

**HIT Standards Committee
NwHIN Power Team
Transcript
June 12, 2013**

Presentation

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thank you, good afternoon everybody, this is MacKenzie Robertson in the Office of the National Coordinator for Health IT. This is a meeting of the HIT Standards Committee's NwHIN Power Team. This is a public call and there is time for public comment built into the agenda and the call is also being recorded and transcribed so please make sure you identify yourself when speaking. I'll now go through the roll call. Dixie Baker?

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

I'm here.

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thanks, Dixie. David McCallie?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Here.

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thanks, David. Floyd Eisenberg? David Groves? Arien Malec? Marc Overhage? Wes Rishel is present; he's just on mute at the moment. Cris Ross? Tim Cromwell? Ollie Gray?

Ollie Gray – Research Program Manager - TATRC

Here.

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thanks, Ollie. Nancy Orvis? And any ONC staff members on the line, if you could please identify yourself?

Ellen V. Makar, MSN, RN-BC, CPHIMS, CCM, CENP – Office of the National Coordinator

Ellen Makar is here.

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thanks, Ellen

Avinash Shanbhag – Office of the National Coordinator

Avniash Shanbhag is here.

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thanks Avinash. Okay, with that I will turn the agenda over to you Dixie.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Okay, thank you and thank you all for dialing in. Ollie, thank you again for lining up our speakers for today. This is the third of three presentations and discussions that we'll have about potential candidates for standards for additional transport standards. So, why don't you just go to the next slide it's kind of – well, this is the agenda.

We'll start with a reminder of what the task is and then we'll have a discussion of the RESTful Health Exchange or RHEX by the team from MITRE and from TATRC, they'll talk about the development of that standard as well as TATRC's use of it in a pilot.

And then we're going to move into an informal readiness assessment of the three transport protocols we've looked at Blue Button Plus, FHIR and RHEX and we'll close with public comments. So, with that go to the next slide please.

This is the task, again, just to remind you it's to recommend additional standards to support transport of data to and from patients and as David has pointed out obviously the transport wouldn't be limited to that purpose, but that's really the focus of this particular task. David, do want to add anymore at this point?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

No, I do want to add a little, a slightly different comment, first I agree with that, that this might not strictly be limited to patients. Also, to add that it's not strictly limited to transport either in the sense that at least for the FHIR, I'm sorry the – well both the Blue Button and the FHIR we talked a good bit about new ways to represent actual content that could move back and forth. So, it's a little bit of both transport and content I think, certainly in the case of FHIR.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah, I agree with you. I think the task assignment says transport but they mean that as a broader sense than just the transport –

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, yeah.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah, that's a good point.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

It's not the technical layered stacked definition of transport it's really movement of data.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yes, yes, yes, good, good, good addition, thank you. Okay, with that let me introduce Ollie. Ollie Gray is a member of the NWHIN Power Team and Ollie you can – if you don't mind you can introduce the members of your team.

Ollie Gray – Research Program Manager - TATRC

Sure.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

And Ollie has some slides he's going to be – that the team is going to be using.

Ollie Gray – Research Program Manager - TATRC

Thanks, Dixie, it's my pleasure to introduce and I'm sure most of you know already Justin Richer, we have Suzette Stoutenburg who will not be able to join us today and Sam Sayer will be stepping in for Suzette, and talking about the work that our team which included Kim Pham and Chrisjan Master from here at TATRC. We worked last year jointly on a project from the FHA and the ONC using the RESTful, REST to develop what we called RHEX to exchange patient data and while the slides come up I'm going to let Sam get ready. Sam is going to go through the briefing of how we did two different projects one of which TATRC was involved with and the other the New England Health Information Exchange was involved with. So, I'm going to turn it over to Sam for the presentation.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Thank you Ollie.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

And now can we get the slides that I believe MacKenzie just sent these to Caitlin to you right before the –

Caitlin Collins – Project Coordinator – Altarum Institute

Yes, we're working on uploading them right now there is just quite a few of them so it's taking a lot of time but they'll be up momentarily.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Thank you.

Ollie Gray – Research Program Manager - TATRC

Sam, do you want to go ahead while the slides are coming up and give a little bit more of an introduction?

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Sure, so I'm just going to lead off here discussing some of the past pilots and also the present pilots using RHEX and then I'll turn it over to Justin Richer to talk about a comparison of RHEX, FHIR and Blue Button Plus and also to go into a more technical deep dive on OAuth and OpenID Connect.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah, thank you I appreciate that we had asked for that as well the comparison about how the three use OAuth, so, thank you.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Yes, no problem, let's see the slides still aren't up, so excuse me I've only known I've just been doing this for an hour so it might be a little disjointed, but, so RHEX is an open source exploratory project to, you know, take web standards like OAuth and OpenID and apply them to health information exchange. So, taking out what is the current best of breed on the web and, you know, do the certain things that need to be done in order to make it applicable for exchanging secure, health information securely. So, we were sponsored by ONC in 2012 and we're currently working with TATRC following up on some work we did last year.

Okay, just need to go I think to slide 4 please? Thank you. So, the TATRC pilot last year focused on exchanging data between a physician inside the military health system using AHLTA and a third-party provider employed by the reference implementation of RHEX the patient data server. So, used OpenID Connect to allow doctors to log into systems that were sort of outside – that were outside their own domains, so logging in – so a third-party provider could log into a system at a military health system and access the data that they needed to get in order to perform referrals. So, in this case it was exchanging vital signs before performing a knee surgery.

And then the second pilot was done with HealthInfoNet, the Maine Health Information Exchange, this one was a quite a bit different. So, Maine their goal is to connect every single provider practice in the State of Maine up to the Health Information Exchange and because Maine is such a large State, such a rural State they needed to find a technology that they can use that they can connect very small practices, very small family rural practices without a lot of expense.

So, they turned to us, you know, to provide the open standards and the open source and to help them build the open source software in order to accomplish this. So, to install, you know, a very small HTTP client at a small family practice, in this case it was an island off the coast of Maine and they had a corresponding RHEX server installed at the HIE which would then perform an OAuth 2 handshake and exchange information to integrate the data into the health information exchange. Next slide, please.

So, the pilots were both successful, we were able to, you know, do the secure exchange of health data over the web and also to move high volumes of data over the web to support the HIE, I believe in the demonstration last year we moved 100 patients in about 20-30 seconds. We are continuing to pilot with TATRC in FY13 which I'll get into in just a minute. And the MAINE HIE has, you know, we basically turned all of our software and our knowledge over to them and they are continuing to run with it and they're planning to roll it out to all their sites over 2013.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Can I interrupt and ask a question when you get a break?

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Sure, yes, go ahead?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

The scope of RHEX in terms of the use case you've described, is RHEX the focus, is that just the authentication piece of it or is it the actual exchange of data as well?

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

It's the exchange of data and it's the authorization. So, the client is authorized to send data to the health information exchange and through the OAuth 2 handshake that happens the client verifies that the server is indeed legitimate and vice versa to make sure that, you know, data isn't going to somewhere that it shouldn't and to know that data that is coming in is coming from a legitimate site.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

So, if you could position RHEX against existing approaches to the degree that it overlaps and is different from that would be helpful.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Right.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

I was under the impression that RHEX typically focused on the authentication and authorization part only and that you were using other modes for the actual exchange of data, but it sounds like I might be wrong on that?

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

No, we –

Justin Richer, MS – Lead Technologist – MITRE Corporation

We'll cover that later in the presentation.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Yes.

Justin Richer, MS – Lead Technologist – MITRE Corporation

But we will get there.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Yes, we'll get there, but to compare with other –

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

We're very impatient in our group here.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

That's no problem, but just to compare the reason they brought us in was the sort of standard way of doing things for them was to have a VPN connection to a site which is very expensive and very hard to maintain, especially with no IT staff, and a system using RHEX, using OAuth 2, you know, could run with very little overhead and, you know, it's based on, you know, free and open standards so that they don't have to pay, you know, licensing fees in order to use it.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And I need at some point to understand why not just use Direct or why not just use, you know, XDR, in other words what's – just compare and contrast with the existing choice is.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

They – it will take me a minute to remember, there was a specific reason that the Maine HIE told us they did not want to use Direct I believe it had to do with what it took to set up a Direct client site, I guess the client sites were not amenable to having all the Direct software set up at the client site.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Boy that betrays a deep lack of understanding of how Direct works, but we can –

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

So, please take that with a grain of salt, I'm not the person that – I guess I'm not really the person to answer that question, I'm trying to remember back about 2 years at this point.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

And that's not a comparison we're making right now anyway.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

No.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

I think we should just go on and hear more about what RHEX does.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Thank you. So, for HIMSS 2013 we did an additional pilot with the Maine HIE so now that we had data flowing into the HIE from several sites we wanted to sort of reverse direction and use OAuth 2 to allow access to data through an iPad App so an App called – which was developed at MITRE, we set up the OAuth 2 server at the Maine HIE which allowed, you know, a patient to release their own information to themselves onto the iPad App and, you know, view their own health data over a secure connection and that showed – and we're hoping soon to have a pilot with the VHA in rural Utah to support a plan of care for rural veterans and support updating and creating plans of care across the web instead of, you know, over fax and phone which is the way it's currently done. Next slide, please.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

So, that would be machine oriented or people oriented in the first one?

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

It will people oriented. The system is – there are home care nurses that are currently, are sort of the engines of the process and who have to correspond with doctors, and you know, as I mentioned before they had to basically call back or send faxes of documents to – back to the VA and the process is a little unwieldy for them.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Thank you.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And at some point contrast that to Blue Button, etcetera, but I assume you're going to get there.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Okay.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Just to clarify we've asked for them to compare only the use of OAuth 2 in the three we haven't asked for a complete comparison of all protocols.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

No, but Dixie –

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Later I think Justin is going to address that, but it's contained to the use of OAuth is what we've asked them to do.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah, and Dixie and others I think that all of that will make a lot more sense if we just hold tight.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Good, okay.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Just give me two more slides and I think we'll be there.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Great, thanks.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

All right. Next slide, please. So, we also plan on doing additional pilots with TATRC. Our current pilot we're sharing images between AHLTA and third-party provider systems so exchanging DICOM images between two third-parties and an AHLTA system and then in the future we hope to do another hReader pilot but this time using data that's in AHLTA and then finally look into using RHEX to securely migrate health data from AHLTA to VistA to support seamless exchange of records, you know, when service members become veterans. Next slide, please.

So, this is just a simple graphic demonstrating the current imaging pilot we're doing with TATRC, it's very similar to the one we did last year but we sort of – we've added an additional step of including an extra third-party and also exchanging image information over the web which is a little bit more difficult than exchanging XML and JSON.

So, now in this pilot a PCM will have access to, you know, a radiologists – images the radiologist takes and also radiologist report and can then refer a patient to, you know, a third-party surgeon for a procedure based on what comes in from the radiologist and the radiologist or excuse me the orthopedic surgeon can have access to all the relevant imaging and report data in order to perform the surgery. Next slide, please.

Okay, here's the moment everyone has been waiting for, I'll turn it over to Justin to talk about the comparisons between RHEX, FHIR and Blue Button Plus.

Justin Richer, MS – Lead Technologist – MITRE Corporation

All right, hello everybody, my name is Justin Richer with the MITRE Corporation and I am actually personally directly involved with two of the three projects on the slide here so I'm very happy to be able to come in and compare and contrast kind of what's going on between all of these different efforts.

Now the first thing I wanted to point out is that all three different efforts really have a different driving focus, they're really looking at solving a different part of the larger overall problem. The driving impetus to RHEX was to build something, to actually take a bunch of existing technologies such as hData, OpenID and OAuth and build those together into a working prototype that we could put onto real systems.

FHIR is a set about defining a set of resources that represent the health and healthcare information and how you get those resources in and out, so all of the – sort of the RESTful protocols related to them. RHEX itself was never specifically set out to actually define what those look like but the RHEX systems that we built and deployed have all used the hData protocol to actually handle that.

Now, Blue Button is taking the approach of enabling patients to access their own data in both human readable and machine readable formats and to be able to share them so that this leads to a very different type of environment and very different drive for each project. Next slide, please.

So, the network style for RHEX, since we are focusing on building it in terms of these pilot systems, is that we kind of know all of the providers and end points, and the applications, everything is sort of, you know, spun up on an on-boarded out-of-band not really a dynamic network which is not to say that RHEX can't handle that it's just out of scope and undefined much like it is for FHIR, they don't really care kind of what the network of who is talking to who and who is sharing data, they're concerned with, on their project, what the data looks like, how you get it and how you manipulate it.

Blue Button on the other hand is explicitly set out to solve the problem of a highly dynamic and distributed environment with registry components and trust bundles and all sorts of other things to actually make this happen. What this means is that each group is either making certain assumptions about how the world works or they're explicitly saying, well we can work in a bunch of different places such as with FHIR we can work in a bunch of different ways but we're not trying to solve that particular part of the problem. Next slide, please.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Can I interrupt and ask a question here?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Sure.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

This is David, on the – when you are referring to Blue Button here are you referring to the original Blue Button or to the view, download and transmit version of Blue Button or to the Blue Button Plus/Pull that we heard about from Josh Mandel a couple of weeks ago? I did like the pull.

Justin Richer, MS – Lead Technologist – MITRE Corporation

The later.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

The later, okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah, the later, it's the Blue Button Plus/Pull though the work that we've been doing in the pull working group is also feeding into the push working group in order to allow the Direct system to be more dynamic and useable by actual people.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, because I was puzzled by the reference to registry components because that's sort of an optional part of the Blue Button Pull model that you could actually have a service that registers trusted outs but is really not a required part of it so it's not – to me it's not really what Blue Button Pull is about.

Justin Richer, MS – Lead Technologist – MITRE Corporation

That's incorrect. So, the registry is absolutely required as part of how Blue Button is doing things what's optional is whether or not your provider or your application has to be in a registry in order for it to work with other things that are Blue Button compliant.

We expect, we fully expect for the registry components to be the anchors of all of that and there are policy and trust decisions that base off of that and honestly in my own professional opinion where we need to go with the Blue Button Plus model both on the Pull and the Push side is to go into a world of trust frameworks and dynamically allocate trust frameworks that allow it to work in environments like this where some decisions are anchored in a registry and/or a trust bundle type of situation or they're based on just the dynamic aspects of the moment and I would still like that next slide, please. Thank you.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

I also think you're mixing apples, oranges and pomegranates here.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes that's kind of my point.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

It's a little bit misleading.

Justin Richer, MS – Lead Technologist – MITRE Corporation

No, I'm actually trying to point out that these are apples, oranges and kiwis, and pomegranates everything else.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

These are very different projects that are solving different problems that was what I was trying to set out to say.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

They're approaching different layers of – yeah different problems too, it's not just different problems but it's different parts of the stack, but keep going, this looks good.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah, I do have a comment about the previous slide of the network style, you know, maybe you meant something different for network style from what I have in mind, but the FHIR specification does include the RESTful transport.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

As part of the specification.

Justin Richer, MS – Lead Technologist – MITRE Corporation

That's not the network style.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah, that's right; I'm not sure what a network style is.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, by network style I mean whether it is pre-defined, whether it's federated or whether it's distributed and those are all on a sliding scale of what end points you know, who is going to be doing what ahead of time and what trust model is associated to them. So, I don't mean network in terms of the – you know, the – 7 layer model of network.

I mean network as in there must be clients and there must be servers somewhere and somebody has to know where they are. And how much you know about any of those before a single bit goes across the wire determines the style of network that your data has to flow through. Does that help clarify?

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yes.

Justin Richer, MS – Lead Technologist – MITRE Corporation

All right.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

So, what would be the network style that's used for e-mail today over the Internet?

Justin Richer, MS – Lead Technologist – MITRE Corporation

If you're talking just pure e-mail that is a fully distributed system. If you're talking about S/MIME that is not, because there are trust anchors and certificate routes.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Thank you.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, anyway, the security model and the security questions that you're actually trying to solve on each of these end up being different because of the different network style. Now, first and foremost all of these use HTTP over TLS/SSL as the transport mechanism and what is important here, and I know Dixie and I have spoken about this previously, is that this is what is called a one-way TLS where the client checks the server certificate when it makes the connection to the server but it is not a two-way mutually authenticated TLS connection and this is actually very important, because that's the way that the worldwide web as we know it today actually works in practice. That's the way that actually scales to truly Internet scale and that's explicitly what all three of these are trying to do.

Now, RHEX as you saw in what we've been trying to pilot, actually approaches the authentication problem, the authentication of end-users and providers to remote systems by using OpenID Connect which is an authentication and log-in protocol. It uses OAuth 2 for the authorization layer for getting data from side-to-side and also incidentally for getting identity information about users from side-to-side. This is all due to the fact that OpenID Connect is built on top of OAuth 2 which if we go into the technical deep-dive of how these two protocols work in just a few slides here you'll see how that actually happens.

And so RHEX defines profiles for how to use OAuth 2 and how to use OpenID Connect within a RHEX system, kind of nails down certain optional parameters and optional systems. Now the profiles that we have written for RHEX were never meant to be all inclusive nor were they ever meant to be the only profiles that would only ever be written for use. They just happen to be the first ones and they happen to be ones that fit the use cases for all of the pilots that we've been doing with RHEX and I'm aware of some work for getting a FICAM Profile of both OpenID Connect and OAuth 2 that would be – it would be the kind of thing that a project like RHEX would be able to say, oh, yeah we can use that, you know, we can use that profile for these use cases.

FHIR on the other hand they don't define really what the authentication is nor do they actually define authorization because they're concerned about what the RESTful protocol looks like and the truth is if you're going to secure a RESTful protocol these days the far and away runaway best answer is to use OAuth 2 to do so. So, that's why under – my understanding that's what the FHIR group is actually looking to do. They're suggesting the use of OAuth 2 if you read through all of their related documentation but they haven't come down yet and said, you know, use OAuth 2 in this manner for this way, because that's not specifically the problem that they're trying to solve right now.

Now Blue Button Plus on the other hand, the authentication in Blue Button Plus is again undefined, is assumed to be local to a provider much like in most of the FHIR world is but the authorization model is actually a distributed OAuth 2 model which goes beyond what RHEX's profile does and this is where the registry components actually come into play. So, we have dynamic registration and trusted registration and other mechanisms that allow OAuth 2 to work in a very, very widely distributed world like it does on the public Internet but without it being a complete free for all and so it's a very, very powerful use of the technology that has looked at how RHEX did things in sort of the more static environment and sort of pushed beyond that.

Now in working with the Blue Button Plus group I think that there is a lot of room where the authentication side of the house can actually also be used, but again, that's not a problem that the Blue Button Plus is actually trying to solve right now, it's a different aspect of it, to borrow the previous analogy these kind of are apples, oranges and pomegranates, yes they're all still fruit but they all do look pretty different when you slice them open. Next slide, please.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Well, and this is David, just to –

Justin Richer, MS – Lead Technologist – MITRE Corporation

– the slide.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

David just to pick up on the fruit metaphor, you know, they're not mutually exclusive either.

Justin Richer, MS – Lead Technologist – MITRE Corporation

No.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

So, for example –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Well, and actually hold that thought because the slide that was supposed to be on the screen right now was actually going to show that. So, if I may speak through what should have been on your screens, because I think this will address this, if you don't mind?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Go ahead.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, if we could go back a screen, so the – you can think of RHEX as taking hData and adding in OAuth 2 and OpenID Connect to sort of bring everything together and so hData as the RESTful protocol, OAuth as the authorization and OpenID Connect as the authentication it really kind of adds all of those together and moves all of that forward.

Now FHIR on the other hand is really analogous to, and in many ways compatible with, explicitly compatible with hData much like hData doesn't really define what the authorization or the authentication layers look like, it defines all the rest of that stuff.

So, a project like RHEX could very easily show up and say, we want to do authentication and authorization in this way and use FHIR as our transport mechanism, which brings us to Blue Button Plus, specifically the Pull protocol and that is doing just that and is using the FHIR protocol for all of this sort of API end-points both the document end-points and the searching end-points so sort of the various things that the hData project started but FHIR is bringing to sort of a deeper fruition.

Blue Button Plus is directly using FHIR and explicitly using FHIR to do all of that while Blue Button is also defining how you authenticate or how you do authorization across domains, how you do registration, how you do tracking, what your policy is for different kinds of servers, clients, data and users.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right which is pretty much what I was going to suggest so I'm glad that I got a chance to have that in my mind before reading your slides because it makes sense to me that FHIR could replace the use of hData, which I understand why you used hData because FHIR didn't really exist at the time RHEX got started.

Justin Richer, MS – Lead Technologist – MITRE Corporation

It didn't exist at all as I recall.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right and Blue Button uses FHIR already. So, you know, at the building block level, you know, the building blocks, the core share of building blocks here are OAuth 2, OpenID Connect and FHIR, and of course HTTP, TLS, etcetera, but I won't list them again.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes, absolutely.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

These are three experiments that combine those building blocks in different ways.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

All of which are reasonable. I mean, you might do RHEX differently; you might use FHIR with RHEX if you were starting today, but other than that it seems pretty logical.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Exactly and if you look at FHIR, if you actually look at how FHIR is defined it really borrows quite a lot of its basic model from hData, there are explicit transformation rules in and out of hData's previously defined end-points and document formats.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah and we didn't –

Justin Richer, MS – Lead Technologist – MITRE Corporation

– a goal of the FHIR project.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, we didn't get a chance to ask Lloyd McKenzie when he spoke to us at our last meeting about hData it was on my list and he just unfortunately had a limited time span, but do you know if the hData work, which I think originated from MITRE originally, right?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Is that comfortably to be replaced by FHIR, is that going to be a competition or is this just a transition?

Justin Richer, MS – Lead Technologist – MITRE Corporation

I believe it's just a transition. So, I'm not really the best person to speak on that particular set of standards.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Justin I can jump in here.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Go right ahead.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

So, I'm actually – so I'm still involved with hData so it's currently a DSTU in HL7 and we're currently still refining the spec to move toward the full movement of standard, however, you know, FHIR does seem like for what it does it will be the way forward, but sort of the target for hData going forward is, you know, if your use case doesn't fit in with FHIR you can build your own service with hData and then FHIR is, you know, built on top of some of the building blocks of hData.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay, who was that speaking by the way?

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

This is Sam, sorry.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay, yeah, I don't – we probably should come back to that later, because that will get into too deep weeds but that's very helpful and that will be an – Dixie that will be an interesting discussion point obviously.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah, I agree, yes.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

At some point maybe, maybe we're too early to conclude on that, but okay, go ahead, this is great stuff by the way and the interruptions mean that we're really paying attention, so, hopefully you take it the right way.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yeah.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So at this point that's really all I had for the comparison between the three projects so if you go to the next slide if there is enough interest I can give a pretty quick but still, you know, deep enough to kind of grackle at the different pieces overview of how OAuth actually works, because it's a different security model than what a lot of people are used to in sort of our standard security models that we see with like, you know, mutual certificate chains or, you know, traditional PKI systems or things like that. OAuth turns a lot of that up on its head and still comes out with a very secure and as it turns out very, very, very useable system.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

So –

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Well – I'm sorry.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

This is David, I was just going to say I would like to get a little bit more technical detail but I want to preload you with some questions about the last time I looked at RHEX, which was more than a year ago, so I'm way out of date, there was a lot of debate about the actual profiles to be used and which of the many building blocks could be woven into the way you were using OAuth 2 and OpenID Connect, so I'm curious to know about settling on a set of profiles for healthcare whether that's feasible, whether you consider that an achieved state or not or whether or not that's just not even a valid quest that there shouldn't be a settled set of profiles for healthcare.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So –

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And then the second question is a broader question, which was that the OAuth community was roiled by internal descent for a while and I'm just curious have they settled that and is it now a balloted standard or is it on track to be a balloted standard or is going to blow up again?

Justin Richer, MS – Lead Technologist – MITRE Corporation

All right, so I will answer the second question first in that OAuth 2 is now the IETF RFC 6749 and 6750, and yes I have those memorized. And so those are actual, honest to goodness, full on IETF standards up there with IT and TCP, and HTTP for that matter. So, yes they are real honest to goodness solid, stable, in use widely deployed, widely used standards.

The OAuth 2 – so what you're referring to was a very unfortunate incident where our editor did not like the direction that the spec had been going and felt that power had been rested out of his hands and quite honestly he had to kind of stomp out of the room and throw a fit on his way out and slam the door, which was the blog post that is now somewhat infamous.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right.

Justin Richer, MS – Lead Technologist – MITRE Corporation

The thing is though, there was a very large consensus of the group of very many of us already that the spec was already pretty much finished by the time that he did that and if you'd like some time, you know, give me a call and I can systematically go through all of the complaints that he raises in that blog post and show you like what either the community has done since then to address them or where he was a little bit, you know, off in describing what was actually happening.

And so the one that I will specifically call out is that the OAuth specification shipped only with what are called bearer tokens and Eran, our editor, really didn't like bearer tokens and thought that they were this horribly useless insecure thing for the web, it's honestly not the case, but that's neither here nor there, the truth is that we have non-bearer tokens in a specification that Eran decided not to edit anymore and no longer contribute to the conversation. So, there is a lot more to the history and to the political dynamics of this group that I would be more than happy to go into but not during this call.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

So, let me ask you one question though that is a perception question it's not a technical detail question, but as I recall his blog post or some of his concerns where that OAuth 2 had, I'll use his language not my own, but something like degraded or devolved to being no better than the cumbersome enterprise WSI style standards that it was intended to replace. You obviously don't believe that's true that it is in fact web appropriate and simple enough for use in widespread environments?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah, absolutely, it is astoundingly simple and honestly we have had so many people from so many different groups hammering on it over the 3 plus years that it was in the standardization process that – and this was all done in the open unlike the WS Star work that had largely sort of jumped full form out of a particular enterprise vendor who shall remain nameless. There were so many people hammering on this and so many people actually implementing it that we knew that there was really something there and so, yeah, I think that he's just – it's just wrong and it's really not a good way to categorize things.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Now what is true about OAuth is that OAuth 2 is at its core a framework for how to do delegated authorization and that was a design decision that Eran never particularly cared for either, although he was one of the instigators of it, which confuses me as well, but one of the key things is that unlike OAuth 1 we split out the notions of how you get a token from a server and how you use the token at a server and we explicitly decided to define multiple ways to do each.

The reason for this was not because we wanted to have something that was, you know, just like existing enterprise systems that we could just slap a new label on and call it OAuth not in the slightest, the reason for this was actually deployment experience with OAuth 1.0, which it solved everything in a single sort of monolithic protocol that was never quite the right fit for any of the various use cases that it was being thrown at.

If you used OAuth 1.0 for just authorizing one website to another you still had to deal with the request tokens which makes no sense in that use case. If you were using it to have a native client talk to a web server you had to deal with the client secret which makes no sense in that use case. There were all sorts of things that were added into OAuth 1 in order to support different use cases that just did not make sense.

So, what we did with OAuth 2 was very explicitly call out the fact that there are multiple different ways that you could use to get a token and no matter how you got that token you had a handful of ways that you could use to present that token.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And this is where the profiles come into play which is to my other question which is, is there a healthcare profile or does that just need to be deferred until you have a use case in mind?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Precisely where the profiles come in, so we have defined, like I mentioned previously, we have defined profiles for – actually for both the RHEX Project and the Blue Button Plus/Pull Project we have profiles, OAuth 2 profiles defined for both of those and are those the only profiles that should be used in healthcare, no I don't think so, but I think they're a great place to start and I think that by appropriately describing what the profiles are good for and sort of how you can actually use them you can really start to make some sense out of it.

So, for example in the RHEX profile we do use bearer tokens, we do call out that this is the token type that we're using, but our bearer tokens have to be formatted in a particular way using the JSON web token and JSON object signing and encryption order JOSE suite of specifications to format the token itself. OAuth 2 doesn't care at all what is inside the token it's just an opaque string it could be a random blob of x digits for all it cares.

But for our profile and for parts of the Blue Button profile as well we said, no this token must be this particular format, it must have these particular properties, it must be signed using this strength of encryption or better using these encryption algorithms or better and going forward from there.

And I think it's important that we do profile things in such a way that says that developers and users, and people that are deploying this have a very simple recipe that they can follow, that if you're doing everything at a certain level and you're – or however you want to stratify it then, you know, support these OAuth flows for how you get a token, support these flows for how you use a token, support these things for what is your token for example even going as far as to define things like the rights and scopes, and resource owners for each particular token and each particular kind of resource.

So, that's one thing that Blue Button is doing that RHEX didn't and that's defining what we're calling dynamic and structured scopes around all of these tokens that allow you to define in a way that's predictable and repeatable what the token is actually good for once it gets to the resource that it's protecting.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

This is Wes, I wonder if I might ask some questions on a somewhat more prosaic level?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Sure.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

So, early on or a while earlier you said OAuth 2 is now widely used can you characterize that wide use so we can understand, you know, what wide is, wide is effectively a comparative word. And then I'll ask my other question.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Sure, have you heard of Google, AOL, Yahoo those are really kind of the three biggest ones, Facebook uses something that is very, very close to standard OAuth but they sell off the standardization bandwagon just before it crossed the finish line, so they are slightly out of compliance.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

I don't know that sarcasm is necessary, but could you describe –

Justin Richer, MS – Lead Technologist – MITRE Corporation

I'm not being sarcastic I'm asking if you have heard of all of these companies and I'm assuming that you have they are all using it very deeply and vary widely in all of their protected APIs both between services that they run, between services that they expose to other companies, services that they expose to their end-users and to their developer populations.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Okay, so you're saying Google, AOL and who was the third?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Google, AOL, Yahoo, E-Bay, E-Trade, PayPal it's pretty much –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Okay, all right –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Name an Internet company and you'd be hard pressed to find them not using OAuth, Twitter and then –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

I'm asking specifically about OAuth 2 now.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes and that's who I'm speaking of.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Okay, so how did that usage roll out so quickly given the relatively recent completion of OAuth 2?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Most of the companies were tracking the standard as it was being finalized and their input from their early deployments was actually incorporated back into the standard itself. So, the OAuth standard has been stable from a technological stand-point a good 9 months before it was an IETF final RFC. That last 9 months was a lot of just editorial changes and clarification of what was really meant when you said that the client must do this or a server shall do this.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Okay, thank you and in the case of very large server organizations working with very numerous probably smaller client organizations certification boils down to it works with the server or not, right? But there is no industry-wide certification or anything like that for OAuth 2 is that correct?

Justin Richer, MS – Lead Technologist – MITRE Corporation

That's correct although MIT has recently announced an interoperability testing that they're going to be running this fall as part of the MIT Kerberos and Internet Trust Consortium Conference for different implementers of OAuth 2 to literally get together in the same room and start throwing code at each other.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, we think those are important progress. The third of my two questions now is one of the things that you said was that the work done to date has been put into the public domain; I think if there is to be widespread adoption in healthcare there needs to be some ongoing coordination and evangelization, and so forth where do you expect that would come from?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, the public domain is not quite the right classification but I'm probably getting too deep into the semantics of that phrasing –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Open source, open source, you actually said open source.

Justin Richer, MS – Lead Technologist – MITRE Corporation

An open standard, actually, so the standard is owned and controlled by the IETF, the Internet Engineering Task Force.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

All right, so let's just focus on healthcare profiles now as opposed to the standard itself.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Okay, so if you're focusing on healthcare profiles it's whoever writes and publishes the profile.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

That's correct.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, we did that underneath the auspices of ONC and FHA, and we published them as part of the readout of RHEX last fall and they should be available on the S&I Framework wiki pages as I recall, also on the Project RHEX repositories.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And, Wes, this is David, that's where I was headed with my questions is, you know, is there a set of healthcare profiles and in particular is there a set that maps reasonably well to the way healthcare has tended to use SOAP in the past, is that a good idea, a bad idea, in other words should the bearer tokens for healthcare be SAML assertions, I mean, that's where I was headed with this question as well.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah and I guess I'm headed in a little different direction which is that when it comes to for example Direct into HISP transactions the open source probe and associated intellectual property was pretty complete it didn't necessarily require a lot of interpretation of other standards in order to get a working implementation. I'm not sure that I'm getting the same vibes around these healthcare protocols for OAuth 2 and it really has just sort of the typical Internet speed roll out question of how complex do you make the implementation for particularly the smaller sized players and stakeholders in the game?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, that's actually a really great question and that was one of the driving forces behind the RHEX Project itself. So, as I said back on the focus of the different projects slide RHEX was not just about sort of talking about how it could be done, we actually built the system as open source so that a very compelling answer for how you do this is you run this server that makes all of the decisions for you. The profiles were codifying all of the decisions that we made in the build so that if you wanted to go in and build it from scratch you could still come up with something that still talked to the server or to the client.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah and I understand and appreciate that, what I'm concerned about and probably the question is better directed to ONC, but what turns that work that you did in New England into a living body of work as opposed to a body of work that did exist in open source?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Well, I think, you know, Wes, this is David, you know, my take on that is something that we need to wrestle with as the Power Team as a whole which is in the past the IHE style approach was to rigidly define complete profiles, I mean, that's what they are is a profiling body, that were supposedly so rigidly and sufficiently defined that anybody could quickly implement them.

In this Internet world of lighter weight building blocks and what we heard from Lloyd around FHIR is that the profiling is a much more fluid and dynamic process and that there may in fact be profiles that eventually become so stable and so widely used that they need to be, you know, memorialized in some official way, but otherwise there is a lot more dynamic construction of these things and the three projects that we're looking at here kind of typify that to some degree, they are using the same building blocks but they're using them in quite different ways none of which would have fit an IHE profile of your and maybe that's just the nature of the beast and we have to be comfortable with that in healthcare.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

You know, Justin you have a lot of slides in here for this technical deep dive that some of which might answer some of these questions that are coming up.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Quite possibly, I'd be happy to go through those.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Yes, I think that might be useful. I know there are a lot here, so we might want to try to go through them rather quickly and also we ask –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Okay.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

What we really asked originally was for you to describe the differences in how OAuth was implemented in the three protocols and you've mentioned one difference between the RHEX and Blue Button Plus but if you see opportunities to point out such differences we'd appreciate you doing that.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Will do.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And one question though before you leave your last statement, I agree let's go with what Dixie suggested, you talked about MITRE's efforts in creating OAuth to create a service that people could use is the code available if they just want to integrate the code themselves or do they have to use an appliance to run it?

Justin Richer, MS – Lead Technologist – MITRE Corporation

You mean RHEX or OAuth?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, yeah, RHEX, the RHEX use of OAuth.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Oh, the RHEX use of OAuth, yeah, the code is available.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

It's all open source under Apache license and it's available on GitHub and it has been since the project started.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, that was my impression and I thought you were saying something different which had me worried; no I'm glad to hear that, thank you.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes and an interesting thing just to follow-up with this, on the Blue Button Plus/Pull side we are currently actually also building what's loosely being called a reference implementation of Blue Button Plus, different components of Blue Button Plus not to necessarily give people a project that they can sort of drop in and deploy although they could do that if they like but more to show people how it could be built.

And what I think would be very helpful to the larger community on the Blue Button Plus side is that if we had a dedicated engineering team to build an open source freely available, liberally licensed reference implementation for that, you know, just pick a platform, build it and go I think that would actually be very useful to the wider community, because we've seen how much having a – like a solid engineering team behind RHEX actually helped move sort of this whole notion of how you do things forward.

Because I can tell you for a fact that a lot of the conversations that we've had around Blue Button Plus have kind of knotted back towards, well, okay, so on RHEX we – when I was in the conversation or MITRE when I wasn't there kind of did it this way so, you know, what does that mean, so how should we look at it going forward.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, I agree with reference implementations is the new profile.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah and to put some light on what some people might not understand about the IETF process that gave us the OAuth 2 standard, the IETF model for acceptance of a standard is different from a lot of places, they like to say it's rough consensus and running code.

So, if you can show up and say, I think it would be a really great idea if the protocol did this and by the way over the weekend I wrote it and it's actually implementable that way, that's going to carry a lot more weight than somebody saying, well, I don't think that that's a good idea because I just don't feel like it.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, that was the motive –

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

I really think we need to get on, because we do need to have a discussion about these comparisons among the three protocols a little later and Justin has a lot of – here.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah, Dixie, don't worry about the number of slides, I speak very quickly.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

And I meant for these all to go through very quickly, so don't worry about that.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

– roughly half the presentation –

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Try to get all the ones that address questions we've brought up.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Okay.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

That would be good.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, what I'm going to do now is give a technical overview of both OAuth 2 and OpenID Connect and these were both called from larger presentations that I've given in other audiences of how these protocols actually work on the wire so I'm not going to explicitly be talking about sort of the differences and subtleties of how this applies to healthcare because I believe that the fundamentals are the same. As soon as you decide that I'm doing an HTTP RESTful protocol and I want to secure it this is really where you start. So, next slide, please.

All right, so as I mentioned before if you've one of these pop up authorizations, next slide, please, saying that this App wants to go and use stuff then you've used OAuth before, it is everywhere and next slide. It's important to remember when you're talking about OAuth that it's an authorization protocol and it really is a framework for building authorization protocols more than anything when we're talking about OAuth 2, it's made to go on top of HTTP and it's really, it's tied very nicely to that protocol stack but it's also mobile friendly and REST friendly even though OAuth 2, and this is one of the criticisms that people have leveraged, OAuth 2 is not itself a RESTful protocol, it was never meant to be fully RESTful, it was meant to be very friendly to using it with RESTful protocols and we in the IETF felt that was much more important. Next slide, please.

So, these are the key players that we'll be seeing as we go through stuff. We've got the resource owner who actually controls things they're usually interacting through a user agent, some kind of a web browser, they have a protected resource that their client wants to get the stuff that they control and then there is an authorization server which gives tokens to connect all of these together and then we have different kinds of tokens for different kinds of use cases, but as far as OAuth is concerned a token is just an opaque thing that you just hand across. Next slide, please.

So, the real question is, how do we actually connect these together and really which ones that we want to connect. Next slide. Because ultimately we're trying to get just a couple of components to talk to each other. There are many different ways to get OAuth tokens from one side to another but we're going to talk about the top one for the most part today and that's the authorization code and this is one of the things that all of the profiles of OAuth that we've been talking about so far have done.

They've explicitly said if you are doing say an end-user to end-user and you have this kind of client use the authorization code flow. In Blue Button there is a use case for the implicit code flow, there are Apps, you know, in browser temporary Apps that can use that. On RHEX we actually have use case for the client credentials flow when there is literally no user involved what so ever and all of the trust decisions are made by policy that's been previously set up, but today we're going to talk about the Auth code flow. Next slide.

And there are different ways that you can do stuff even more; we're not going to talk about these today. The point here is, next slide, that OAuth 2 is actually very flexible. So, we're going to talk about the Auth code flow. Next slide.

The players that we primarily care about here are the resource owner which in our case is going to be sitting inside of a web browser, we're just going to simplify things, keep in mind that this could be on a mobile application or, you know, there are different ways that you could deploy this in all the different parts still look pretty much the same, but assume that the user is inside of a web browser, we have the client that is trying to get stuff from the protected resource via the authorization server. Next slide.

So, our goal and what you actually want to remember is that we're trying to connect the client, over on the left, to the protected resource on the right. Next slide. So, the way that this starts out is that the end-user initiates some action on the client that says, hey go get my data it's over there, now OAuth doesn't care how that starts, how the user gets authenticated to the client itself, how that session gets set up nor does it care how the client figures out where the authorization server and protected resource is. OAuth literally does not care about that. Now our protocols such as RHEX and Blue Button Plus do care about that.

And for example FHIR would define what the protocol for actually getting data from the protected resource is. OAuth doesn't care because it's running at a different layer. Now – I did not mean the ISO 7 layer network stack but it is at a different application layer and so, next slide please.

Once the user has initiated something the client sends the user over to the authorization server. Next slide. And once the user shows up there they authenticate directly with the authorization server, now this is where the users actual credentials come into play, so they either log in with a user name and password or they have a certificate or there is some other kind of magic that lets the user in the front door at the authorization server not at the client, keep this in mind where this key is being passed, because that's important later on. Next slide, please.

The user tells the authorization server – so the authorization server has at this point generally asked the user, hey this client says that it's trying to get your stuff do you really mean for that to happen, and the user says, yeah, that's cool. Next slide, please. So, the authorization server mints a temporary credential what's called an authorization code and I'm representing it by a paper ticket here that sort of captures the notion that the user was there asking on behalf of his client to access – for permission to access the protected resource, so that's all kind of bundled into this notion of what the auth code is. Next slide, please.

The Auth server then re-directs the user back to the client, once again this is an HTTP-based flow, so the user shows up at the client with this temporary credential and then hands it to the client. So, think back what happened here, the client told the user, hey I want to go get stuff, get me something that I can use to go do this action on your behalf, because I can't get it on my own. Next slide, please. And that's that ticket.

So, the client uses that ticket and its own credentials to talk directly to the authorization server. Now, you'll notice here that this does not, does not go through the web browser and that's very important, because the client has its own set of secret credentials that go across this back channel we're calling it between the client and the authorization server directly. Notice that the client also does not and has not yet, and will not ever have access to the users credentials that the user used to log into the authorization server and we haven't even gotten the protected resource involved yet. So, next slide, please.

The authorization server looks and it validates that the client is who they say they are because of those credentials, it validates that ticket was issued to that client and then the server can then mint a couple of tokens that it can hand back to the client, next slide, please, that represent that authorization decision actually being made. So, now here's where the cool stuff really sort of hits the road, next slide.

The client can then take those tokens and use them at the protected resource to get things. At this point the user doesn't even necessarily have to be present anymore because what's fundamentally happening here is the user is delegating authorization to the client to access the protected resource on the user's behalf. What's important here is that the user is able to do that without exposing the user's credentials to the client and also giving the client this very limited, very traceable token that can be scoped to a single resource for a single time or a single, you know, very small window, you can limit what that token can do very, very easily and all of that power comes from the fact that you are separating the three different credentials from each other.

There is the credential of the user logging into the authorization server, that's one level of authentication. There is the credential of the client logging into the authorization server when it goes and hands that ticket and the auth server comes back with a token that's another level of authentication. So, we have authenticated the user and we've authenticated the client, we've done so in different context and tied them together and now we have used that to convey an authorization decision which is bound into this token that the client then hands to the protected resource and the protected resource then, next slide, decides what that's good for.

So OAuth is very, very good for avoiding asking users, oh, what's your password over at that protected resource and I promise I'll be really good with it, you know, I'm not going to steal your password, you don't even have to think in terms like that anymore because getting access to a particular end-user's information over on a protected resource is now no longer directly tied to that user's identity on that resource even. You don't care how the user logged in when they got to the auth server you just care that they did and that you can tie it back to your session and tie it back to the right user locally and then move forward.

So, there are literally – I gave the list earlier of some of the biggest sites and the biggest providers all of them have many, many API end points that are protected with OAuth 2, but if you go to a completely blanking on the name, but there are sites on the web that actually list API end points, and Sam maybe jump in if you can remember the name of this, I want to say MetaFilter but I know that that's wrong, that list how many hundreds of thousands and thousands of different sites are using OAuth 2 today and it's growing.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Justin?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Could you go back to the previous slide of your players and just make sure that I've got it in my head correct, I think I've got it, but –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Sure, previous slide, please.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Previous slide and let's assume I'm running TweetDeck and I want to give it access to my Twitter account can you map where – who's me, who's TweetDeck, who's Twitter?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Absolutely. So, you are the little dude in the blue shirt your TweetDeck is the client and Twitter in this case is both the authorization server and the protected resource.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And that's commonly the case right? Those two are often times lumped together even though they're not required to be?

Justin Richer, MS – Lead Technologist – MITRE Corporation

Exactly and this is one of the things that the OAuth 2 sort of data model and protocol model changed from OAuth 1. In OAuth 1 the auth server and protected resource kind of had to be the same box, as it turned out people weren't always deploying it like that, Google being one of the biggest cases where they didn't deploy that, but that was really, really hard to do and Google had to jump through some pretty crazy hoops to make that happen with OAuth 1.

With OAuth 2, OAuth 2 doesn't define directly how the protected resource validates that token is any good and as it turns out you've got a couple of really good ways that you can do that, one of the ways that we've done with RHEX is to define that the token itself is a signed JSON web token that the protected resource can then go and look up the appropriate keys and check the signatures, and check the expirations, and all of that stuff there.

With Blue Button Plus we're allowing for not only that but also a method called token introspection in caveat emptor I am the editor of the token introspection draft in the IETF as well and that makes sense in the Blue Button Plus world because it's a much more dynamic world and you might now know who your authorization server is until somebody shows up with a token from them, but you might want to still actually trust them.

So, you're absolutely right that the most common case is for the AS and the PR to be in the same box and checking that token is literally coking the database to see if there is a database row that has that token value in it. I've written that code a dozen times, but that's not the only way you have to do it.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

So, just to follow-up to that question, the user is the guy with the blue shirt, the client is the fragmentary cubes and what is the rectangle that says "Google" in the middle? How is that distinct from the user?

Justin Richer, MS – Lead Technologist – MITRE Corporation

That's the user's web browser so that's a screen shot of the Chrome web browser, so it's the user –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, all right.

Justin Richer, MS – Lead Technologist – MITRE Corporation

And the user agent. So –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

So the user agent is not the client.

Justin Richer, MS – Lead Technologist – MITRE Corporation

That's correct.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Then this is used in circumstances where there is no visible browser, I'm just having a little trouble understanding what that – is it that this represents some presentation part of the client or how does it – I'm sorry I'm just confused.

Justin Richer, MS – Lead Technologist – MITRE Corporation

It is the user's browser and –

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

No, it's whatever is interacting with the user typically a browser, it could be a piece of software different from a browser, it could be an EHR.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes, that's absolutely true, what's important though is that it be able to present the authorization page at the authorization server and be able to handle the re-direct, which is going to be an HTTP 302 re-direct in such a way that the client can actually get that code back.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Good point.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, so for example if it's an EHR it's likely that the EHR and both the browser can do that part of the user interaction or something like that?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah and –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah, that's how most native clients actually work, so like I said at the beginning this is one particular use case.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Where the user is actually in this case interacting with the client through the browser as well. So, to go back more to the concrete TweetDeck example though what TweetDeck would do is they would actually spawn your system browser using a system call and this is very, very common on the mobile platform, to say, open this HTTP URL that I'm going to give you. That HTTP URL is going to go to the off server in whatever browser you happen to have and then the user is going to log in, they're going to authorize, they're going to do all of the stuff like that and then the auth server follows the re-direct URL which has more than likely been preregistered and tied to that particular client, you know, I've glossed over a lot of the – sort of the fine details of the protocol.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yes.

Justin Richer, MS – Lead Technologist – MITRE Corporation

But the auth server is going to send the user back to the re-direct URL with that auth code as one of the query parameters and so then it's up to the client to kind of have a means of listening for that call back URL. Now if the client is already a website that's pretty easy. If the client is a – you've got actually a few ways that you could do it, you could either run a web server on local host, which I've seen people do, it's very useful, you can do what we did with the hReader Project and actually register a local handler, I forget what they're called on IOS, but on Android they're called intent handlers, where you say, like, you know, hReader://O is opens my application.

And so your call back URL actually has hReader:// instead of, you know, HTTP://. So, when the system browser gets that URL it hands it to the system and says, I don't know how to handle this, you go hand it to some application that does.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Got, it, thank you very much.

Justin Richer, MS – Lead Technologist – MITRE Corporation

And there are a couple of other techniques but that's basically what's going on.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

But this is one interesting difference that I hadn't thought about between say SOAP and OAuth in that it really – the OAuth model really is bound pretty tightly to the HTTP world and the assumptions around things like re-directs and –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Absolutely.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

URLs, so, that makes great sense but it is more narrowly scoped than some of the earlier profiles of distributed authentication – authorization I mean.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes, I'll say that on a sort of philosophical level I completely agree with you, but I have yet to see SOAP actually work on anything but HTTP.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, yeah it's a totally philosophical point.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Exactly.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Irrelevant to – I mean, think this is what works in the world that we're in now.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Right and again the standardization process in the IETF is rough consensus and running code, this is what people built and so we standardize what people were building.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Well, earlier Justin you mentioned that RHEX uses bearer tokens and Blue Button Plus doesn't I think, is this token that we're looking at a bearer token or is there – I mean, this looks like what I think of as a bearer token, but is there something else that – is there a difference?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, it doesn't actually matter. So, the means by which you get a token doesn't care what the token actually is and so what I'm showing here could be a bearer token, it could be the signed request MAC token, it could be a JOSE-based holder of key token. There have been a couple of other proposals. The only one that's been standardized and finalized is the bearer token and so given both that and sort of the security needs for the RHEX Project that's what we standardized on. Now Blue Button I didn't say that it doesn't use bearer tokens I said that it doesn't standardize on a particular token type.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Oh, I see, yeah, yeah.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Because with Blue Button you've got a slightly different model of who your authorization server and who your protected resource is and in Blue Button we're actually sort of bundling those two loosely together in what we call a provider. In RHEX we were a little bit looser about how we actually connected those two and so the protected resource it's going to get a token and that token is going to be signed, it's going to have a particular parsable format and all of this other stuff that would tell the protected resource what to do with that token.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Thank you.

Justin Richer, MS – Lead Technologist – MITRE Corporation

I do think that in terms of building a widely usable healthcare profile, I think, you know, picking token types and picking token formats is a really great first or second step.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, I think that's going to be Dixie – I'm, you know, tipping my hand a little bit, I think that maybe one of our recommendations to come out of this is that we run the risk with the flexibility here of lots of different healthcare variations that could blunt the benefit of the fact that they're all using the same standard. So, whereas – you know, we don't want to make the mistake of over profiling and being overly constrained we might want to identify really common use case patterns and say these are appropriate profiles to use –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, that's –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Exactly and I was just going to say – oh, go ahead, Wes.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

That's why I made the comment about enforcement coming from sort of 800 pound gorilla versus 800 screaming monkeys kind of business choices and somehow we have to deal with the propriety of that kind of dominance in the healthcare world.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah and it's – yeah, it's a policy decision more – it's not a technology decision. The good news is that technology is flexible, the bad news is the technology is flexible.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, somebody is going to flex it right?

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, somebody is going to flex it.

Justin Richer, MS – Lead Technologist – MITRE Corporation

And that's been one of the ongoing and honestly perfectly valid criticisms of OAuth, some people wanted OAuth to be exactly their use case and nothing else, and some people wanted it to be so loosely defined that you couldn't really do anything. I think that we've come down with the IETF in a pretty happy middle ground.

So, you know, how the client talks to the AS that's completely standardized. How the client talks to the protected resource that's completely standardized within the scope of what kind of token you're using and how the protected resource talks to the AS that hasn't been fully standardized yet, but, you know, it's going to be people needing to standardize that that's going to drive that and that conversation is – honestly, that is happening now.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah and it's sort of consistent with – but let's move on.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Okay. So, next slide, please. And so one of the things that OAuth has interestingly been used for a lot on the web is sort of authentication and sign-on protocols. So, when you collect Facebook Connect or log onto a Twitter or signing with Google you're actually do OAuth under the hood. So, next slide, please. This leads to an interesting conundrum. So, is OAuth an authentication, sign-on, log in protocol? Next slide, please. It's absolutely not. Next slide. It really, really isn't.

So, I'd like to bring up a wonderful metaphor that I blatantly stole from the blog that's linked at the bottom of the screen here of chocolate versus fudge. Next slide, please. OAuth is chocolate it's great on its own, you can just eat it directly and it's a wonderful thing, but it's also a very, very useful ingredient. You can use chocolate in lots of, lots of different recipes, lots of different things and it can even be used to make some really good fudge. Next slide.

When you're doing sign-on you're making fudge, it's a confection, it's made with several different ingredients and, you know, one of those ingredients could be chocolate but if you're making fudge you need to have more than just chocolate even if you're making chocolate fudge and if you're making fudge you don't have to use chocolate at all you could use peanut butter, you could use potatoes, I've actually had potato fudge it's surprisingly good, next slide, please.

So, what we really need is the recipe for fudge that uses chocolate, in other words a recipe for how to build sign-on with OAuth 2. So, how about we create an identity API and we make that identity API the protected resource and we're going to standardize all of the user profiles, your name and e-mail and picture and everything like that. We add in important stuff like session management and levels of authentication and levels of access but still keep compatibility with very basic run of the mill OAuth 2. Next slide, please. That's a great idea. Next slide.

So, why isn't anybody working on that? Next slide, please. It turns out they are. That is the genesis and the driving force of OpenID Connect which is building a distributed identity system on top of OAuth 2 at Internet scale. Next slide, please. So, what we're really doing – so some basic background of Connect, it's a new identity protocol that's being built right now on top of OAuth 2 but based on experience with OpenID and OAuth, and SAML, and Facebook Connect and a bunch of other things over the last, you know, decade and more of knowledge and sort of known quirks and weirdness of deployment is going into OpenID Connect, it's being developed by the OpenID Foundation you can see it there and I will say right here it is actually currently in implementer's draft due to be final by, hopefully by the end of the summer. So, it is actually technologically stabilized at this point. Next slide, please.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Well, I thought Google used, this is one point I've been confused about here, I thought Google used OAuth 2 for both authentication and authorization is what they're really using OpenID Connect?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, that is a very, very interesting question. So, what's really happening is that they have been slowly adopting components of OpenID Connect as they become more and more stable. So, it's Google's goal to be able to say that Google is fully OpenID Connect compliant as is the goal with, you know, Microsoft's Azure Systems and E-Bay's log in systems and a bunch of other providers, but Google does not want to come out and say that, and sort of stake that flag into the ground until OpenID Connect is a final standard and it is not yet a final standard. So, like I was saying before though it's disingenuous to say that they are doing authentication with OAuth, it's that they are using OAuth in an authentication protocol.

Dixie B. Baker, MS, PhD – Senior Partner – Martin, Blanck & Associates

Okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, you need to have other things added to OAuth in order for OAuth to be an authentication protocol, it's really, really not on its own and again this goes back to the chocolate versus fudge metaphor and that's that if you have just plain old OAuth and if somebody comes to you expecting just plain old OAuth you can delegate an authorization but you cannot get authentication from that you need to have other things in place, because OAuth doesn't tell you anything on its own about the user that logged in, it doesn't even give you a distinct user identifier, it doesn't even give you any information about the server that you're talking to, all very, very basic building blocks when you need to do authentication.

So, what OpenID Connect does is it starts with OAuth as we see here and JSON web tokens and a bunch of other stuff and it builds on top of that using the motto making the simple things simple. So if I just need to log in a user that should be really easy for me to program, especially as a client, but make the difficult things like stepping up to higher levels of assurance and doing encrypted, and signed requests, and all of that other fun stuff possible without breaking basic compatibility. Next slide, please.

And so, what you do inside of OpenID Connect and this is what Google's login with Google does today is you use plain old OAuth 2 and you get a regular OAuth 2 access token, the access token that we saw previously, but you also get this thing called an ID token. Now the access token in regular OAuth doesn't have any particular format or content or anything and the audience of the access token, so who is supposed to actually be interpreting it, is the protected resource.

The ID token on the other hand is a different animal entirely, so the ID token has a defined format, it is a signed JSON web token and it has a defined content structure so that you can look at particular what are called claims inside the ID token, this should start to feel a lot like a SAML assertion to those on the call who are familiar with SAML and there is good reason for that, because what we've done with OpenID Connect is kind of take the best parts of SAML and re-adapted them in a way that makes it much, much easier to use.

So, you have the ID token to manage your session and sort of your current user information and figure out like who just came in through the front door because the audience of the ID token is the client itself. So, the client is meant to look inside that token and figure out who is there, who did they just authorize and use that for authentication purposes. Now that access token can then be used to go off to what's called the user info end point and get back a user profile, because it turns out and as we learned with OpenID 2.0 for many years just authorizing or just authenticating the user isn't enough. So, authenticating the user with a globally unique identifier seems like that that's all that anybody would want to do when you're working at the security protocol level.

But it turns out that people want to do silly things like say, you know, know the user's name so that they know what to print on the screen, be able to access the user's e-mail address so that they can send them messages, you know, things like that were the first things to be added to OpenID 2.0 as extensions, so that has all been built into OpenID Connect and it uses the OAuth model in the following way.

So, the user authorizes the client that they're – so the user is logging into the client, okay, the user authorizes that client to go and get log in and profile information about the user from the protected resource and that protected resource is the identity API. So, next slide, please.

So, I think this is the picture, no, we'll get to my picture in a bit, of what all those different components are, because you can do some other interesting stuff if you really get into sort of the more complex side of OpenID Connect and it really can get down into the very bottom it's SAML with JSON's curly braces instead of XML's angle brackets if you wanted to. Now the key difference between this and SAML is that with SAML you kind of start way up the complexity curve to even do a simple operation. With OpenID Connect you start way down the complexity curve and it only gets complex once you start to do kind of the more crazy stuff. So, next slide, please.

OpenID Connect has been very useful in getting OAuth 2 and related components actually used and a lot of what OpenID Connect had to invent has been pushed back to the wider community, next slide, please, as well as the interoperability testing. Oh, apparently that slide also got cut.

I think that we must have had a version mismatch of the slide decks, but, so what I was going to point out is that if you remember back to the slide that had the four different pieces on it, so the user and their web browser is still the same, the client that was down in the lower left-hand side that was accessing the protected resource, that turns into the site that you're trying to log into what's called the relying party in identity management parlance.

And the identity provider turns into the two components on the right hand side, the authorization server and the protected resource. The authorization server is what's able to communicate the authentication information and the protected resource communicates the profile information that you have been delegated access to. So, Sam, this is sort of a summary slide for RHEX as a whole, I'd be happy to take more questions on the specifics of OAuth and OpenID Connect here before we wrap up.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

This is David, one question on OpenID Connect how does it compare to Mozilla Persona?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, Mozilla Persona is a really interesting project that I'm not really sure where that's going to be going, because they're using sort of the ID token structure from OpenID Connect but they're bundling it into the browser itself and there is a lot of sort of distribution at scale problems that I don't think it actually solves.

And what I think needs to happen is that the Persona folks and the OpenID Connect folks really need to communicate more, because Connect has a notion of a self-issued identity which is really what Persona is getting at, that Persona could even directly use if not directly influence how it works and how it's processed. I know that the two groups are aware of each other's efforts, but I will say that nobody is directly talking right now.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay and then one other question that's a slightly different angle, where I work at Cerner we are heavy users of OAuth 1 and of OpenID Connect or just OpenID 1, is there some reason why we should change or will those be considered perfectly applicable and useful protocols if they meet your needs? What happens to those early adopters?

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, those are still valid if you still have a good reason to use them. So, the advice that I've been giving to people is that if your code already speaks OAuth 1 or OpenID 2.0 just use that, but if you're building a new system and especially if you're designing a new protocol stack like the ones we're talking about today, definitely, definitely go with the new ones. There are some security benefits for going with the newer protocols, there are some very, very key architecture benefits, so I'm sure that you've noticed in your systems that getting OpenID 2.0 and OAuth 1.0 to talk to each other they're very, very different animals they really don't like to get along.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, yeah, they don't mix in our world so we've gotten away with it but yeah I can see how that would be problem.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Exactly, so what we found with Connect is that because it's based on top of OAuth 2 we can actually use the same authorization system for doing both OpenID Connect transactions and vanilla OAuth 2 transactions and this is actually the basis for the software that's bundled with RHex to handle authorization services and what's being used in some of the initial OAuth or Blue Button Plus reference implementation work that Josh Mandel and others have been working on is that you can do both with it.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Good, Wes did you have a question there I think I heard you?

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Well, I might have, he might have answered it. What I think I hear is that if we adopt, you know, push as Meaningful Use certification standards something based on OAuth 2 then are we in effect requiring vendors to switch to OAuth 2 for their existing authorization usage or just to be able to support OAuth 2 in a health information exchange scenario?

Justin Richer, MS – Lead Technologist – MITRE Corporation

It would really be the latter because the compliance would be with the profile that mandates OAuth 2 and particular sort of flows in components of OAuth 2 and so the question goes back to the previous question, if I may rephrase it, you know, who provides the teeth to these standards, you know, who is actually holding people up to the compliance levels. So, if a vendor wants to ship something that does OAuth 1 and OAuth 2, which I've done and I've got systems here at MITRE that I've worked on that do both, then they're still free to do that.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah, see I – this is David again, I think we've – you know, we are unfortunately given a somewhat vague, you know, mandate from ONC and it's a pretty complex space so you combine a vague mandate in a complex space and there are a lot of different ways to answer questions, but it seems to me that these, you know, OAuth 2, OpenID Connect which is a specialized use of OAuth 2 and FHIR, and hData or maybe let's just focus on FHIR those are all sort of lower level building blocks that can be combined into more interesting and more complex services such as Blue Button, ABBI Plus, which is the combination of those building blocks and some other things that we haven't talked about like the registration service into a more complex profile that's really almost more like an application than a building block.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yes and I would –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

There is a valley, I'm trying to tiptoe around the edge and looking for a guiderail which is there has been discussion about the possibility that modules that are certified together for a certified EHR should have some standard interoperability, I'm trying to avoid us getting into a situation through this that we're telling module developers that they have to re-develop on a lot too.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

No, but what I think we could say, Wes, just hypothetically is we could say to a group that wants to build a new healthcare service let's just take a mobile health App, you know, and say, you know, you want to build something that becomes so widespread that ONC might at some day even consider, you know, building it into a required certification standard, we could say to that group it's safe and appropriate to use OAuth 2 and it's safe and appropriate to use to FHIR as building blocks for your new service, go experiment and when you've got something that you want to sell to the rest of the industry by "sell" you know make it a standard or actually sell it you should feel comfortable that you've picked good solid building blocks. Does that make sense?

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

I mean, it's, you know –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

No, I guess I'm basically looking for and this has nothing to do with the presenter unfortunately, but I'm looking for some sort of scoping of our recommendation in regards to is it about inter-enterprise interoperability or is it about certification of components of an EHR in the manner that they operate intra-enterprise and I recognize that in the long run there should be, you know, as few differences as possible. I'm just trying to avoid sort of backdoor mandates that require re-engineering of systems without specific – engineering.

Justin Richer, MS – Lead Technologist – MITRE Corporation

So, I can actually address that a little bit by bringing up the fact that when we've been piloting RHEX specifically the RHEX Project, the RHEX source code is not itself a full-fledged EHR really when you get right down to it, but it can be used to wrap and tie into EHRs such that your medical records with the appropriate RHEX adapter in front of them are now interoperable using all of the standards-based profiling that RHEX comes with when they're talking to something else that is interoperable.

So, when we went and deployed this over at the Maine HealthInfoNet we didn't make them throw out their EHR not at all. We didn't make them re-engineer their entire EHR. No, what we did was we built a component that could speak to their EHR using some, you know, direct database connections and some other kinds of stuff to actually get the data in and out, I forget the details specifically, but the important thing is that it talked to their EHR in some opaque manner and that the thing on the other end doesn't care at all how those internal connections were made, because the importance of interoperability are always where you have to cross the lines between the different domains.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

I think –

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

So, see this is I think the dilemma that we face is that we could come out of this, our recommendations and easily say, I think we could easily endorse OAuth 2, OpenID Connect and FHIR as being, you know, worthy of use in healthcare applications. To say that RHEX is a profile of those applications, of those lower level components that vendors should build into their products is a much different question, because now you're talking about an actual service capability that uses the building blocks because RHEX is much more close to XDS and, you know, it may in fact be better than XDS but it's a much different animal than saying OAuth 2 is a good building block.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Right, but –

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

And I think the same with Blue Button, Blue Button is a much different animal it just happens to be probably – since it's carving out a new space and doesn't have existing options similar to it, it's a little easier to say Blue Button Plus is a good idea we should go forward with it.

RHEX has a lot of similar things out there and to put the same push behind that is a more complicated decision. I like it a lot and I like what they did with it if they would switch to FHIR for example, but nonetheless it's closer to saying, oh, now you've got XDS, XDR, XCA and RHEX take your pick have we just made the world much more complicated.

Justin Richer, MS – Lead Technologist – MITRE Corporation

This is where I think having a, you know, going back to the if you're going to do this then you do it this way type of recipe-based profiling really, really comes into play and I think that this is a place where the FICAM profiles that have sort of started to be worked on I think that those need to be run to completion to give a baseline for a lot of this stuff and is Deb Bucci on the call by any chance?

Although actually people are muted I think so I'm not sure that can happen, but I know she and her group have been working on both an OpenID Connect and an OAuth profile influenced by the RHEX profiles, and I think that that's something that NWHIN and ONC and all of these other groups can really kind of start to get behind and say that at the very least you're not only just doing OAuth 2 which as we know that can mean a bunch of different things, you're doing this particular flavor of OAuth 2 because you fit this particular use case.

And so by doing that you not only give people sort of the right building blocks but you start to shape the building blocks so that they fit together in ways that actually make sense. Because ultimately interoperability is going to have to be a full stack thing.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right and that's my point, you know, again these things are – it's best to think of them as layers, you could say that OAuth 2 is your framework there are a couple of healthcare profiles that have seen wide use and we recommend that you pick from them if they meet your use case, but since it's a framework you could come up with something different if it doesn't meet your use case and then add all that up and you get a higher level service like say RHEX or ABBI, you know, the higher up you go the more variation there is, of course the harder it is for vendors to implement something that just works out of the box that's our –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Well –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

I'd just like to comment that one of the issues that we have is that we have a choice whether to recommend that ONC adopt a standard certifiable or not. If it's certifiable it's almost – it's very likely to be a complete service, I mean, I don't – certify something that uses OAuth 2 in some way in the – stack I don't think that's a reasonable approach.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right, but our task force yields with things that are at a lower level than are certified in and of themselves. So, I agree we wouldn't say let's promote OAuth 2 as something certifiable, but we –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, I think we've all been looking for building blocks that could be widely used.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

In various profiles, I think that – I guess I didn't quite understand the scope of this group, I should probably come to more meetings.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Well, it's been vague from ONC, it's been very vague.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, okay.

Justin Richer, MS – Lead Technologist – MITRE Corporation

And so I think another thing to keep in mind is that certification and compliance testing really does have to be against the application, but you only have to test the interoperable points that you care about. All vendors can –

Justin Richer, MS – Lead Technologist – MITRE Corporation

Lots of other different things.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Right, we know that part.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

We agree.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

So, I think one of the things that Dixie sent me an e-mail that she unfortunately had to drop off the call for an emergent thing that came up so we've got to land here in the next 8 minutes, so I think what Dixie wants us to do at the Standards Committee next week is to do a preliminary report out on sort of our assessment using our modeling tools that or assessment framework that we put together in last year's sessions on these three subjects and I think what – these three topics FHIR, RHEX and ABBI Plus Pull, and I think what we're going to have to do up front is to make it clear that these are apples and oranges, they're not all at the same level.

I think my opinion is that it would be nice to break it down and say the things they have in common we can make strong statements of support for, but that at the higher level weaving them together into services I think is going to be something along the lines of we think is worth furthering and piloting but they're not ready, they're not stable enough and widespread adoption enough to be considered ready for certification, that's my opinion, I shouldn't probably shouldn't –

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Give my opinion so strongly, but Wes does that make sense to you?

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Absolutely, it sounds like – do we have another meeting before we have to report out? I'm just wondering mechanically how to go to what Dixie wants to have ready for reporting out.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

I'm sorry, say your question again, Wes, I got distracted somebody stuck their head in the door?

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, that's fine, I understand that Dixie wants us to report out some work based on what we've heard on the last two calls and some assessment of it, I'm just asking about the mechanics for doing that. Do we have more phone calls? Do we have to do it without phone calls? That's basically my question.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yes, I think given that we've just got a few more minutes before – and we have to have some public comment time, is Dixie circulated to you and me a preliminary assessment on the previous two and then I think we'll need to complete one on this, on RHEX just a single sheet using our old framework where we score what we heard and I think we can do that via e-mail, because we only have a few days before the Standards Committee meeting. We don't have to be done we just need a preliminary report out and we can say here's our preliminary scoring and then if necessary we can have more open calls where we go into debate.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, I think – yeah, I think that sounds ideal because we would like to get more active involvement in our deliberations here and there's nothing like getting an interim score to get people interested.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

To get people interested, exactly. We can make a pitch to get people to show up for the meeting or to schedule them more further in advance so that – I mean, everybody's busy, but – so, speaking on Dixie's behalf and as the Co-Chair first Justin thank you for a really great presentation with a lot of detail in it and for being patient with our incessant interruptions, really helpful stuff. If you find the slide deck that had the missing slides in it and you want to send that to us feel free to do so, so that if you think there is something we should have seen on a couple of slides you missed, you know, we'd be happy to receive them.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Sure, I'll say we did cover the content but I'll make sure those get passed along. There wasn't much difference mostly just a couple of extra pictures.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay good.

Wes Rishel – Vice President & Distinguished Analyst – Gartner, Incorporated

Yeah, I think the pictures, you know, a picture is a worth 1000 words at least.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Yeah.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Yeah and I agree, I am very much a visual learner and explainer.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

We're visual folks I suspect and Ollie and whoever else spoke whose name I'm probably not remembering correctly thank you as well for setting this up, this is – you know, and for funding and supporting the research that led to these experiments they're really useful for the community to have this work done. We'll all benefit from it. MacKenzie should we open it up for comments in the last 5 minutes here?

Public Comment

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Yes, operator can you please open the lines for public comment?

Rebecca Armendariz – Altarum Institute

If you would like to make a public comment and you are listening via your computer speakers please dial 1-877-705-2976 and press *1 or if you're listening via your telephone you may press *1 at this time to be entered into the queue. We have no comment at this time.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Okay, well, thanks to our guests and to our sparse but loyal attendees and we'll adjourn this call and I assume there will be more discussion on these subjects after next weeks' Standards Committee meeting. Thank you everyone.

MacKenzie Robertson – Federal Advisory Committee Act Program Lead – Office of the National Coordinator

Thanks, everybody.

Samuel Sayer – Senior Software Systems Engineer – MITRE Corporation

Thank you.

Justin Richer, MS – Lead Technologist – MITRE Corporation

Thank you for having us.

David McCallie, Jr., MD – Vice President, Medical Informatics – Cerner Corporation

Bye.

Ollie Gray – Research Program Manager - TATRC

Thank you.