

Health IT Joint Committee Collaboration

A Joint Policy and Standards Public Advisory Body on Health Information Technology
to the National Coordinator for Health IT



Health IT Joint Committee Collaboration Application Program Interface Task Force Final Transcript March 22, 2016

Presentation

Operator

All lines are now bridged.

Michelle Consolazio, MPA – Federal Advisory Committee Lead – Office of the National Coordinator for Health Information Technology

Thank you. Good morning everyone, this is Michelle Consolazio with the Office of the National Coordinator. This is a meeting of the Joint Health IT Policy and Health IT Standards Committee's API Task Force. This is a public call and there will be time for public comment at the end of today's call. As a reminder, please state your name before speaking as this meeting is being transcribed and recorded. I'll now take the role. Josh Mandel?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

I'm here. Hi, Michelle.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hi, Josh. Meg Marshall?

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

I'm here.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hey, Meg. Aaron Miri?

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Good morning. Hey, Michelle.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hey, Aaron. Aaron Seib? David Yak?

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

I'm here.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hi, David. Drew Shiller?

Drew Schiller – Chief Technology Officer & Co-Founder – Validic

Here. Thank you. Sorry.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hi, Drew. That's okay. Ivor Horn is not able to attend. Leslie Kelly Hall? I know Leslie's on. Linda Sanches?

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Yeah, I'm here. Sorry.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

That's okay; thanks, Leslie. Rajiv Kumar? Richard Loomis?

Richard Loomis, MD, CPC – Vice President, Chief Medical Officer – Practice Fusion

Good morning.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hi, Richard. Robert Jarrin? And from ONC do we have...we have Jeremy, Lucia and Rose-Marie Nsahlai I believe?

Rose-Marie Nsahlai – Office of the Chief Privacy Officer – Office of the National Coordinator for Health Information Technology

Correct.

Lucia C. Savage, JD – Chief Privacy Officer – Office of the National Coordinator for Health Information Technology

Present.

Jeremy Maxwell, PhD – IT Security Specialist – Office of the National Coordinator for Health Information Technology

That's right.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Anyone else from ONC on the line? Okay, with that I'm going to turn it over to Meg and Josh.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

Hi, thanks...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Well thanks very...

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

...yup, go ahead Josh.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

I was going to say thank you for the...getting us kicked off and we've got a pretty packed agenda for today and so just very briefly, to lay out the plan. We are going to hear from Justin Richer about OAuth, because it was an area where we thought the workgroup could use some background to make sure that we're using shared vocabulary that aligned with the standards community and that we were able to ask questions of Justin about the specifications and about how OAuth works.

And so we've budgeted 30 minutes for that presentation; Justin will show us some slides and give us some background for 20 minutes and then we'll have 10 minutes to ask questions, and I know it's going to be compressed because we could probably have this kind of discussion for a whole afternoon, but unfortunately that's the time that we've really got available. And then we've got an hour set aside to begin talking through the recommendations that we've begun to form; and thank you all for your help in drafting a set of specific tasks and recommendations inside of the Google doc.

Let me turn over to Meg to see if she's got any other introductory remarks before we get kicked off.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

No Josh; that covers it, thank you.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Excellent. So with that, I think we'll turn over to Justin and we'll note the time and plan for Justin to talk for 20 minutes and we should close up by 11:05 Eastern time.

Justin Richer, MS – Independent Contractor, Founder - Bespoke Engineering, LLC.

All right, sounds good. Thank you, Josh and Megan for bringing me in today and, you know would it be possible to full screen those slides instead of the smooth scrolling? That might be helpful in the viewing for whoever is going through this.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Uh, this is Michelle. Because we had the update at the last minute we can't do that, but if on your screen there's a little arrow thing that looks like four arrows. If you click that, you can see full screen of the slides, it's at the...on the top.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

Right. So the question was not a matter of it being large enough for me to view, it's more a matter of they're being only one slide on the screen at a time. So if whatever viewer you're using, if you could turn off continuous scrolling, that might help or single slide view, depending on which PDF viewer you have; that might actually help. But, anyway, we can move on with this regardless and I'll just say next slide as we need to scroll through these. Sounds good.

All right. So, good morning everybody and thanks for having me out here. I'm here today to talk about OAuth 2.0 and first wanted to say that because this is the API Task Force, I really don't need to introduce you guys to the idea of APIs, at least I should hope not. If I do, you should definitely have more speakers to come talk about that. But much of my data, much or your data and the sort of the functionality of our everyday life these days is available through APIs.

All kinds of different services and functions are publishing things as APIs on the web today and it's not enough for those APIs to just simply exist out there, excuse me, I want my application to be able to access those APIs because that data sitting off in a silo somewhere doesn't actually do me any good. I want stuff to be able to access that data and those functions and do so on my behalf, but I don't want those applications to have to impersonate me in accessing those systems, those APIs. I don't want to have to share the keys that I would use to access that data with applications that I want to access things on my behalf. Next slide, please.

So the idea here is what if we had a special kind of key, a valet key for the web? Now if you're not familiar with a valet key, and I realize that not a lot of cars come with these, especially these days with electronic keys and whatnot, a valet key is a special kind of key for a car that gives limited access to the car. A lot of times that access is only to start the engine; it can't open the glove box or the trunk, for example or it can only unlock, you know one door. Some of the more fancy ones will even limit the top speed of the car or not let it drive a certain distance away from the place where it started, if you get some of the really fancy ones out there.

But the whole idea is that it is a different key from the one that I use in order to drive the car and it gives whoever I hand that valet key to, a limited set of access to my car. So what we want to do is have that type of access but for web APIs. Next slide.

And this is where OAuth 2.0 comes into play. Now what is OAuth 2.0? You guys I'm sure have heard of it by this point. Next slide, please. And from the specification that defines OAuth 2.0, which is RFC6749, you can see that it's a large and often inscrutable, impenetrable block of text. Next slide, please.

We're going to go over just the good bits of this, which is that it is a protocol that allows a third party application to get limited access to an HTTP service, an API, on behalf of a resource owner. In other words, next slide, please; OAuth is really a delegation protocol. It lets people allow applications to access things, such as APIs, on their behalf and it's all about delegating that right of access. So next slide.

Who's involved with this process? The OAuth ecosystem has four different major players. The first of these is the resource owner, and we'll also call them the end user at some point. There's also the client application and the protected resource, and we'll see how these connect together in just a few minutes. And finally we will introduce the authorization server and how that brings everything together.

So first off, next slide. The resource owner; this is the entity in the system that has access to some resource, or API, and they can importantly delegate access to that resource or API. Now this is usually a person, they usually have access to a web browser, that's kind of the canonical...case. And there are extensions and adaptations of OAuth that bend that a little bit, but what we're going to talk about today and what you see most often in the wild is a user in a browser acting as the resource owner. Now it's important that we keep in mind throughout all of these that all of these slides that these are roles that different entities can take within the system. And that means that each entity could possibly be playing multiple roles in a real OAuth deployment.

But for the purposes of this discussion and for the purposes of defining the specification, we're going to keep all of these roles separated because in truth, you can actually deploy an OAuth system with all of the different pieces separated. Thus our resource owner we're going to say is the person that has access to a web browser and they want some piece of software to accept the client.

Now I will also say as a side note here, the resource owner is not necessarily the data subject of the API that's being accessed. The resource owner is simply a person who has a right of access to that API; it doesn't mean that the data is about them, though often it is and it can be they simply are ones that have access and have the ability to delegate that access. Next slide, please.

And what they're delegating access to is known as the protected resource. This is the API and it's not...just not that it's an API; it's an API with a specific set of security controls. This is something that not only stores things and provides functionality for the resource owner; it also shares that functionality on the resource owner's behalf.

So it's not enough for this to just be a data silo somewhere that allows anybody to access stuff, and nor is it enough for it to be a system that allows nobody to access stuff; it's just, you know, it's kind of a data storage locker sitting somewhere. It really does have to have this aspect of elective sharing and it's at the request of and under the direction of the resource owner. We'll see how that works in just a few minutes. Next slide, please.

The client application is the thing that wants to access the API. This is a software application doing something on the resource owner's behalf, and it can do any number of things, whatever the API actually allows. Now this is where the term client tends to trip people up because this client application could actually be in a web server and that is, in fact, the original and canonical case of OAuth where the client application is running on a web server somewhere and its accessing an API, which is of course a web service running on another web server somewhere else.

It is a client in OAuth parlance because it is a client of the API even though it itself is running on a server. It could also be a native application running on a mobile device or a desktop or a JavaScript application running inside the browser. Each of those has different deployment patterns but they're all known as OAuth clients because these are the things that are accessing the API on the resource owner's behalf. Next slide, please.

And that's fundamentally what we're trying to solve; we want to give the client access to the protected resource on behalf of the resource owner. Now in the past there have been a lot of different ways that we solved this, lots of kind of simple ways that people have done and the most popular of which, which as a security architect is kind of terrifying, is to give...is to just have the resource owner hand over their username and password to the client and the client will just replay that username and password over to the protected resource, thereby impersonating the resource owner enough to get access to the stuff at the API.

And now this is terrible because it leaks the password, it leaks all of these security credentials to the client so now the client can do whatever it wants to as that user. And the client can actually even send that username and password to somebody else to have them login as the user, because the client has access to it. And all sorts of terrible things can happen when you let that happen.

Furthermore, in order to get rid of the client's access, what you need to do is change your password. Changing a password is a really, really onerous thing for an end user to be able to do. The resource owners really don't like to do that and so...but the thing is that's the only way to revoke access to a client once you've shared your password with it.

What we're going to try to do instead, next slide, is introduce a new component that allows us to do the OAuth process, and that is the authorization server. This gives us a mechanism that lets us bridge that gap, the bottom of this diagram, between the client and the protected resource in a way that doesn't expose the resource owner's credentials. Next slide, please.

The authorization server is the thing that generates tokens for the client. It also is responsible for authenticating the resource owners, authenticating the clients and managing the authorizations of the resource owners to the client. So when the resource owner says, "Yes, this client can access this stuff," that's all handled by the authorization server. Next slide, please.

And that is all represented by an OAuth token. Now an OAuth token is simply the end result of the OAuth delegation process and it represents a certain set of authorities that have been granted to the client by the resource owner in order to access the protected resource. Tokens are issued by the authorization server. They are used by the client, but the client doesn't actually know what's in the token, nor does it care to know what's in the token.

So if you're familiar with WS-Trust or SAML or Kerberos or one of those other protocols where the client actually has to process the tokens going through the system, OAuth doesn't work like that. OAuth very explicitly uses opaque tokens, as far as the client is concerned, in order to carry all of this information across the network. But fundamentally it is consumed by the protected resource, which takes a look at the token and decides whether or not the token, as presented by the client, is good enough for the request that's being made. Next slide, please.

OAuth token can look lots of different ways. So, this is something that I found confuses people stepping into this, especially if you were coming from a SAML or WS-Trust background or even a Kerberos background and you're like, so all right, what is an OAuth token? Well it could be a hex blob. It could be a UUID. It could be this crazy structure in the middle, which is called a signed JSON Web Token or JWT. It could be a high entropy word list; it could be all kinds of different things.

The thing is, OAuth as a protocol doesn't actually care. The whole idea with an OAuth token is that it's representational of that transaction of the client being delegated access by the resource owner in order to get to the protected resource. And we'll take a look at how exactly that process works, in just a little bit. So next slide, please.

The OAuth process in a nutshell has the user and the client authenticating for themselves separately, and the user authorizing...the resource owner authorizing the client to act on that user's behalf. The server, the authorization server then generates a token that represents that authorization and hands that to the client. And the client then presents that token to the protected resource, the API, in order to gain access. This might all seem fairly complicated but it's actually a relatively straightforward, very deterministic kind of process.

And next slide, please. The truth is that you've actually used OAuth. If you've ever seen a screen pop up that says, "Hi, some application is requesting access to the following kinds of things, do you approve this?" Chances are you are actually using OAuth; it's been in use for quite a number of years.

OAuth 1.0 was invented around 2008, OAuth 2.0 was ratified in 2012 and it is in wide use across the Internet and across a wide variety of different systems and services. And if you've also, I will add, ever used a Facebook application, those all use OAuth for their authorization in getting your Facebook information. If you've ever logged into Google, you're actually using OAuth under the hood.

If you've ever used the Steam or Spotify desktop applications, those actually use OAuth to connect into their backend services. And finally, if you have ever used a Smartphone to do nearly anything, I will almost guarantee you that you have used OAuth in that process as well because it is embedded in lots and lots of different services because it works very, very well in today's web API environment. Next slide, please.

We're going to be talking about the different bits of the OAuth system now and we're going to talk about one way that you can connect them all together. Now OAuth as an authorization framework is actually a family of protocols, but we're going to talk about the most canonical, kind of the most core version of this that separates all of the different pieces. Everything else in the OAuth 2.0 system is kind of an optimization of the process that we're going to be talking about today.

And I will say, next slide, please, we're going to be getting into some technical guts here, but I'm going to try my best to keep it very kind of high up and I definitely have plenty of time at the end for lots of questions and discussion. So we're going to be talking about the authorization code flow. This is really the canonical OAuth transaction and this is where OAuth was really kind of built, where it grew up.

Like I said, there are other ways to do OAuth and those are available at the end of the slide deck, if you guys like, to download and check those out, but those are all really simplifications and optimizations of this flow for very, very specific purposes. Next slide, please. The authorization code flow connects all of the different pieces in a very specific order and in a very specific way. So, everybody's got this diagram? Great; next slide, please.

We are going to talk about, don't worry, I'm going to get into detail, I can hear the panic coming over the phone. Umm, before we get into the detail of how everything is connected in that diagram with all the lines and arrows and whatnot, we need to talk about two different ways that these commun...that these pieces can communicate to each other. Next slide, please.

The first of these is known as the front channel, and this is everything that goes through the browser, basically. Now this is something that uses HTTP redirects in order to communicate between two systems that are not directly talking to each other. So, this is a really interesting means of communication because it means that the resource owner and their browser can actually be involved in the process here without having to share anything to components and get them to talk directly to each other.

So the resource owner is sitting in their browser and they can get a redirect, an HTTP redirect from for example, the client, except that the URL of that redirect will actually be a URL that is served by the authorization server. And the URL that is sent in that redirect will have a set of query parameters on it that are used by the client to communicate certain pieces of information to the authorization server.

The important thing here is that as the client and authorization server are on separate domains, they are in different origin policies for the browser meaning that their sessions, that their information, that all of that type of stuff that browsers do every day to separate sites from each other, that all comes into play and works for you. The authorization server can also send messages back to the client by sending a redirect to a URL that is hosted by the client. And the client can likewise serve that URL and pull the parameters off of the query on the way in.

All of this happens without the sessions between the resource owner and client or the resource owner and authorization server crossing. Never went across the street; ghostbusters thought of that a long time ago, and it's still true today. Next slide, please.

This isn't sufficient, though, for really securely communicating between everything in most cases because as you can imagine, lots of stuff leaks through the browser there. So OAuth also makes use of the back channel, which is a direct communication, a direct HTTP connection between the client and the authorization server, the client and the protected resource. And this is, if you've ever done HTTP programming, this is kind of the more traditional HTTP stack. You make an HTTP request, whether it's a get or a post or a put or what have you, and then you get back an HTTP response and that response will have a...object or an XML document or really whatever needs to come back from that response.

OAuth makes use of both the front channel and the back channel at different times and in different ways in order to communicate different pieces of information, and importantly, to keep the information separate on all of these different transactions. And that separation is the basis for the security model of the OAuth protocol, which is why it's important when building an OAuth system that you use the right OAuth grant type with the right use of the front and back channels and all of these different pieces that we're talking about, when you're building out your application. So, next slide please.

We are now going to take, as promised, a step-by-step, high level walk through the authorization code flow of OAuth. Next slide. Step one, the client says to the resource owner, "Hey, I would really love to be able to get that protected resource for you but I can't get do it on my own. I need you to go to the authorization server and come back with something for me that I can actually do something with, that I can actually move this process forward, in order to eventually get stuff for you at the protected resource."

Now remember this takes the place of the client saying, "Hi resource owner, I would really love to be able to access that protected resource, please tell me your password for the protected resource and I promise I'll be nice and I'll go and do it for you." Instead the client says, "Hey, I don't want your password, I...you may not even have a password, I don't know and I don't care; I'm going to send you over to the authorization server who I know can eventually get me the stuff that I need." So we're going to start that off and this...remember it's going through the browser, this was part of the front channel. Next slide, please.

So the resource owner, in their browser at the authorization server, can then authenticate to the authorization server using whatever password they want to with the authorization server or it could be certificates, could be any number of things. The important thing here is that the resource owner authenticates to the authorization server, but not to the client. The password or whatever they're using to log into the authorization server never connects to the client, it is only in the session between the browser and the authorization server so that the authorization server knows who's there.

Because of the request that came across, next slide, please, the authorization server also knows which client is asking for stuff and what they're asking for. So what generally happens in most OAuth systems is that the resource owner gets prompted and asked, "Hi, this client wants to go get stuff for you, is that cool? Is that what you want? If so, click okay, if not click cancel;" that kind of thing happens at this stage.

And so at this point the author...the resource owner can say, "Yes, that's...I was just at that client, I want them to go get my stuff or maybe I don't want to give them access to less stuff than they asked for and they'll be mad at me, but that's fine." They can...this can play out in lots of different ways; the important thing is that the resource owner has a potential opportunity to make a decision here. Next slide, please.

Now the reason I couched that and say it's a potential opportunity to make that decision is that that decision can actually also be made by a centralized policy or a sort of traditional enterprise controls and those can fall into what I like to call the...a three-tiered layer model, a three-layered model of a whitelist, greylist and blacklist.

Now the whitelist is your traditional enterprise security federation type stuff where you say, these are all the people that we know about. This is the stuff that we, by policy, by contract somewhere, have in a dark back room decided is okay; so we're going to do all of that. And we're not actually even going to ask the end users anything, it's just going to go through.

We also keep in that same central control a blacklist that says; this is stuff that is never going to happen. Importantly in the middle we have the opportunity in OAuth to have the opportunity in OAuth to have greylist, which is the entities that we don't know about and we're going to actually prompt the user to make a decision for us and use a mechanism known as "trust on first use," and use is supposed to only have one "U" in it, I missed that typo.

It's trust on first use or TOFU, and you can think of this type of decision like the kind of markings that you see on a fire door where you're giving the user information in the context that they're about to make an actual decision. So I walk up to a door and I need to know, am I able to walk up to this door, do I actually want to walk through this door given my current situation.

If I walk up to a door and it has big red letters and it says, this is a fire door, alarm will sound, I will ask myself, "Is there currently a fire and do I want that alarm to sound?" If the answer is yes, then I'm going to go through that door. If the answer is no, then I am not going to go through that door. And you're giving the user, the resource owner that same type of control here. And this is, of course, not in lieu of, but instead in addition to the traditional whitelist and blacklist models.

So, we made an authorization decision, next step, please; next slide. Thank you. And now the authorization server mints a temporary credential, known as the authorization code, and sends it back through the browser to the client, again using the front channel. Now this is a one-time use token, or a ticket basically...next slide, please, that the client then sends in the back channel to the authorization server. And it's important here that the client can also send its own credentials, it can authenticate for itself at the authorization server. It's not authenticating for the resource owner, doesn't have the resource owner's credentials, it just has this authorization code and it has its own client credentials.

So the authorization server can take a look at this, look at the authorization code and the client credentials and say, "You know what, yeah this is the client who asked for it, I know the user that approved this, I know what it was approved for." Next slide, please. I can now take that information and

mint that down into an OAuth token. That token represents the entirety of that delegation process that just happened.

All of those decisions that the resource owner made back when they were actually here at the authorization server, those are referenceable from the authorization code. So I can then take that and carry it forward with the token that I then mint and hand to the client, which the client can then...next slide; use to call the protected resource to go in and do stuff for the resource server. They are accessing the resource using the access token, and the client importantly doesn't actually know what's in the token. Next slide, please.

It's completely opaque; the client doesn't know or care what's in the token itself. They may never know who the user is, they just know that somebody approved it and this process worked and I got a token and the token works. The client doesn't actually get to know any additional information about who's there, you need other systems in order to communicate that to the client if you want to, but that's talk for another time.

The resource server, the protected resource does need to understand the token and it needs to fundamentally know who the token was issued for, the resource owner, which resources it's good for and what types of actions it's good for. Because it needs to be able to map those rights of access to the request that's being made by the client using that token, and it is fundamentally up to the resource server to decide whether or not that token is good enough and sufficient for the request that is being made.

And that in a nutshell is OAuth 2.0. Next slide, please. Thank you very much. There are lots of other slides in this deck, backup material that talks about different ways to do OAuth and how to build an authentication, a login system using OAuth, because OAuth, as we saw, doesn't actually tell you anything about the user, the client, doesn't actually ever know anything about the user, but there's some information in the back of this deck for those that want to pursue it more. And with that, I will take questions; thank you.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Hey Josh and Meg this is Michelle; do you want to use the hand-raising feature or do you just want to let people ask questions? It looks like Aaron has his hand raised, which is why I'm asking.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Umm, that's entirely up to you how you want to run this meeting.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Well Aaron took his hand down, so...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, let's at least give Aaron a chance and so let's, I think we weren't actually that clear about the timing but let's take up to 10 minutes from now for questions and we'll start with Aaron and if folks want to raise their hands, go for it.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Justin that was an amazingly helpful and very valuable presentation, thank you for your time today; I had a couple of questions. There were two points during the presentation where you talked about what I think might be best practices that we should be recommending and I just wanted to get your thoughts on specifically the charter of this task force with regards to best practices around the OAuth tokens being shared. You had a slide that had a couple of examples; is there any preference to JWTs or other tokens over multiple entity strings and so forth, from a security perspective?

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

No, not really. So, not from a security perspective as long as it's a high enough entropy that the tokens cannot be guessed or generated by a malicious party.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Okay.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

Because OAuth tokens, at the least the bearer tokens that we've talked about today, which is far and away the most common use of OAuth...

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Right.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

...are effectively a shared secret. They're a limited shared secret and there are provisions for them to be automatically rotated and expired and all sorts of other things like that, but you've really got to treat them as passwords. However, since they're not a password that any person needs to type, they can have a whole lot of entropy crammed into them in a variety of different ways, because they're just being replayed between systems.

That said there are some benefits to using a signed JWT in that you can carry some information in the token itself which may be helpful downstream at the resource server as it tries to figure out the JWT. So, for example, the approach that we're taking in the OpenID Foundation's HEART Working Group, the Health Relationship Trust working group is that we are actually mandating that all tokens be signed JWTs, and specifically asymmetrically signed JWTs such that any resource server that gets...any protected resource that gets added to a HEART compliant authorization server will be able to read-in the token and know which server issued it and be able to check its expiration and signature directly.

If you actually add in a separate protocol known as token introspection; that gives a mechanism for that protected resource to call back to the authorization server and say, "Hi, I just got this token; I need somebody to tell me what it's good for." And you can do that without packing information into the token itself.

So again, what we're doing is we're actually prescribing a use of both of those together such that there's kind of a first order of validity information baked into the token itself using the JWT signature format, but then in order to figure out, you know, what client it was issued to and what scopes it has and all of that kind of stuff, you need to call back to the authorization server using token introspection to get that information.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Is there any...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

(Indiscernible)

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Is there any infrastructure needed to support that token introspection or is that something that an unknown, you know, client can employ with a resource owner, any resource owner?

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

So the token introspection involves neither the client nor the resource owner.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Okay, good.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

If we can actually go back to a...let's just go back two slides, please, I just want a diagram to point to; thank you. So, introspection is really about the connection between the protected resource and the authorization server, it's once that token comes across the wire, it's kind of a, you know, one way to do step eight here, is the protected resource needs to know what the token's good for.

Token introspection is a web protocol that allows the protected resource to call the authorization server and ask, "What is this token good for?" And it can do that very quickly and it's a very, very simple protocol. You posted the token...

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Okay.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

...like a JSON object; so this is a service that's provided by the authorization server for use by the protected resources. The resource owner and client are not involved in token introspection. The reason that it's one way to do step eight in this process is that if the authorization server and protected resource are very tightly bundled, you can actually skip that and just do a database lookup, you know, when the authorization server creates the token, it stuffs it into a database along with a row of information about, you know who its issued for, what scope it's good for and all of that other kind of stuff.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Mm-hmm.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

If the protected resource is bundled very tightly with the authorization server, it can just go look it up in that same database. Umm...

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Got it.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

...or there's lots of different things you can do.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

And my other question, if we could go to slide 28, I was just...it wasn't clear to me how the client, umm, maybe slide 26; I'm sorry, two more back. On here you show the client redirects the resource owner to the authorization server. How does the client...so the client...how does the client know which authorization server to redirect them to?

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

That's outside the scope of OAuth and its assumed to be a tight configuration coupling between the protected resource that the client is trying to access and the associated authorization server. Now the traditional OAuth world where OAuth was really born, there really was no ambiguity here because there is...you're probably talking to some proprietary, you know sort of single instance API like the Facebook API...

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Right.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

...so to talk to the Facebook API, you have to talk to the Facebook authorization server. Same goes for Twitter and Google and all of those others.

Now where things start to get interesting is when you can have the same API, such as OpenID Connect for identity or FHIR for medical records or any number of things, available on lots of different protected resources that may be protected, and likely will be protected by different authorization servers. So the means by which the client makes that binding between the authorization server and the protected resource is out of scope for the core of OAuth and is really kind of a discovery process.

OpenID Connect and User Managed Access, which are protocols that are built on top of OAuth, both have discovery systems that allow you to do that and sort of bootstrap that process. And of course you can always just code it into the client itself and say, you know your token end point is here, your authorization end point is here and your protected resource is over here.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Go find it here, yeah.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

Yup, exactly.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

And just one last kind of brief question, as far as the...you went to some, you know talked a little bit about front channel, back channel and being careful about what you do through the front channel. Is it helpful to the community if this task force makes recommendations about what should never be done in the front channel? What would your recommendation be for...like ours?

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

So for this task force I think it would actually be more helpful to have very concrete guidance on the types of applications that people are deploying, both in terms of sort of the clients and the, you know the nature of the resources, and which flavor of OAuth to use in those situations. And you see this type of recommendation from both the SMART on FHIR Project, which Josh runs, as well as the HEART Working Group.

And this type of recommendation actually goes all the way back to the Blue Button Plus REST and RHeX Projects from quite a number of years ago where we basically say, if you have an in-browser JavaScript application then you have to use OAuth in this particular fashion. If you instead have a web server-based application, then you have to use OAuth in this other way, which is different from that. If you have a native application, you have to use OAuth in this other way. If you have an application that is not acting on behalf of an end user, you have to do this other thing. And so it's that type of recommendation, it's the, "if you're trying to A, then you must do "X" type of formula..."

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Right.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

...that I think is most useful, as opposed to blanket restrictions on, you know never send this over this channel or what have you.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

That's helpful, thank you. It would be...Michelle, it would be helpful if we get somebody from the staff, the support team to pull together some of those things that Justin just referenced that have been done in the past or, you know kind of help us produce the type of output that he's suggesting.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Sure, I'll defer that to Rose-Marie Nsahlai. Umm, Aaron, if you're done with your questions, Leslie Kelly Hall has a question.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

I am totally done.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Okay, thanks.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Hi, thanks so much for that. And this is a follow on to Aaron's question but, and may be for Josh. So as the recommendations stand under the Blue Button specifications in the HIT Certification, does that meet the full and complete need for this task force with regard to OAuth recommendations or do we have further clarity or recommendations that need to be provided?

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

I think I'll let Josh field that one.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, that's on my side the...as the regulations stand, it just says that you have to offer an API and it has to be secure and you might think about using OAuth as the way that you secure it; that's about what the regulations say right now. And frankly I don't know that this workgroup is going to actually recommend a lot of specifics about the security protocols.

I think our purview is to be able to highlight examples where these protocols are used successfully and securely. I don't think that we should be in the business of trying to decide things like what goes in the front channel and the back channel? Which profiles do we use and what token format? I mean, I think that just goes well beyond our scope and our expertise and it would be hard for us to get that stuff right.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Yeah, I agree with that, I was just wondering if, what Justin what you mentioned, is there existing specifications in the standards under the...Blue Button recommendations under certified HIT. So there isn't anything beyond then what you just mentioned Josh in that specification?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

That's right.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Thank you.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Sure. So I know there may be other questions for Justin and maybe what we can do is if there are, please share them with us over e-mail and we'll make sure that we share them back with him. At this point we've got a few more topics we want to get through in the rest of today's call so, thank you very much, Justin for joining us. I really appreciate the overview and I'm sure the other members do as well and I wouldn't be surprised if we get back to you with a few more questions after people have a little bit of time to digest.

Justin Richer, MS – Independent Contractor, Founder – Bespoke Engineering, LLC.

Absolutely, thank you for having me out and glad it was helpful.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Excellent. So what we'd like to do at this point is move on to some of our draft recommendations, and Meg and I have highlighted a few topics from the document where we've been drafting them, that we wanted to...that we thought would be especially valuable to discuss in person. So I think we picked out four topics, between the two of us, that we wanted to highlight. Justin just sent me a note, Justin you are welcome to stick on if you would like, but at this point I think it's going to be sort of task force business.

So in terms of the four issues that we wanted to discuss; one of them is an issue that Meg had brought up which was a question about blacklisting, and this actually ties back nicely to some comments that Justin made; Meg, is it fair for me to ask you to introduce this topic, and at the same time on the screen share, if it's possible for us to bring up the set of recommendations underneath the topic of limitations and safeguards on sharing? It's the one place in the document that the word blacklisting is mentioned; that would be helpful to pull up this as well.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

Yeah and so thanks Josh and while we're moving to get that forward; we've talked in the workgroup about the need to address potential security threats by the API developer. And then kind of anecdotally, this blacklist concept, and we definitely just heard a little bit more. The greylist was I don't think we discussed that in the group yet, but certainly we're talking about the blacklist which says, you know the API technology has identified a security threat with the App requesting access and will blacklist or refuse to share data with that App moving forward.

And so just wanted to talk a little bit about the text as we have it now in the recommendation, there is a note in here that just, so like, let me rephrase, there are actually two things that I wanted to raise with that. The first one is do we want to provide some clarification around when it's appropriate for an App to be blacklisted?

And just to kind of carry that thought forward a little bit, thinking about, you know that there actually has to be an action and it can't necessarily be something that the vendor just doesn't like about the App, you know, if the App meets all the right technical specifications and hasn't introduced any behavior specific to a threat to the system, then it should not be okay for the App to be blacklisted. So looking for a little clarification there; but then the second part of that is, when the App is blacklisted, the notification to the consumer that the API is refusing to connect with that App, what options are available to the consumer? And I think there is one suggested here that the consumer should always have the ability to override that.

So just wanted to have a little bit of discussion around, okay, if you're a consumer and you want, you really, really want to connect with this blacklisted App, what are your alternatives? What are your options available to you? Could you appeal this? Could you broker an appeal on behalf of the App, you know, what are the right clarifications and recommendations that we could make here? So, to Josh's point that was, you know, what I had wanted to talk through with the group a little bit.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

This is Leslie; just a clarification. So is this App blacklisted because of the way it behaves on its own and uses data or is this App blacklisted because of the way it can potentially connect to the API and introduce malicious behavior there? It seems to me that in the prior case, if it's just a bad actor, that's buyer beware, the consumer should have to override it. If it's something that actually is causing harm to the API and the connection to PHI, the provider should not have to be required to whitelist it again?

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

This is Yak; could you move to the part of the document that mentions the blacklisting and display the comments that have already been put there?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, if you search...

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

I think its page 16.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

...this document for the word blacklist, just...

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Yeah, it's page 16 under recommendations.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

So Leslie, this is Meg; to your point. I think that's a really good clarification around the two scenarios that you just outlined there, you know, the bad actor, how much responsibility should we place on the provider or the API technology to understand when an App is a bad actor? You know, certainly I think the second case is very clear, if they're introducing a threat and it's a current threat, an existing threat, that there's behavior that can be associated with that, absolutely I think that's very clearly within the blacklist realm.

But if you're trying to, you know, reconcile an App that is...purports to have a certain policy but then its actual behavior doesn't follow that, how much responsibility does the provider or the API technology have as far as trying to vet that out and then place it on a blacklist. So again, I'm not necessarily, you know, offering a proposal here one way or the other, I just, for the group, would like to at least get that figured out and what our recommendations are, because we should be pretty concrete here.

I think the whole blacklisting concept and if you follow that down to its logical conclusion, could potentially introduce, you know, API vendors who simply don't like an App, they feel that its competitive, you know, there could all of these subjective reasons so we really should be able to objectify that. And not have to worry about falling into that information blocking realm, right? There shouldn't...there should be a threshold where they don't have to worry about that additional liability.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Meg, this is Aaron, I have a question that I think will help me understand better. Umm, the notion of whitelist, greylist, blacklist, is that bound at the resource level or is that somewhere else in the ecosystem? Or can it...may it be somewhere else in the ecosystem?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So as far as what healthcare providers allow their patients to do, at the end of the day if the healthcare provider, you know, working with their vendor can turn stuff on and off. Now they may decide to make decisions independently, so you know, Boston Children's Hospital can have its own blacklist or whitelist for that matter...

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Umm.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

...or they may decide to try to share that with a third party and take advice from an outside company that they trust to make those decisions for them. But at the end of the day, I don't know if that's...when it comes to patient access, I don't think it makes sense for us to say that that solves the problem. At the end of the day, the provider's the one responsible for sharing data with the patient.

Let me see if I can break what Meg said into two parts. So one was, what does the provider have a responsibility for? So is the provider responsible for figuring out what Apps are sharing my data against my will and doing bad things? And then the second is what if the respo...what is the provider allowed to

do if they feel like it? Because we might say that the provider's not responsible for protecting me, but the provider might still feel like trying to protect me.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Mm-hmm.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

And maybe it's worth describing one scenario, which is sort of realistic, which is a provider has a thousand patients who have used a particular App and it's been good and patients have loved it and they've been using it for a year and then suddenly the provider starts noticing signs that this App has been hacked. It starts issuing API requests in a new and different pattern and faster and requesting more data; all within its scope, all within what it's authorized to do, but the provider basically suspects or actually is quite sure that the App has been compromised.

And that might be a case where the provider might want to blacklist the App, even if they don't have a responsibility to blacklist the App. And then so the question is, well can they? And you know, what if they're wrong about that decision or what if in another sort of slippery slope argument, what if the providers decide to blacklisting just because? That's the sort of scope of issues that we're talking about here.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

So just to repeat it to make sure I heard you correctly, what is the provider obligated to do with regards to these lists and what are they permitted with regards to these lists?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah and it seems to me, you know, the provider is not obligated to provide any...in the case where a patient has said, "Yes, share my data with this App; I'm willing to take the risk." It seems to me the provider is not obligated to give them extra protection. They might choose to, but it's...ultimately the risk is on the patient, the provider cannot have an obligation to further protect the patient. That seems pretty clear to me; does anybody sort of dispute that idea?

M

(Indiscernible)

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

This is David Yak...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Go ahead, David.

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

Yeah, I guess the corollary I'm thinking about is e-mail. So if I subscribe to an e-mail service, almost all of them provide a blacklisting capability within their e-mail filter, but they don't delete messages, they'll just either quarantine them or warn me or put them into a bucket. Now I know that's message by message, but I wonder if there's something along those lines where we don't recommend that we actually block access, but rather we just provide different levels of risk awareness to the consumer and let them make the decision. I really don't know who would be the authority to make a decision on behalf of a patient to say absolutely no to someone.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, I don't know that there's a good quarantine analog or analogy here. The only analogy I can think of is turning off access for a while when you're worried and letting the consumer turn it back on later if they decide to. But e-mail has the wonderful property that you've got...it's asynchronous and you can queue it up...

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

Yeah.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

...and decide how to treat it later. And of course...quite that...

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

It would be more of a risk...it would be more like risk awareness or risk levels or something like that to just inform the patient as opposed to taking action on their behalf.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Right. And so I think that we spell this out reasonably well right now in our discussion about endorsements and hopefully we'll have time to talk about it a little bit. But let's assume that there are risk levels and when a patient makes a decision, it's a decision that's informed by a reasonable discussion of the risk level. Still there's the question about what happens a year later when suddenly the risk estimation changes overnight because of, you know a new pattern of activity that, you know maybe the App has been hacked, you know, is it okay for the provider to just turn off access for a while in that circumstance? That's what we're sort of grappling with.

M

Josh, this is...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

In that case Josh...this is Leslie, in that case then, is the provider then understanding that that App has been hacked and the harm therefore may be to the data they protect or...which is a no-brainer, turn it off.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

No.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Is that they see that it may be hacked and it might provide harm to the person who's using it?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, the latter.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

And if that's the case...the latter, okay. So if it's the latter, I think all you can do is inform, but it's still a buyer beware situation. I do like the e-mail analogy, that some way...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So you think it would...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

...it's like having to block that.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So you think it would be inappropriate, for example, for the provider to turn off access and send the patient a note and say, "We noticed you were using this App, it looks like it may have been hacked. We've temporarily suspended it to protect you but click here to re-enable it;" you think even that would be going beyond what a provider could do, would be allowed to do?

M

I would...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

I think that that's a reasonable approach; I am concerned about that going down a slippery slope of...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah.

M

Me, too.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

...benev...unintentional benevolence that ends up being actually says, "Well we don't trust you unless you're a HIPAA certified...we don't trust you to last."

M

I also, yeah, maybe the inverse or the same argument is as a provider, now do I become liable for policing what Apps my consumers use?

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

That's exactly right; this is Aaron Miri and as a provider CIO I can tell you that that is a slippery slope and I completely agree with what everybody is saying on this call. We need to be very careful about that because I do worry that this could be misconstrued and taken a different direction.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

There's this...this is Aaron.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So is there a consensus on...

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

I have one...I can consent to it, or you know, support it 100% if we make an explicit statement that the underlying assumption is the provider themselves are not at risk of harm when they follow certain behaviors or patterns.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Well I...again, this is Aaron Miri. I would worry that, we need to get clarification from the OCR on that and how that works because again, there is a fine line there of what we are responsible for, especially

when it comes to patient data. So I would be curious what the legal interpretation, what the OCR interpretation is of that, because that's what the providers are going to listen to.

M

Right.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

And I agree with the idea it's a buyer beware market. Now is it...I don't think the provider should be obligated to police; however, if you could self-police and say, I'm going to report this App to such and such. So just like we have the testing certification process that we have, is there a role where ONC can play that we can build into these requirements that when someone does want to blacklist or greylist an App they're using, it can be reported to a consumer-friendly website that indicates where the bad actors are? I'm just concerned that the provider in...with all of the best intentions will ratchet us back to a place that nothing gets connected or be overwhelmed by policing the gazillions of Apps that a patient might use.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

And I think we get to that position if there's, and you know Josh you've kind of asserted this before, that there are ways for providers themselves to protect their system from any risk that an API may present to them. And that's what I understand from OCR, there's two barriers, you know two kinds of things to get over is first, the provider's required to share in the means used by the consumer, if it can be shown that it will cause no harm to the provider system, right? Which is a...I'm just saying, that's a pre-condition and, you know however you expose these APIs, you're safe. And then the second is that it adds no cost to their operations.

So, you know we've talked about before how a properly exposed API will essentially protect the provider's system from harm by a malicious API. And I just, I'd want to make sure that we make that clear. If you cannot be harmed by any API that comes and knocks on your door and properly authenticates the user, there's no reason why you would need a whitelist or a blacklist for them to prevent denial of service attacks, I guess.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

So Josh, this is Meg; that last part of the sentence that you just wrote "or if the App poses a threat;" do we want to clarify what type of threat? Is it existing, it's active; it's a high enough threat? I would think that we would want to put some comments in there just to kind of objectify it. And I'd also like to see a statement in here around the API developer as well; so if, you know, if Cerner for example becomes aware of an App that consistently violates the documented use of our API, I think that it's important that the developer also have a mechanism and protection for blacklisting that App as well.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Mm-hmm.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, so I'm not really that happy with what's written down here now, I'm just trying to capture the discussion to reflect back on it. And so here's what I think we've heard. I think we've heard that, you know first of all this is all assuming that providers aren't liable when Apps go wrong downstream. So assuming that that is true, you know, what we're saying is that ONC would clarify that providers are not

obligated to protect patients by blacklisting suspicious Apps. And furthermore, API providers aren't even allowed to blacklist an App for the protection of the consumer; if that's their reason, that's not okay.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Agreed.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Is that what I'm hearing?

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Only if it's going to harm them, like...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

However, the App providers may suspend access or blacklist an App that has breached the API provider's terms of service; this was language that was here before. Do we think that's the case or an App that poses a threat to the provider's own system and, you know, we have to say what that would mean.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Yes, I'm sorry, that's exactly right, or that's what I would...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

It's to their own system or the PHI-based attacks within that system.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Yeah.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Or, you know...

Multiple speakers

(People speaking over one another).

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

What if we become aware that the App has a faulty...then, for example; that's different, that's not a privacy issue or a security issue, that's kind of a quality issue; is that a comment that we can make here as well? I just...I'm trying to fig...if we're trying to scope out what's reasonable and what's not, to try to use the best language here. I'm not sure I'm comfortable with saying unequivocally the provider or the API doesn't have the right to blacklist unless there is a high enough existing security threat. I just want to; you know talk through that a little bit.

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

Hey, this is Yak; a little bit of thought about this question of why e-mail is different. But you know as I'm thinking about it, isn't it true that these questions do need to be asked on every single data access that an App makes to an API and that the API may well be adjusting this blacklist? And so that it does have to be transactional, I don't think this is a one-time shot.

The other thing I...that in a blacklist, right, is that as a consumer I can choose to use the blacklist service that's provided by the e-mail or turn it off and say, "I don't want to use your blacklist, thank you very much." But even if I have it on, I can override it and I can say, "Always allow, always deny or put it into the quarantine." Now I don't know whether quarantine is the right model, but I'm thinking maybe transactional is more valuable here than sort of that one time, initial consent that's being granted or denied.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Well, so let me make just a quick comment about the logistics there. If we say that a patient has to say okay every time an App is accessing their data, what that means is that you can no longer actually connect your App to the provider's system and keep them synchronized and up-to-date in the background. You no longer have the convenience of establishing long-term access and...

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Well Josh, there is a feature that says, always all...there is a feature that says, "Always allow," so a consumer can say, always allow and then the blacklist action will not have any effect for that user for that service.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, I'm just afraid that...so we want to be clear about what the properties are of the system that we're specifying. We're not going to be able to actually define the technology specifications for all this stuff and I don't think vendors will be able to do it in a standard way either because there's a lot moving parts in the kind of system that you're describing.

So, I mean one thing we could say is maybe, provider...API providers are allowed to blacklist an App, you know if a patient opts in to their blacklisting service, right? I mean the bottom line there is what we're saying though, what that means is that they're not allowed to blacklist an App in a way that patients cannot override. At the end of the day, the patient needs to be able to override it. Is that sort of where we are?

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

(Indiscernible)

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

I'm not...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Can I comment back to Yak? So I think on the e-mail analogy the difference, you know those are binary and asynchronous and they are transactional but what we're talking about with the API, this is a query to bring up large results and I agree, we can't make a transactional level commission each time; I agree with Josh on that. I do think that the idea of having a way for a person to identify a blacklist and turn something on is the way to go. If the provider wants to advise people, "Hey here's a suspicious App we know of," that's a nice courtesy.

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

Okay, just...

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

And this is Meg; maybe we're kind of looking...maybe we need to discuss some semantics. So we just heard in the presentation this layered trust model, you know white, grey, blacklist and I think that we need to define blacklist, you know absolutely without...beyond a shadow of a doubt known bad parties, attack sites, the provider and the API technology should be able to shut those down, right?

So maybe the next...maybe this discussion is really around this supposed greylist or something that, you know where we begin to look at sequestering or some other type of mechanism. But I don't want to get too far away from saying the providers and the API technology need to be able to protect certain types of App activity.

M

Meg I think...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So, I mean from...

M

So...oh, sorry Josh.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

...my perspective, let's just say from my perspective I think it would be reasonable for API providers to turn off access for certain Apps, as long as the patient can undo it. And when the patient undoes it, they can also say, "And, you know, and don't make this decision again; I want it and I want to keep having it and, you know stop trying to help me."

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

I agree.

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

Yeah and I think by the nature of...I agree with that Josh. I agree and I think that makes it transactional, I think that means it has to be current as per request. You can't just do this one time; it has to be an ongoing surveillance. So I think I will stand on the belief that this has to be transactional.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

I think we may be using those terms a little differently so I'll try not to use the word transactional. All I'm saying is that the consumer should have the ability to share data with an App over time and at some point in time the provider may actually turn off that access based on new information. And that should be okay, as long as the patient can override the decision and they can't...the provider can't keep going back and forth and say, "I want off," and a patient says, "No, I want on," and then off and on back and forth again.

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

Yeah, yeah.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

So Josh, if this App...if this App has been blacklisted because it is an attack site, for example, like the worst of the worst, it's been...it's malicious, it's been hacked, it's been taken over; we still want to be able to recommend that they...that the patient can override that?

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

So what site is being hacked in your mind Meg?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Yeah, I mean at the end of the day I...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

If it's the provider API that's being under attack, it's clearly their right to turn it off. If it's the myhealthcareinmygarage.com App, or App that I use and it's now being taken over by China and being used to all kinds of stuff, that's my problem. And you might have on this whiter site list, hey, myhealthcareinmygarage is really lousy, don't use it and you think the data usage policies are not being adhered to. And so the patient says, "Well yeah, but I like it, and I'm okay with that," then they have the ability to override it.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

But that's different, so...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Put an attack on the API is...the provider has a right to protect themselves.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

Yeah, you talked about two different things. The first one was the actual hacking, right, the bad App and then the second one is we don't think that they're following the data usage policies; those are two different levels of threat, levels of risk. And my point is that when you're at that highest one, if you caught this...if you know for sure that this site has been hacked and, you know is doing really bad things, I'm not sure that it makes sense for you to allow for the consumer to be able to override that.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

It's their data, it's their risk. How could even...how could...first of all, how would Cerner know that or the API provider know that that particular site over here has been hacked in a way that the patient doesn't...maybe the patient's okay with Apple Blossom.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

Umm...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

I just don't see the provider being the police of the Apps; I think that they're the police of the API.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

And I agree you know so data usage policies aren't being followed and I'm actually okay with that, you know, that's a different level of risk. But again, I'm kind of going back to, I think that it might make sense to carve out the situation where, you know all hell is breaking loose and the provider and the API technology says, this is, you know no. And we can look at it later, we can try to reach out and contact

the App developer, you know maybe it's this appeal process where we try to mitigate and figure out what's going on. But while this is really bad news, you know we need to be able to protect our system.

And I recognize, and maybe this is a part where we put in a statement that says, you know this read-only access is not the same level of threat to a system as a write access, and we recognize that. But I'm not sure that we want to get, you know again, right after this, we're kind of setting up a foundation for how this is going to grow. So if the next App that's created is a write-only access, do we really want to set up two different, you know oversight mechanisms that if the provider says, "Oh wait, this is a read-only API, so...that's threatening my system and I'm not going to be able to blacklist it," I'm just not sure that that's the right idea I guess.

So maybe we just need a few clarifying statements in here that just kind of carves out when that really, really bad risk is and then what the processes are to mitigate that that really is not the consumer overriding it. It is access to his data, but it's also the provider system and it's, you know the API technology may have the ability to detect things that the provider can't. So for both of them I do think it's important that we carve out some protection.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So I wonder if we might have consensus over this formulation then which says that we want to make it clear that API providers aren't obligated to protect patients in this way, but they may choose to, and they can blacklist an App, you know, if and when it appears to have been compromised or if it's posing a threat to the provider system. But in general, and the patient can still override that kind of suspension, unless the reason was...unless the reason for the suspension was that the provider's own system was at risk. Is this something that we could agree on as a set of recommendations here?

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

This is Meg and thanks, I'm okay with that.

David Yakimischak, MBA – Senior Vice President, Information Systems – Surescripts

I could see that.

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Yeah, this is Aaron Miri, I'm good with that.

M

Josh, can you be a little more explicit about the provider not being liable and just add the phrase for downstream or something to that effect? The provider would be liable up until they disclose, yeah?

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Right. Yeah, so that's fine. So, one second...

M

(Whispering)...whether they suspend API access to it...breach to API provider terms of service...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Hi, sorry about that, I'm borrowing some conference space in between meetings here. Um, so let's highlight this and move on to another topic for now. Of course we can always come back to this, but I think this is a fairly nuanced approach and I'm happy with where this has gone.

Let's talk about one other topic, and this is a related topic that I'd like to spend at least 10 minutes on before we move on to public comment. I don't think we'll wrap it up in 10 minutes, but I want to sort of present explicitly where we've gotten to on this topic and then get some feedback. So this is the topic of endorsements and it's called the process of endorsement/certification. That's the name of the topic; it's on page eight in the document, if someone could scroll to that page, please. And what I'll do is just describe this very briefly; it's really just a page and I'll almost read it here on the phone.

So basically what we said here on page eight is, for background, there's going to be this diverse ecosystem of healthcare Apps in which some are "more trustworthy than others." And we recognize that trustworthiness is a very broad concept and it has lots of facets; means different things to different people. But it includes stuff like: Is this App giving me good clinical advice? Is it protecting your privacy? Are the App servers well secured? Does it actually provide value for you? Is it worth the money that it costs? Is it stable? Is it going to be around in 18 months when you want to check it again? Is it developed by people with a strong reputation in this space?

And this is not an exclusive list, but this is just designed to give you the sense that there's more stuff in this conversation than, you know, than we could possibly regulate. Nevertheless, these things are quite important. So the point is patients are going to have lots of choices to make in the marketplace and we want to make sure that patients can learn any information available about which Apps are in fact trustworthy.

So in terms of findings, to summarize what we heard during our testimonies, you know we heard from a number of healthcare providers who expressed concerns about what it would mean to allow unknown patient-designated Apps to connect to their APIs. And these healthcare providers would in general be more comfortable if they knew that patient-designated Apps were somehow certified or had a Seal of Approval or some kind of endorsement from an outside party; that would make them more comfortable.

And at the same time we heard from patients and from consumer representatives who expressed a concern that if Apps need to be certified, that's going to unduly restrict consumer choice. And these representatives on the consumer side were worried that these kinds of restrictions would eventually add up to a violation of a patient's right to access. So that's sort of what we heard.

And I'm wondering if we're...okay, able to scroll to this topic here, it's the process of endorsement/certification? It's on page eight. I just want to make sure that folks are able to line up what they're seeing with what they're hearing, or at least know where to look in the document.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

I'm going to have an epileptic seizure soon.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Oh no, umm, well I'll just keep going here. Listen to the description here at least, it's the one that starts with this bulleted list; it's not this, it's on page eight, the page that starts with the topic, "Process of Endorsement;" there's a link to it from the...process. Okay, so...

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

There we go.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So that's a little bit of the background about what we're trying to do and what we heard and then I want to describe, you know the sort of two paragraphs of recommendations that I drafted over the weekend. What we want to recommend is that ONC should push toward an ecosystem where we don't have centralized certification or testing of Apps in a sort of government mandated way, but instead we want ONC to pursue the kind of regulations that are going to enable a secondary market in endorsements.

And in this kind of secondary market, what it means is that there's many different kinds of organizations, might be EHR vendors and security experts and consumer advocates and professional societies and, you know hospitals. All these different kinds of organizations might want to give their own endorsement to a given App and there should be a publically visible way for them to do that in a distributed, non-centralized process.

And then these endorsements can help inform a consumer's choice. So when a consumer is thinking about which App they want to use, they can look at the endorsements that are available. And if we have an ecosystem like this, it'll allow aggregators to come in and actually roll up all the endorsements that have been made about a given App and allow consumers to browse the available Apps.

And a consumer could even go to one of these aggregators and say, "Show me all the Apps that are trusted by Consumer Reports," or "Show me all the Apps that the National Association of Trusted Exchange has done a privacy policy review on and make sure that NATE says that they're not going share my personal data with advertisers." We could have all kinds of very granular claims like that and consumers could filter based on the things that mattered to them.

So we'd have a discoverability property and then at the same time for provider organizations, when they're actually allowing a patient to approve one of these Apps. The point is that you can't use the lack of endorsement as a way to block the registration of an App or to prevent a patient from sharing data with an App. So it's not that endorsements are necessary, and yet they can be helpful. So if a provider organization happens to trust a few different endorsing bodies, they might prominently display endorsements from those bodies to the patient at the time of App approval.

So for example when a patient is asked to confirm the decision to share data with a specific App, the healthcare provider that's hosting the API might display endorsements from the American Pediatric Association, for example, if that's an organization that the provider happens to trust. Or they might display a warning to the patient saying, "This App doesn't have any endorsements that we know of, or doesn't have any endorsements from anyone we trust, so you can still proceed, but it's at your own risk."

So from my perspective, this was sort of the best aligned approach that we could describe for ONC to pursue. But I just want to take at least three or four minutes and get reactions to this and see if this is the kind of thing that we might be able to drive towards consensus on here.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

Josh, this is Les and I agree with all of it in principle; my concern is in the further discussion on then if the covered entity recommends an App or offers and App in their App Store, that somehow that commercial relationship indicates that that is now a BAA and not just simply an App that the patient can link to, which becomes just another form of keeping the patient getting either a trusted App. If they just really want the provider to say yeah, I really like it versus having to go through all of the BAA.

So I'm okay with this as long as it doesn't lead to an assumption that there is a commercial relationship, well there could be a commercial relationship, but there isn't necessarily a data relationship that would mandate the App developer being a BAA.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Okay, I don't know where you see that in the text that's written here; if there's something that we could change...

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

It was in the other document...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

...or speak to or try to clarify.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

It was the document you sent me a link to the OCR site use cases and then I commented back to the link on those. So that was...that's my only...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

I got it. So that was on the topic of when is a BAA required. So I think that's an important topic and I hope we'll have time for it, but I think its separate from this endorsement issue, unless you think they're linked in a way that I'm missing.

Leslie Kelly Hall – Senior Vice President of Policy – Healthwise

I do think they're linked because at least the way I read it, it said that if the provider was providing the App on their App Store, that somehow that commercial relationship either through endorsement or hiring that App provider made that then be a BAA relationship. So, I agree with what's here in this as long as it doesn't imply that that endorsement makes it a relationship that mandates a BAA.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Okay, I...it certainly doesn't imply that for me and if you think there's language we could add that would help to ensure that it doesn't imply it for you or for others, by all means I think that would be useful. Are there other reactions or responses? Is this the kind of thing...is anybody just totally worried about this whole scheme and thinks this is a terrible kind of proposal because we should be doing something very structurally different?

M

Josh I think it's great. A friendly amendment to sort of echo what Leslie was saying is perhaps a footnote that explicitly says in no way does this imply that a business associate relationship exists with the provider.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Sure, I'll just add that in right now when we talk about provider organizations.

M

That would be great. Is...Leslie?

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

So, but...this is Meg; just for clarification. We could link to the OCR guidance that makes that statement authoritatively, I don't think the task force should define whether we, you know, this says or doesn't. So if it's not addressed Leslie specifically in the guidance, if it is in the guidance document, then let's link to it, but if it's not, let's ask for clarification that we can link to.

Multiple speakers

(Indiscernible)

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

I appreciate that, this is Aaron Miri; I appreciate that wholeheartedly.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Josh...

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

So I added a footnote saying we don't think it...that...it's not implying that there's a business associate agreement, but it's not implying that there isn't; just it's a totally separate issue. We've only got five minutes and I think that we need that time for public comment. If anybody has a really quick remark, you know 30 seconds, please go ahead.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Thirty seconds, should we prohibit any kind of negative; you know dis-endorsement such as...blacklist.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

I don't see how we could; it feels like a free speech issue. I don't see how we could prohibit organizations from telling people not to use Apps.

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Good. Okay. I just wanted to bring it out.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

No, I mean I think that there's actually going to be a market in those, too...

Aaron Seib – Chief Executive Officer – National Association for Trusted Exchange (NATE)

Yeah.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

...you know things that the security experts think have been hacked. It seems like that could be an interesting...

Aaron Miri, MBA, PMP, CHCIO – Chief Information Officer – Walnut Hill Medical Center

Yeah, I mean this is Aaron Miri again, last 10 seconds here let me just state this, regardless of how we proceed on this path, the last thing we want to do is encourage any ability for folks to somehow exclude from the marketplace relationships leveraging APIs and give any type of provider organization, such as myself, any kind of wiggle room in providing access to patient data. Because I guarantee you, free access is the key and any type of fear factor will be leveraged in multiple ways. So what Leslie's saying, what I'm saying, what others are saying is just let it be an even playing field for everybody.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Good, I highly agree. If you see something in here that suggests otherwise, by all means come back and suggest how we could clean that up. And with that, let me pause and I think turn back over to Meg, to see if Meg has any closing remarks and then move on to public comment.

Meg Marshall, JD – Director, Government Health Policy – Cerner Corporation

Yeah, no...not...no closing comments; thanks, this is a great discussion. So just from a process, we've got two meetings before our draft recommendations need to be presented to the committees. I think that it's important to note that we're not going to be able to have this level of discussion for the entire 26 page document.

So I echo Josh, please take a look, go through, review, put your comments in there and let's really flesh out how to make use of the three hours of face-to-face time that we actually do have before the recommenda...the draft recommendations need to be done. So as much as we can do offline, let's see if we can.

Joshua C. Mandel, MD, SB – Research Faculty – Harvard Medical School

Perfect. All right and let me turn it over to Michelle then with that.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Umm, thank you, Josh. Jaclyn and Marcus, can we open up to public comment?

Public Comment

Jaclyn Fontanella, MPH – Digital Project Manager – Altarum Institute

Sure. If you are listening via your computer speakers, you may dial 1-877-705-2976 and press *1 to be placed in the comment queue. If you are on the phone and would like to make a public comment, please press *1 at this time.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

While we wait for folks to call in, there were a couple of comments in the chat. We'll e-mail these out to the group but I will share a few of them, too. Joel Arnold from Aetna, "The way API Gateways work is that Apps need to be authorized to access the APIs. This is not a consumer/user level. Allowing an App access or not is a decision made for the App, not the user."

And then we have another comment from Luis Maas from Direct, "Good progress on blacklisting section but as written, patient can override disabled access for App that violates terms of service."

And it looks like we also have a public comment from Lucy Johns from DirectTrust. As a reminder Lucy, public comment is limited to three minutes and please go ahead.

Lucy Johns, MPH – Vice Chair, Board of Directors - DirectTrust

Thank you. A public comment of this task force March 8 asserted among other points that DirectTrust "has been a disaster" for patients. DirectTrust takes this opportunity to respectfully disagree. DirectTrust, in creating a voluntary security and trust framework and supporting a large and growing

network for health information exchange, is mindful of and honors the importance of patients and consumers and their personal health information.

What distinguishes DirectTrust's role in the HIT space is commitment to privacy, security and trust in identity, consistent with HIPAA. Given the increasing rate and number of breaches of PHI, and the huge cost to patients and providers and payers, DirectTrust hopes there is little disagreement about the need for security controls adequate to protection of HIE. We believe that without privacy protections and assurance of security, there can't be, and in fact should not be, trust in HIE.

DirectTrust's high level of security and trust in identity is a voluntary option for organizations and individuals who value security and privacy of PHI when engaging in HIE. We are working on a comprehensive strategy that accords patient the same opportunity to protect their PHI as we make available to providers. We have members now designing the Direct exchange experience expressly for patients and consumers.

DirectTrust encourages and supports policy that extends the voluntary option of security and privacy protections into expansion of HIE technology, such as APIs and their associated Apps. A medical record is alleged to be worth \$50 on the black market, compared to about \$0.50 for a mere social security number. Patients and consumers should have the option to try to ensure this doesn't happen to them. Thank you. Please see the remainder of this statement in your packet.

Michelle Consolazio, MPA – Federal Advisory Committee Program Lead – Office of the National Coordinator for Health Information Technology

Thanks, Lucy. And that's the end of our public comment. So thank you everyone, we appreciate you staying with us for the minute that we went over. And our next meeting is on Monday, March 28 at 10:30. Thank you everyone, have a great day.

Multiple speakers

Thank you.

Public Comment Received During the Meeting

1. Joe Arnold: Hi Joe Arnold from Aetna. The way API Gateways work is that apps need to be Authorized to access the APIs. This is not a consumer/user level. Allowing an app access (or not) is a decision made for the app not the user.
2. Luis Maas: Luis Maas EMR Direct - good progress on blacklisting section, but as written patient can override disabled access for app that violates terms of service.