Cybersecurity: How Far Up the Creek Are We?

October 5, 2020

>> Jonathan Zittrain: I see an odometer of participants joining us that is sky rocketing. We have crested 80.

>> James Mickens: This is great. I feel like this event is going to be wood stock for the modern era. I think people who are not going to show will tell others that they were. It's going to be generation defining. That's at least my hope for the meeting. Hopefully as a moderator you can lift us all up to reach that goal.

>> Jonathan Zittrain: I enjoy setting low expectations and barely meeting them and you're not helping with that. It is strange that Zoom appears to be adding people one at a time. What type of turn style does that represent?

>> Lis: It is very strange. The crowd seems to be slowing, so I'm going to welcome everyone to what is already the beginnings of a great conversation between Jonathan Zittrain and James Mickens. We will be talking today about Cybersecurity: How Far Up the Creek Are We? Which just seems like a great question for this moment right now.

Before we begin, we will not have a chat function going on today, but if you would like to pose a question for the question and answer phase of this conversation, please use the Q&A function in Zoom.

I'd like to welcome James and Jonathan to the conversation today. James is an associate professor of computer science at Harvard and we are excited that he is also the newest member of the board at the Berkman Klein Center, so welcome, James.

Jonathan Zittrain is the professor of the international law at Harvard, as well as many other roles that he plays at Harvard. He is the cofounder of the Berkman Klein Center and also a faculty director, and I welcome you both.

>> Jonathan Zittrain: Thank you so much, Lis. Thanks for having us today. And thanks everybody else for joining. We know you have the choice in Zooms at any given instant and we appreciate your choosing this one. James, welcome, welcome, welcome, so glad to have you both on the Berkman Klein Center board and here today to gauge just how far up the creek we are on cybersecurity.

A question, if you had been asked that ten years ago in, you know, approximately 2010, what would your answer have been then?

>> James Mickens: Well, you know, in the current era time seems to flow so much differently than it used to. It sounds like you're talking about the mid ooefl era. I would say that look, there's never been a time at which cybersecurity has been great. There's never been a time that we would say look at the landscape and say yes, we're dotting all the Is, crossing Ts, everyone gets to go home early and spend time with the family.

I think the challenge getting worse over time is that computers are becoming more ooh bibbing did yous, right, whereas before back in the old endays, even back in 2010, you know, you didn't have the things like IOT. You didn't have the pervasiveness of machine learning algorithms being used to think about who should get credit, mortgages, who should go to jail or not go to jail, whose applications to a job should be passed on to the next stage.

And so, you know, because of this increasing ooh bigity of technology it's incumbent upon us to scope this cybersecurity more broadly than it used to be. What cybersecurity sort of used to be was loosely speaking can someone hack into my system, can my data be stolen and also sort of binary yes or no way. But I think now with the technology becoming more pervasive we have to think about these bigger questions of not just can someone

tack into my canned but if they can just access my system in some seemingly benign way can they gain it in some way, influence it to create societial outcomes that are not as easy to quantify as yes, I was hacked, no, I wasn't. But we still might have huge societial impacts

>> Jonathan Zittrain: So back in the good ole days, we'd worry about data ex-filtration, having something on a platter on your machine and then somebody gets, quote, into the machine and makes off with the data and then does something with it or shares it further. We'd have so-called privilege escalation where some piece of software that's just supposed to show you a dancing hamster instead is able to get into the machine and muck with all sorts of other stuff.

And you're putting out the machine today might not just be some laptop precariously perched on a shelf but could be a refrigerator or fit bit or some kind of, I don't know, skatea system that controls whether a dam opens and closes. But it sounds like you're even going farther than that, the definition of cybersecurity itself is broad.

And for that then, would that apply, I'm coming up with examples here, to like hacking the college admissions process; is that a cybersecurity issue that's different from the old, what is it, Matthew brod Rick in war games who logged in and changed his grades?

>> James Mickens: Well, I think that, you know, sort of it speaks to this larger issue of allege rims being pervasive and computational systems being pervasive and what does that mean when potentially untrusted or, you know, participants can submit things to those algorithms, submit things to those systems. And I think that the problems are getting worse

in part because these systems that we've created to take in this data and to compute on them and then give us some answer, increasingly we don't really understand how those things work. And this has always been a problem. Like I said, this mythical time you talked about, the 2010s, even back then when we look at let's say operating systems, for example, you look at Lynn next, win doze, Mack OS,

there is no single person that understands every single line of code in those systems. Even dealing with this sort of problem of are constructions transcending our ability to understand them, that's been happening for a very long time. But I think that what's happened with technology as we become more ooh big did yous is that people, or certain segments of the population have not been as concerned about this as they should be.

And they sort of look towards computational systems as these sort of magical opaque answer boxes. They say oh, well, you know how are we going to determine how to, you know, admit students into a job or into college. Oh, we'll just use a computer to do it, you know, because that's what computers are good at doing. They take in input and output answers. But in reality, there's all this sort of underlying complexity in terms of

are these systems secure sort of in that old school cybersecurity sense, and also are these systems more secure in sort of like the new school cybersecurity sense in that are they gamable or can you influence them until ways that were, you know, not envisioned by the original creators of the algorithm.

>> Jonathan Zittrain: So the old school way of defending was some combination of trying to be extremely alert one's self, like this is a link and it looks like my utility company but don't touch it or it's all over, and having some, you know, good code to defeat the bad code; I'm running Mack Afq although it seems bonkers, so I'm running some reliable Russian thing, whatever it is, icelandic.

What's the equivalent today of doing that? What virus definitions by metaphor am I updating? How do we defend against the new generation of threat you're talking about, whether in theory or by example?

>> James Mickens: Well, you know, we've reach a very walk ward point in the conversation. I wish that I could tell you look my friend JZ just go to the app store, download this app, it's called security, it's great, 4.7 stars, people can't agree on everything. But sadly such a thing does not exist.

And, you know, I think that one reason why security sort of broadly written is increasingly more difficult to achieve is that it's not -- it's not easily definable in a sense of I just take this checklist and do these things, if I do these things, therefore my system must be secure. I think instead of trying to achieve high security, it's somewhat of a design attitude where at every level in your system design you're sort of thinking about what are the possible things that could go wrong, what are the ways this system can be influenced, and, you know, what are sort of circuit breakers that you might have in place in case something unforeseen happens. And that sounds like a very vague answer in a certain sense because it is, you know.

Their isn't a magic way to do stuff. What I frequently find, for example, as a computer science professor, is that sometimes people will want to rely purely on quote unquote testing to ensure notions of security and safety. They'll say hey, I tested my code with these 15 different test cases, surely it must be ready to push to production.

And the problem is that, you know, typically those tests they don't think adversarially, broadly speaking. They don't think about, well, here are, you know, for example, certain political goals that certain who uses your system might have and how might those goals influence how people use the system. So I think that it is really more about changing the way that we talk about the design of our technical projects.

And the same way we say, oh, there's no simple way to figure out if our system is going to be used ethically. We think about security, we have to have a similarly sort of broad attitude to say these sort of fundamental questions which are ambiguous and which have no clean answer, you know, what is security, how do I make my products secure, and so as a result, we just have to be more imagine active than we are right now

in terms of defining how we test our products for security

>> Jonathan Zittrain: So I somewhat see what you did there which was interestingly, I had asked a question without even thinking about it one way or the other, that was more about from a user perspective, what do I go get in the app store to secure my stuff and what's the equivalent of that for my fridge; do I need to buy an extra add-on so the ice maker doesn't start spitting fire.

But you were shifting to the supply side before even putting that fridge off the semably line or more systemically before cutting the ribbon on a new system at large for college admissions you need to have a more imagine active approach to security. And I don't know, then, how much, does that mean we should be licensing or otherwise scrutinizing or having some regulatory even overlay

on people producing code? Because, you know, the innives are such racing things to market you can erase the bugs later, what would in sent, if the benefit is going to have to be applied on the supply side, what's going to in sent the suppliers to worry about systemic risk that might not be traced back to them?

>> James Mickens: Those are all great questions. Sadly, I have no spiritually satisfying answers, but because I am professor, I have learned how to fill la buster my way up to the next question. So I think, like, there's one side of me which is the citizen side of me which says, yes, certainly

we need regulation to force these tech companies to quote unquote do the right thing, because, you know,

evidence suggests that the current arrow of late stage capitalism is not pointing towards these tech companies sort of doing the right thing in many cases. That being said when I look at it through the computer science or engineering part of sort of my job, I think my goodness, I get worried about what might happen if the legislation that comes out, if the regulations that come out are, you know, technically inarticulate,

if they're written in a way that doesn't understand underlying technology. You and I have talked about the GDPR which is a great example of how things can go well and poorly --

>> Jonathan Zittrain: GDPR is European and it stands for God damn privacy rules.

>> James Mickens: That's right. That's baishlly you should type into Google if you want to turn more about that, sturn the safe search off. So the GDPR is this set of rules produced by EU that among other things give users several rights that at least at face value seem like they're good, you know, the right to, for example, have your data be innumerable. You can actually go to a service writer and say, what are all the things that you have collected about me.

You get this right to be forgotten. You can go to a service provider and say, hey, all the stuff that you have that belongs to me, get rid of it. I don't want to be known by your service anymore. So at a high level, that's great. But almost immediately you start seeing all these corner cases and all these subtleties for which the idea of like what is the right thing to do, it's not clear.

So, for example, what happens when you upload some data, let's say from a fitness tracker or something like this, and then a service provider runs a machine learning model over that data and then derives some insights from that data, like here's maybe a better exercise routine you could do based on our understanding of your own unique physiological profile.

Well, when you exercise your right to be forgotten, what happens to that model? Is that yours in some sense because it was derived from your data? Well, what if it wasn't derived solely from your data, if the company did some type of met at that analysis of data belonging to a bunch of users and distill down a plan specifically for you. So the GDPR doesn't really speak to a lot of these southern of thorny issues that arise in practice.

So when you look at how companies try to comply with the GDPR, what I hear sort of off the record a lot of it is sort of prayer based because they don't understand exactly what the GDPR is asking of them and then furthermore, from the sort of hard tech side, there isn't a lot of good tech support in terms of like operating systems or things like this that would help people to comply with these laws.

So it's a bit of a mess. So I both simultaneously say, yes, I do think that we need to have more regulations but I also think the only way they are going to succeed is if we get more buy in from people actually making the tech. Of course, that's a double edged sword because when we talk about self-regulation, people say why -- in the same way people say why should we trust the oil companies to write environmental law, why should we

trust tech companies to write laws involving data privacy, for example

>> Jonathan Zittrain: Well, it calls to mind, I feel like there are basically two laws of Internet governance, that if we could just abridge them or figure them out we would be set. The first is we don't know what we want and the second is we don't trust anybody to give it to us. If we just had a better idea of what to do and trusted anybody, this is what you were just talking about, like any governmental entity to responsibly

implement that vision and align people towards it, we'd be set. In the absence of those two things, what do you foresee the trajectory here? Fast forward, suppose, are we still going to be on Zoom in ten years or is ee lon mus ck going to have put implants in or something? What's this conversation going to look like ten years from now? Would it be like, oh, we were on the right track, but at last we solved it?

Is it going to be like, no, we thought it was hard then, but oh, gosh, it's even harder now? What's the direction this is going?

>> James Mickens: Once again, another unsatisfying answer. Every direction. We're going every direction. I mean, I think that this issue that you touched upon about, you know, who do we trust, that's an issue that pervades a lot of these questions about cybersecurity. You know, for example, the debate over encryption, when should encryption be used, should back doors with put in, so on and so forth.

I think that like at a high level, encryption seems like a good thing. Why would I want someone looking at stuff that wasn't intended for them. Then you look at issues of who actually uses some of these encryption methods act, who is using TOR, what's actually being communicated with using these technologies? And any of this tension between quote, unquote regular citizens wanting to not be surveilled and also

us not wanting criminality to flourish. It's tough. It once again gets down to trust. And so I don't really know that those fundamental tensions are going to be resolved, you know, cleanly in the next ten years because I think that, you know, sort of what's ended up happening is that the rate at which some of these new technologies are being introduced is outpacing the rate at which we can understand the implications of these things.

So to a certain extent, you know, I think a lot of the current state of cybersecurity is actually driven by how quickly do new companies get formed and how willing is the stock market willing to, sort of private equity, willing to fund these companies. Because, for example, we could imagine a world which is certainly not our world, but we could imagine a world where the people with the mon kels who fund start-ups,

they basically say slow down there, young company, we're actually not going to give you money, an additional round of funding until you think deeply about, you know, how your technology might be exploited by hackers, on how your machine learning models might have bias in them, so on and so forth. We could certainly imagine such a world because it's at least describable using a human language, take English.

But, of course, we do not live in that world right now, and so one of the big problems I think with cybersecurity is that currently, you know, the funding model for a lot of start-ups is one that does not emphasize things like security, for example, and emphasizes things like user growth. And so if that's the situation that's going to continue, that's going to continue to create these I am balances in terms of what these companies prioritize, imbalances.

>> Jonathan Zittrain: Well, one, I guess I'll call it a theory rather than hypothetical, but one sensibility I have about the past 10 or even 15 years of consumer facing technology has been a movement originally from what I call owned, which is to say you're running Microsoft Windows on your laptop and then going on to compu serve if you're online and if there is a problem online, compuserve has an 800 number and you call and yell at them

and if you want to regulate them you go to Ohio and you know where to find them. And then it moves from that owned nature, where some vendor is responsible, to unowned, namely Internet, and now I'm just double clicking on stuff and downloading it on to that windows PC but it's just running and Bill gates doesn't have anything to say about what you're doing online.

And that creates this profusion of start ups and services that aren't thinking about security and you run it all anyway. And my theory had been sir ka 2005 to 2008, that that was going to create its own backlash because people were going to find their experience, oh, insecure at multiple levels that they would demand a return to the compuserves of the world so there would be some vendor responsible for being the umbrella over everything.

And then fast forward from 2005 to 2008 to today, and it feels like the world is a lot more owned, that when we're online we're spending our time on just a handful of apps that may or may not even be websites that we advise sit, and they might not have toll free numbers but they have CEOs and they are, so long as the regulators are willing, a big asterisk, regular able.

I don't know if that means how it plays into your story, but it does it mean some Dodgey start up the way it moves today is it gets bought early by FaceBook which has early radar of a start up that could compete with it in ten years so they buy it up, at which point, okay, I know whom to call if I've got a problem.

So some of the story that if a competition, antitrust story would be would you know of worry, consolidation, from a security standpoint is that actually a green chute? Is that a, well, it's not so chaotic out there as it was 10, 15 years ago?

>> James Mickens: Yeah, that's an interesting question. I mean, it's definitely true that, you know, if you empirically look at sort of like the start up landscape now, particularly in tech, yeah, a lot of these sort of young tech companies that could threaten the offshoots they get eaten up. They get bought out by these larger companies and then, you know, sometimes that's the last we hear of them. Sometimes that stuff

gets merged into the mother ship. It depends. As to whether that's good for security, though, it's not entirely clear because when you start having these sort of lij data hedge mom's, it's not natural that that in sent vices people, by people I mean companies, to sort of do the right thing. And I would also say there's this interesting aspect

to, you know, sort of the waled gardenness of the modern computational experience because on the one hand, particularly in the Apple ecosystem this is exactly what apple wants. They say okay, you buy our Apple box and what happens, you only plant app seeds in the apple box that we have blessed. If it hasn't gone through our review, met our standards of quality, you get kicked out.

In a certain sense Apple wants to live that experience there. But, you know, if you look at let's say Android, for example, you know, the Android app store is comparatively super wide open, and if you look in terms of what apps people are running, yeah, the FaceBook app is popular, yeah, maybe the New York times app is popular, but my goodness the long tail on that app store is insane.

From a security perspective, if you look at how people get hacked -- this happens to me yesterday, not the hacking part, I main my credentials. I was on dual Lynn go which is an app to teach you languages, because I'm mizerly I don't pay for the ad free version, so I saw this ad for this game and this game had clearly been designed in a week and it was basically you're trying to redirect the flow of water to make sure

that a fish gets water so it can breathe. So I'm looking at this app and I'm like that's mal wear a hundred percent of the time. I look at the app reviews on Android, half of them are clearly written by bots. It's like this game, the best it is for sure, 18 stars, out of five. You know. So that's on the app store. And that can be downloaded. So despite the fact that we could look at Android and say the security of the Android platform

in and of itself might be good, we might say the Google provided apps might have high levels of security, when you allow an open app store that's where you allow a lot of vulnerabilities there. So I think that, you know, in my opinion it's not clear that we're definitely going down this route whereby, you know, you can't side load apps, everything has to be blessed by a central authority.

At least in the Google world, in a certain extent the Microsoft world as well you can still load things on your computer that might not be good for you

>> Jonathan Zittrain: Yeah. It kind of seems like the worst of both worlds; that what most people see and are offered, unless they're not bothering to do a lingo premium, are very mainstream things from the usual suspects, and yet, they're still the story link that could creep in and you or your kid or whoever can click on it and then everything is terrible.

And in the analogue counterpart world, if we're thinking about stuff that affects human health and flourishing, there are some standards what I can buy at a super market or what's available at a hardware store and whether the light bulb I screw in is going to blow up when I flip the switch. And it does seem like we've long ago given up not even started any form of scrutiny of that sort; that we're just relying kind of on pinningtons, on commercial

vendors to serve that role. Now, maybe -- it just, even as I say it, I hear the harassy that it represents, I'm not looking for a government panel to judge every applet or extension on a prouzer, but at the same time I'm not Jonesing to give up that kind of scrutiny on product labeling or on super markets. So I'm just even trying to explore my own inconsistent tee, is it just rank status quoism?

>> James Mickens: Well, I think -- so here's a thought experiment; right? So you go to the app store and let's say that you want to buy a flashlight app. You just want an app that's going to turn on the flashlight on your phone. You're a simple man. You like simple things. So you download the flashlight app and then it says here are the permissions that this app is asking for.

One of them is the permission to turn on your flashlight, all is well in the kingdom. But another permission it asks for is the permission to look at your contact list. Seems curious, right. Why should my flashlight need to know what my grandmother's phone number S so like to us right now in this sort of clearly laid out sort of discussion, it seems obvious something is fishy there. But at a higher level,

who would prevent or decide that a flashlight app having contact permissions is wrong for some definition of wrong. There are these kind of interesting questions of scale when it comes to regulation and certification and things like that. And in part it's because, this gets back to something we were discussing earlier, you know, how do we concretely and probably automatically if you want to set the scale in the app store sense, define what it means

for somebody to be secure or for something to have too many permissions. So the reason why I think that the flashlight app example is pretty funny is because it very clearly identifies, you know, permissions that should not be given to an app, and yet, it's not entirely clear how we would sort of adjudicate such a thing. Getting back to what you were just saying, are we to have some council of trusted elders who sits around and looks at all these things and say, Zeus told us a flashlight app should only have rights to the flashlight. On the other hand, if we don't allow any type of sensing or regulation or anything like that, we get into these very clear problems

>> Jonathan Zittrain: Well, I think like I'm really wanting to take that question very seriously and it makes me start thinking that, all right, as between government and some industry, whether it's the same industry producing the stuff or some industry that springs up to do the monitoring, I can maybe see now why if it's about sending speccers to slauter houses around the world or the country, you might need a government for that

because there's a lot of physicality involved, there's a lot of economies of scale for that that you only achieve when you're doing all of them at once. And it's a common public good. So that Augusters towards government expertise, whereas here, if you're talking about an app store, maybe that's not as much the case; that the government isn't in particularly better shape to go look at the flashlight app than Apple is

or somebody else to do it. And here, at least from your example, we do know what we want. We don't want flashlight apps that can look at your contacts, there's no reason, unless it's about some obscure funding model, the only reason that flashlight is free is because it's selling grandma's phone number, which now we're just arguing with the Kato institute.

So if we know what we want, then it's just whom do we trust most to give it to us? And if it's nobody, is there some new institution or institutional relationships we could create? I mean, do we trust Wikipedia's vetting of the many contributions offered at any given moment? I don't know, trust is a big word, but we might not say we do, but my guess is we all when we're looking up something and Wikipedia is the first hit

or Serie knowledge is slurpg it right from Wikipedia, if that's going to be how many clause a crab has we're going to trust it. And similarly, I pose, for those systems running Lenox or something, there's a bunch of people supporting code to it and a council of elders right that own different tributaries of that. And I don't know if any of those examples of kind of hybrid or novel governance are scaleable, but

at least if we try to hold constant for a moment the definition of the project of cybersecurity and its boundaries, the 2010 definition such as it was, if there are enough best practices emerging, we do know what we want and then it's like, all right, do we use a free and open source software model, do we use an industry council, do we use government? We can just start to try to answer it.

Now, as we move twoordz an ever larger definition of cybersecurity where there isn't best practices anymore for these larger societally implicating systems, I find myself a little more at sea again.

>> James Mickens: Yeah. I think that's correct you know, it's tricky because I think as soon as we start looking at, for example, you know, the government's role in things like security, do we start caring about the government's role in performance, for example? Do we start looking at the government's role in accessibility? Is your server accessible to people who are blind or who can't hear, things like that.

>> Jonathan Zittrain: In America there is a government role for that, right, there is always the spec tore for those who are designing and not thinking very carefully of ADA requirements kicking in. And for performance, I guess there's at least enough government regulation that says you shouldn't lie about the performance. If you say you've got a quad core 16-thread 18-piston processer, like that had better be inside. Right?

>> James Mickens: Yeah. There are definitely sort of like sort of analogies or precedents we can draw with sort of existing technologies; although, I think that you know, a lot of the things that we're talking about with respect to, let's say, cybersecurity aren't so easily quaut fibl. So it's one thing for me to say I'm going to build an elevator, and that elevator has to be, you know, 14-X load capable such that if you overload it by some enormous amount, then

nothing bad is going to happen. But what would it mean, for example, for me to say your app must be 14-X secure in terms of like, you know, hacker resistance? So I think in part one of the problems is that some of these security metrics we have are qualitative. And even if we all kind of agree that these are some best practices, like the extent to which someone satisfies them

can sometimes be subjective. So let me give you an interesting example of this. Thing about W way. For some of this equipment people found not that there was an explicit back door that literally said, hey, communist Chinese party come in here, we left the door open for you. Instead in some of this Wa way equipment it was using outdated libraries, outdated code that were known to have some security vulnerabilities.

Now, one could interpret this observation in several ways. One could say, well, you know, Wa way just wasn't using best practices when designing this router or what not and they got unlucky and they can always change this. Another way to interpret this, which is like what many people in the American government currently believe, is that this was not sort of a mistake of coincidence; that this was done intentionally precisely

in a certain sense of laundering away the back door capability, and then the Chinese government could say no, anyone could have taken advantage of this problem. So imagine this came up in front of a litigator, you know, and let's say that you had certain different types of laws, which one of which was for sort of like negligence -- and by the way for all the people in the audience, I'm not a lawyer. When I say negligence, I mean this as a layperson would say it,

but maybe one of the laws says you're just negligent, whereas another law says you have specifically aided and abetted a foreign combatant, you know, that's a much more sort of overthe top, an aggressive charge, what would you do in this case? You know. Thick there's arguments to be made on both sides. That's why I think that sometimes looking at this from the regulatory perspective, although I believe it is necessary,

there's a lot of gray years there. Maybe it has to come to proving intent. You could speak more of this than me, but these questions are ambiguous.

>> Jonathan Zittrain: Yeah. It's a common question, it has been for years, of people just rolling in, say, to law school about why there aren't huge damages owed for building vulnerable software that then is quite predictably exploited with horrible consequence, when there is for, you know, putting bad soup on the super market shelf. And the weird answer turns out to be the happenstance, just say in American law, common law,

that purely economic damages or dignitary damages usually aren't recoverable for mere negligence; that if somebody does something they really shouldn't have down that falls below the standard of care that could hurt you physically, but it doesn't happen to hurt you physically, it only makes you deeply upset and traumatized and reasonably so, no case. It just doesn't go.

And then, of course, we teach the exceptions to it but the exceptions are rare. Now, that could always be changed. And I've always assumed that a big reason why that hasn't been changed is not on the loss side of the ledger but on the technicaling side of the ledger, not on the law side, but figuring out blame when there are so many bugs to go around. It's funny to think of Waway saying the case the catastrophic bugs are merely that.

It's not even negligence, what do you expect, it's a router, of course it's vulnerable rather than intentional. It's so common and often the mistakes are the results of multiple problems that once being exploited that we wouldn't know upon whom to pin the blame. Now, it's possible to do it and I suppose to the extent that there's an umbrella over it, like an app store, you could blame Apple for any bad apps that work their way in.

It would just have been by design the predictable consequence of having Apple switch from it's prohibited to -- to it's prohibited until it's permitted, and whether we want that and the hit in novation on that, I don't know. But it raises for both of us maybe the broader question of do we need to have some transformation in our thinking around cybersecurity for things to get any better, or

if there were a big enough check to write, would you know how to spend the money and to whom to kind of fix the problem?

>> James Mickens: Well, I think that one way to look at that question is to say, well, maybe trying to come up with sort of like a crisp and fine night enumeration of things that should be down or otherwise you're going to get sued, maybe that's sort of a fool's errand and maybe what instead we want to regulation or incentivize the use of a good process. And that's always, you know, whenever someone uses the word process like that,

distrust them immediately, unsubscribe from their mailing list. It could be, we want to see evidence that you engaged in sort of a process of war gaming what might happen if things go wrong, thinking about unintended consequences, and if you go through that process, then we will say, well, okay, bad things could still happen, but at least, you know, you were able to sort of do what we consider to be due diligence.

I think that might be an interesting model to look at. I think, though, that the constant challenge that you always bump up against, and it's not clear to me how to sort of adjudicate this, but things like regulation, in my opinion, they objectively throttle innovation. You have to jump through more hoops. You can not do things as quickly as you might want to do as an engineer.

As an engineer, I am personally fine saying I'm willing to take that hit. Like the food pills and jet packs may not be coming for five years, but we're not killing people with dangerous food pills. But that is the tension there. And there's sort of different countries, I think, will come up with different sort of ways to balance these different issues. But I think that unfortunately what's going to happen is that

there's going to be some huge disaster that's going to take place, you know, some part of the power grid is going to fail or some big chunk of hospital infrastructure will fail due to somewhat we will see in retrospect will be some preventable cybersecurity issue. And then there is some legislation that comes out better than nothing but not optimal and we'll have to refine it. So to my mind a lot of reasons

my research focus is on sort of like tools for developers that allows developers to try to make their code more secure, some definition of secure, is that I want to sort of try to make things better and give developers power to do so before that disaster happens and we have to sort of have some tragedy and look backwards and say, oh, if only we'd done this, that and the other

>> Jonathan Zittrain: Yeah. Your example of a terrible thing happening makes me wonder if a division that seemed cleaner in 2010 than it does now between industrial systems and consumer facing systems could be a division that says how much regulation there is, if it's something controlling a power grid, it's not clear, it needs to be able to, like, be tethered to the Internet at all times or be usable on Android or something; whereas, if it's

just my laptop, what's the big deal? Or if if there is enough interdependence that, no, you add up enough laptops and what is deep inside the guts of a Tesla but a laptop at the end of the day, that might be a distinction that's harder to maintain. And I don't know if if you have thoughts on that. I was also thinking we could turn to some of the questions, too, that have been rolling in from the world at large.

>> James Mickens: So maybe I'll briefly address the last question you asked and we can look at the audience-submitted questions. I definitely think toos a good idea. I definitely think that there should be differences in what, let's say, kriber physical systems for power grids have to do

in terms of regulation, versus the proverbial freak anyone gentleman or something like that. I think the scope for harms is different in both cases cln. And I think conveniently for some

of the cyber physical stuff security is more narrowly defined, which is convenient from the perspective of sort of regulatory type things. So I think that's a good idea.

>> Jonathan Zittrain: All right. Well, turning to some of our questions, one of them is whether you think that standards like NIST and ISO are a way to formalize trust and best tras for new technologies? How much do you buy the alphabet soup of organizations that have stepped forward and said, well, we'll come up with some process or some other form of label?

>> James Mickens: I'm not against them per se. I think NIST does some great work. I would say, though, that, you know, standards aren't going to save us all, though, because at a high level, many people on this call have probably heard about this concern that there's going to be this splinter net, that basically at some point China and aligned companies are going to form their own standards and define their own notions

of interopen tablt and go do their thing and -- so it's interesting to think what happens when there are competitions amongst standard bodies. Because it once again boils down to the basic question of who do you trust. So on the one hand I like the fact NIST can weigh in and say certain things about this kript toe algorithm is good or bad, but NIST has not been anointed by the gods as a single standards body. And so, you know, if we're

looking sort of at security more broadly, so, for example if we care about securing communications that travel between multiple countries that may have multiple different standards bodies, some of which are competing, then the issue becomes more subtle. But, yeah, I think getting back to JZ, something you said earlier, different standards of regulations for different settings, I think NIST style certifications or standards are particularly valuable

when things like cybersecurity can be defined in a crisp way, here's a checklist, do this, this and this and then things will be roughly fine. I think for complex stuff, how do we know machine algorithms have bias, I think NIST in sort of my reading is less qualified to comment on those things, in part because some of these questions are questions that are cross-cutting interdisciplinary. If you want to say something is like a machine learning algorithm bias, get

historians, things like that, and at least historically NIST has not been sort of -- they haven't had that wide enough set of expertise

>> Jonathan Zittrain: You want to put the hist into NIST noof that's the name of the rap album right there. I hope everybody heard it. That's the mix tape.

>> Jonathan Zittrain: All right. Exactly on that note, one of our new fellows says that in every cybersecurity training I've had to take they tell you the weakest layer of the security system is the social layer, eg the person that presses the link in the e-mail. It's the people that make

things awful, that's my editorization. In your opinion how does the social layer and the vulnerabilities associated with it change given the expanded

definition of cybersecurity that you've explored?

>> James Mickens: It's true, humans are oftentimes the weakest link, and that's sort of like one of these really dark realizations you come to. It's like you realize oh, man, bad things happen to good people and this observation here slowly follows shortly thereafter. User education, that's always a tricky thing, you know, because many problems that society faces could be solved with better user education.

I think that one of the problems with cybersecurity that we're currently seeing that is very relevant is that, you know, look at misinformation, for example, which I would put in the sort of domain of cybersecurity, even though that's not strictly speaking can I be hacked or not, can my pass words be stolen. The question of, you know, what is misinformation, who should decide what sources are trusted, this is a political question.

And so, you know, the idea of user education ends up being throrny because, oh, if you talk to a lot of political conservatives in America, they think this question is being analyzed incorrectly by Twitter, by FaceBook, so on and so forth. So maybe it's interesting to think about what types of user education are noncontroversial, here are the signs of a fishing e-mail, in terms of what parts of education are more ambiguous, like should this

particular FaceBook ad be treated as true or not. But I do think user education is part of the problem. Part of the reason it's the problem for many products security was not thought about since the first order design principal since the beginning of the project and things change rapidly and it confuses users. If you want an example of this, like a homework assignment, go look up the history of what the Google chrome browser shows when you browse an HTTPS website. This has changed several t

to an HTPS website you'll see a green lock up in the upper left-hand corner of the URL bar, this has changed at various different visualizations depending on whether Google thought they should call attention to the fact that you're on a good site, HTTPS site or should they call attention that you're on a bad site and let the steady state be so on and so forth. So I think in terms of UI issues it becomes difficult for users even well intended is --

>> Jonathan Zittrain: Yeah, it also suggests that some of these things we really do wish that the elders of science could just fix, and we don't think of it as important to living in a free society that we understand how our refrigerators work and the fact they might be wifi aware shouldn't change that. But when we talk about miss and disinformation or again, some of the broader social things you're bringing into the

rubric of cybersecurity it seems having people consider that is innately part of it. Maybe there's a way to try to fix information so that all we see is the truth, but your point that people are going to disagree about that, I don't know, it suggests somehow that if somebody in 2020 is saying, I want to go in to cybersecurity, it sounds like by your definitions what they are going into is a field that's going to be quaunta more broad than what they thought

they were going into if they were joining the field in 2010. It kind of calls for am interdisciplinary center of some kind

>> James Mickens: Right. Somehow if we could have a clearinghouse of people from a variety of different backgrounds. If only, let's brainstorm about that afterwards. I think you may be on to something, the beginning of a beautiful friendship.

>> Jonathan Zittrain: This is where people learn this is not the ad free version of the web cast because they didn't pay for it, there is a ad for the Berkman Klein Center and society. Let's see, other questions, with an increasing rate of companies signing up with the three major cloud providers for their back end or to host their website, I imagine, AWS, asher, Google, what are your thoughts with the cybersecurity issues that arise

with this heavy concentration of information centralized on three main services?

>> James Mickens: Yeah. It's a problem. I think one of the reasons why, you know, back in let's say the late '90s and early 2000s why you saw so many attacks being launched on Windows was even after Microsoft got serious about security, because they were the monopoly it made financial sense if you're an attacker to focus your malicious criminal energy on windows.

>> Jonathan Zittrain: It's reminisce sant of the will Lee sut ton quote, when asked why he robbed banks, he said because that's where the money is.

>> James Mickens: Exactly. And so it is true that in general when you have more consolidation, that's sort of -- what does that do? Well, there's a good thing is that maybe by consolidating that company gets access to more did he haves, more extensive security team and so forth. It is true, I would definite say the big tech companies, Microsoft and Google, have better security shops than smaller start ups, for example.

But it is also true the eye of soreon, tilts its gaze towards those companies. We've seen this. We've seen problems where AWS goes down and then sites that you as the end user would not associate as Amazon sites, they now disappear. They do not belong in the same material universe as you do at that point because the () went down.

Now, what you are starting to see some companies do is try to diversify across multiple platforms for two reasons at least. One reason is for security or availability reasons, so they want to say, oh, well, if Google's data center goes down, at least Amazon's will still be up with high probability. And they also do this to sort of prevent vendor lock in and have negotiation leverage. They can always go to the other data center provider and say, hey, you're cool, but I'm getting these cool, you know, signal messages from this other data center provider so give me a deal on the next contract we get. So I think that's actually a promising way to try to improve the security, by intentionally designing your distributing services such that you store data in multiple providers

>> Jonathan Zittrain: One of the long time advisors to our assembly program, which you are also an advisor to, that's PKMLA dot org, HTTPS to get there, for those watching, asks, is poor cybersecurity just like the rest of tech today, it just works poorly most of the time, dropped audio, reboot, web page does not load, order does not go through, and will just take a long time for things to normalize. Automobiles started out pretty unreliable --

I should add my own observation, thanks to the application of tort law they got a lot better after payments for featuring an ornamental spike on the steering wheel of the early pinto. But are people's standards just too low when it comes to tech. So that suggests that just wait ten years and somehow it will have figured it out the way that we have with automobiles?

>> James Mickens: I like this question a lot because I personally think that software quality has gone down over the past five years. And I think a major reason for that is because I think that a lot of companies now they have been very inspired by the model of the web; right. So back in the day, so for those of you who are old enough to remember this, back in the day there were physical stores you'd go to if you wanted software.

If you wanted the new version of Windows you would go to best buy or office depo, you get a device, small, shiney, CD, compact disc and that's how you got your code. Then once every year or once every six months, let's say, but not frequently your computer would grind to a halt while it downloaded this huge security update and that's how things used to work.

Software these days often uses this model called continuous integration, so it's the basic idea that your software is always downloading a bunch of tiny feature updates and a bunch of tiny security updates all the time. And this is very similar to what happens in the web ordinarily, where every time you go to a web page, to a first approximation, you're sort of fetching a bunch of new content and that content can change.

So when we think about what version of Amazon's web page is running right now, it almost doesn't make sense to say that. I forget if this is like ancient Greek or Roman mythology, you have this ship and one planing is changed every day. It's like at what point do you have a new ship. That's sort of roughly speaking how modern software works today.

So I think this has been a mistake. What ends up happening, I agree with the questioner, now software is much more flakey and it's hard arer to understand how it geg grates with the outside world and with itself. And in part people went towards this continuous integration model because they feel that to not do so would be to seed the race to get new features, to companies that do perform continuous integration.

But like the classic example I give is like if you go to, let's say, FaceBook dot com, that web page, it never completely works at any time. Right? I mean, you'll hit page down, it will give you the night writer thing, no more posts to show. My friends are not dead, I know they have been posting stuff. FaceBook is broken. I get a washed up on the beach that's supposed to be a photo of my friends. That's incorrect.

Why is that happening? Because they're pushing out new features. All websites are like this. So to get to the question, I think that unless companies sort of decide to deprioritize pushing out new features and sort of prioritize stability and security more, I don't personally think in ten years software is going to be better. I think some companies have managed to do this sort of rapid release strategy well, like Google chrome is a great

example of that, but I think other soft wears do not do continuous integration well

>> Jonathan Zittrain: Well, an hour has flown by. It's perhaps to be understood as a sign of our times, that it's not as if we were able to come to a whole lot of answers, but I would love to pop our tape, to use an old metaphor, into a time capsule and revisit it in ten years and see how much we're still talking about the same stuff or, gosh, how naive we were back then when the real problem was, cut to the Soprano's ending.

I'm delighted at the prospect of our being able to continue to work together through the center and elsewhere, and I gather you do take graduate students; correct? So if anybody out there is wanting to work on these things, slots available, apply now and continuously?

>> James Mickens: That's right. Just constant background radiation of applications, nothing I love better waking up, you know, 89 messages in my in box, yes.

>> Jonathan Zittrain: Wonderful. All right. Well, James, thank you so much for this conversation. And I'm looking forward to the rest of this year and to being able to conduct at some point in the not too distant future events like this actually in person rather than through a so far reliable but no doubt has its vul inabilities technological intermediary.

>> James Mickens: Yeah. Thanks for the invite. I enjoyed this conversation and had a great time. I too look forward to the year 2525 when we're back on campus and we can chat about the doom that is cybersecurity person-to-person. Thank you.

>> Jonathan Zittrain: Very good. Thank you everybody for turning up and for your questions. All right. Lis, are we done? Is that it?