

MICHAEL AISENBERG: It's certainly not too early – even though it's only April, the call for speakers will probably come out in June or July – but it's probably not too early to think about RSA 2018.

ALLAN FRIEDMAN: I'm still recovering.

CRAIG: I think we actually spoke at RSA (inaudible) this year. I think if we really try and get back to our core (inaudible) that we talked about before, where are the device managers going? They're not necessarily going to RSA, as much as we all live there. Well, that's exactly; and some of us also spoke at CES. So be thinking about that engaging.

I'm not sure if anyone from CTA is here today; I don't think so, but that's a core area. So thinking outside the box, they aren't necessarily security focused companies. Maybe it's the National Retail Federation, their events, Shop.org, other areas that they go to – like an Internet retailer conference. So be thinking about those types of venues. That's to get beyond the choir.

ALLAN FRIEDMAN: NRF was on the line earlier. Yes?

RALPH: [Adir] and I are representing SMRP, and we'll be talking with you tomorrow and obviously 5,000 members that are in the world of applying these devices.

ALLAN FRIEDMAN: Can you make sure everyone in the room knows what SMRP is?

RALPH: Society of Maintenance and Reliability Professionals. It covers the world of industries. Look forward to talking with you tomorrow.

ALLAN FRIEDMAN: Fantastic. All right, this is pretty good. I think we've hit most of the items on the agenda. one of the things we want to talk about is next steps – how do we make progress? One of the things that NTIA heard from you, from the initial get-go, is frequent meetings are good; frequent travel is bad. So at the very end of January/beginning of February, we had a virtual meeting, which was a first for us. I think it worked reasonably well.

Sense of the room – would you like the next meeting to be a virtual meeting? I realize we have a biased sample here because a lot of you are based in D.C. I want to say is there anyone on the phone who wants to chime in enthusiastically, either in favor of traveling someplace or having a virtual meeting.

MALE SPEAKER: Virtual meeting.

ALLAN FRIEDMAN: Virtual meeting, says the gentleman from New York.

Harley?

HARLEY: Virtual meeting.

ALLAN FRIEDMAN: All right, so the question for those of you who are involved in the working groups is six weeks – is that a good time frame, six weeks from now? Think we'll have made some substantial progress that we can come together? Do you need more time? Can you make it happen in less time?

MALE SPEAKER: Because of some vacations, (inaudible) eight weeks kind of thing?

ALLAN FRIEDMAN: Eight weeks might put us into the middle of the summer, which is what I'm a little worried about.

(Multiple voices)

ALLAN FRIEDMAN: Okay, towards the end of June – middle to the end of June.

Craig, your light is on; but that could have been an oversight.

CRAIG: I'm just actually looking at the Consumer Technology Association's site, and we might want to look at whether it's the next meeting or some other to come up to a year – it was a year ago Austin that we pulled that together. But we should be thinking about that again, and that made some drives and milestones there.

ALLAN FRIEDMAN: And in fact that relates to the next thing that I was going to say is in terms of if we want to actually get some planning on the calendar. One of the challenges we have in the Government is it takes a little while to schedule a meeting and to get all the contracts and have this fantastic camera crew and this great space for you. It turns out it takes many, many weeks. So if we wanted to start planning for an early fall meeting, the question I have for you is can we build on an existing community -- whether it is what happened last fall, trying to co-locate with the Consumer Technology Association -- or are there other events where you think we will have a great critical mass of folks who can be involved, or at the very least where a lot of you will be?

Does anyone know that they're going to be in a room with at least one or two others of you in early fall?

Well, so the default to have an event is in Washington because it saves you, the taxpayer, a little bit of money. But on the other hand – and I know that a lot of you are here in Washington, we'll try to save people from getting on a plane. On the other hand, if people are going to be co-located or if there's a constituency we really want to get, then I think it might make sense to entertain having a meeting co-located outside of Washington.

I'm not sure if this is one of these issues where we say we absolutely need to go to Silicon Valley for legitimacy. I think this might be something that is truly diffuse enough that wherever we have a meeting, we're going to have it (inaudible). The last process we had, we did one in California; and the Washington folks complained. We did one in Washington, and the California folks complained. So we did it in Chicago, and then *everyone* complained, which at least felt more fair.

So we'll look around and try to find a date in early fall. I think for those of you who are involved in building this out, we want to have in the back of our mind that might be a great time to actually have the "this is it, this is what we have, and now how do we get it out to the wide world" type of event. It might be six weeks from now that we'll actually be close to that point, but we really want to sort of have this simultaneous mission. One, really drill down on the work we have in front of us; but, two, keep in the back of our mind the work we're going to have to do once we have documents to make sure they have the impact we want.

Any further thoughts? I'd hate to end early and give you time back on this lovely afternoon..

OPERATOR: We do have two questions on the phone.

ALLAN FRIEDMAN: Thank you for chiming in, Angela; we really appreciate that. We've got Kent and then Claude.

KENT: Okay, from my perspective, having it at the end of June is much better than the middle of June since we're already doing a presentation at first about this, which might be in the middle of conflicting. If we're going to have a meeting in the fall, we really should try to associate it with one of the targets that are outside of the cybersecurity phase. We kicked it off that way, and I think we should do that – maybe not the same group, maybe a different group. But doing that, I think, gets more people that could use it into the room.

ALLAN FRIEDMAN: Great, thank you, we'll continue to work on that. Hopefully, you as a community can sort of think about good stakeholders that we should be working with and reaching out to and stakeholder core communities.

Kent, I assume you're okay with a virtual meeting?

KENT: Yeah, absolutely, a virtual meeting is just fine.

ALLAN FRIEDMAN: All right, and Claude?

CLAUDE BAUDOIN: I wanted to say the second half of June is better than earlier for the simple reason that the IIC is meeting in Berlin the week of June 12th, so it would be good if it was right after that. The OMG is also meeting the week of June 5th, so I like that schedule, either the week of June 29th or the week of June 26th.

In terms of September, I can just offer to advocate potentially to co-locate with the Object Management Group meeting in New Orleans the week of the 25th. That may be a little bit later in September than you wanted, but there would be a lot of people involved in security of real time IoT protocols at that meeting; so that might be interesting. It's the week of the 25th to the 28th of September in New Orleans.

ALLAN FRIEDMAN: Thank you, that could be some promise. We'll have to talk about that.

I love it when we actually get such wonderful emerging consensus. Usually it's just half the group can't make the early June; half of you can't make later June. This is great; we've solved the problem. We're going for late June for the next meeting, which means that those of you who have been involved in the working groups, we have some deadlines. Those of you who haven't, this is great! We get to pick working groups, enjoin, and roll up our sleeves and really get some work done.

We're going to send out an e-mail probably tomorrow with the notes from today. They will be sort of the raw, unnetted notes that the fantastic Megan who, by the way, has been doing security for a lot longer than I have and is just fantastic and really is just helping us out immeasurably by documenting this discussion. But she actually is our cybersecurity/national security expert in the office, and we all owe her thanks.

[Applause]

So we will send out some notes. What I will do is for each working group, we'll make sure the presentations get sent out. And for each working group, we're going to have one ask that we're going to send out to all 500 people that get these e-mails, saying, "Here's the one ask for this working group, for the point of contact, so that we can continue the ball rolling, continue the conversation going."

So anyone have any last words they want to chime in before we basically say—

KENT: Question.

ALLAN FRIEDMAN: Kent, go ahead.

KENT: One of the things that we talked about today as part of this was the fact that we're looking at more of a holistic document that will cross sort of the boundaries of the four individual working groups. It might be good if we formed a maybe ad hoc, maybe official, group of folks who are looking at the work across the different boundaries trying to see how we can actually put together a deliverable and get that process going so that if we're shooting for a September or October time frame, that we have not only the individual workings groups, but we have some idea of how we're actually pulling it all together. I think the sooner we start that effort, the better we're going to be; and the easier it's going to be to get done.

ALLAN FRIEDMAN: That is a great idea, Kent; thank you so much.

Are there volunteers from this group that say, hey, this is really important; we need to have this cross-cutting view? Maybe there are some folks who aren't already track chairs who might be able to take some time and really have this broad-reaching document that's going to be easy to understand but really helps communicate what we've done.

[Pause for responses]

Matt, fantastic – Matt, start the ball rolling.

Do we have one or two people that will help out Matt?

Ralph – we've got the cable world well-represented. All right, we will solicit some help from the broader community; and I think this is also something that I'm going to ask the already overburdened working group chairs to help out. At least they'll sort of have the situational awareness of what's going on to make sure we can have forward progress.

Kent, would you like to be involved in that effort?

KENT: I sure would, yes.

ALLAN FRIEDMAN: Fantastic, thank you, I appreciate you giving that up.

Harley?

HARLEY: Did I hear you say that you had one ask for each working group?

ALLAN FRIEDMAN: Yes.

HARLEY: What is that?

ALLAN FRIEDMAN: Oh, I'm going to be asking the working groups, what do you think the ask for your working group will be? We can do this remotely, or we can do this right now.

HARLEY: I guess I don't understand. Is this an ask that the working group is making to the larger multi-stakeholder body?

ALLAN FRIEDMAN: Yes.

HARLEY: Okay, great.

ALLAN FRIEDMAN: So we don't end with just a comma; we end with a here's what we do next.

Again, thank you all. This initiative works because you show up; you get on these conference calls; you read documents; you *comment* on documents; and that is how we can have community- and industry-led solutions to complex social issues.

So thank you all very much and have a great afternoon. Thank you to those watching at home.

[Applause]

Angela, thank you very much for your help as operator.