Di **[00:00:12]**: We have talked pretty extensively about identity providers and externalizing user authentication and the benefits of that and how you get that done. I'm hoping now that we can turn our attention to the next phase of the evolution, which is externalizing that authorization decision. In the XACML [eXtensible Access Control Markup Language] architecture, the components are the policy enforcement point [PEP] and the policy decision point [PDP]. Mike or Maury, could you tell us a little bit about CONNECT's experience implementing a PEP and a PDP?

[00:00:57]

Mike: Again, a lot of it is at the technical level and being able to both validate who wants to make a query against the system? Who wants to talk to one of our partner sites? Then, what it is that they actually want to do? And so, kind of what was mentioned earlier: on the one hand, if somebody wants to make a query, say from Nebraska against Alabama, Wyoming, Kansas, we have the centralized architecture. And, some of that centralized decision point needs to know who's talking to it.

And so it needs to know that Nebraska is a trusted partner in the first place. And also that user is a trusted user and has the rights, and the capabilities, and the correct credentials to be able to make the kind of query that it's trying to toss out. And the other state or the centralized portal can kind of monitor those and look at what's going on. So, by presenting credentials, by trying to launch a query, a couple things happen. One, we know that we need to trust the state, and we need to trust the user –

[00:01:58]

– but then also, can we allow that person to do what their putting forth that they want to try to do. And the granularity within GFIPM [Global Federated Identity and Privilege Management], as well as within the XACML stuff, really allows you to do those kind of descriptions and have it be transparent to the user. More reliant upon the user database, what's being maintained by the states as a trusted user authenticator, and what's been done by the state as a trusted authority trying to talk to Alabama or trying to go through our centralized portal.

Di: So can you think of a couple of use cases that you've implemented in the CONNECT queries where there are particular data resources that are permitted to be viewed, but not others depending on maybe

the conditions, or the user's role, or what kinds of attributes are important to you?

Mike: Again, for us, the easy example is driver's license photos. Nebraska statutes allow driver license data to be shared pretty much with criminal justice users.

[00:03:00]

Di: Name, address?

Mike: Name, details, demographics, some things about the driver history. And those can be shared with probation, they can be shared with corrections, they can be shared with law enforcement, prosecutors. But, when it comes to the ***photos,*** there was a concern about privacy with photos. So when we put in the digital photo, we had a couple of centers that were very concerned about both use and misuse.

So the photos were restricted to law enforcement agencies, as well as DMV [Department of Motor Vehicles] agencies. The statute has changed a little bit to open it up to certified law enforcement officers and other agencies. But we need a mechanism then to talk about not just a user coming from Alabama that was . . . . We need to know, could it be a probation agency? No. Could it be a law enforcement agency? Yes.

Now, we haven't done it yet, but if [Maury] can take his county attorney's office or district attorney, or whatever they've got, and they have sworn officers within that agency, and he can credential them and pass them off to Nebraska, we could expose those photos to them. But we need that level of detail. So, we need that both within our state, as well as when we're trying to talk to other states.

[00:04:05]

Di: So matching these various attributes about the target resource – photos – the credentials of the requester – are you or are you not certified, sworn – and then making those decisions. Tell me about how you implemented that policy decision point in CONNECT. What does that look like?

Mike: Basically, a number of . . . . almost like a checklist that has to be handled, going through the credentials that are being passed off.

First off is at the state, and then, for the user, do they match certain "yes" or "no" kinds of things. In my mind, I think of it as a matrix, almost. You can have 1, 2, or 3 things on this level or if you have 2 on line 2, that's okay. Or if you have line 3 and you're a superuser, then that's it as well.

[00:04:59]

So, there can be different conditions that need to match. It's not just necessarily one credential or one type of user. But you have to have that flexibility, and that's what really nice about this system. It has that flexibility to be able to build in multiple conditions, or multiple options, be able to describe a user, as well as – not just the user – also the type of agency. So you can get a lot of different types of granularity.

Di: So my understanding of CONNECT's architecture is that each state is hosting its own PDP, very near the target resource. Did I have that right, Mike?

Maury: Well, yes, there are two sets of PDPs in the system. The first one is actually at the portal level. The initial one: after the identity provider of the state sends the credentials to the CONNECT portal –

[00:06:03]

– the policy decision point there looks to see, first of all, if you even have authority to get there. And, that's that decision. We keep that wrapped up in XACML policy file that is central, everybody shares with a common, it has a common look and feel for everybody in terms of what goes in there.

Then a query is also submitted, and this is transparent completely to the user. That's what's so nice about this, it gets complicated on our end, but they don't realize how much technology is behind the scene to make this work. The query also goes up to the server, but then it carries with it after you're approved at that first decision point back to the data source that you're trying to query.

And at that point, yes, the decision point, the PDP, is back within the state that the query service actually looks to. And that determines, with your credentials that are coming through, if you have authority to look at that very specific piece of data.

[00:07:04]

We also keep the same, well, a similar type of XACML file right there is our administration piece, to let the decision point know whether or not to allow that.

Di: So that makes a lot of sense to me. From the nature of your relationship, these four equal states, that you would have sort of this federated policy that is "you can come on in or not come on in." And then also you would have the dispersed decision points that each of you manages with regard to your own data resources, right? John, have you seen some other ways of designing the PDP that meet different kinds of business needs?

John: Well, I think there's the centralized PDP model for a whole enterprise –

[00:08:00]

– which has the advantage that you've got one repository for all the different resources. And then there is the design where it will make certain high level decisions and then there are multiple PDPs. From a performance standpoint, having multiple PDPs down close to each of the resources, if you're using a commercial product, making sure you've got everybody using the appropriate rules and you're keeping all of that in sync is possible. And that would provide a more robust performance environment.

So I think the distributed model is a good one, but you definitely need a commercial tool to make sure that these are all talking to each other. And when a change is made at the master or down below, that they're not getting their policy separated out.

Maury: And again, a competency level back at the data provider –

[00:09:00]

Since it's distributed, everybody has responsibility involved with that. So there's an administrative piece that's got to be upheld. But it also gives control back to that data provider, as well, that's very tightly connected to them. So there are some advantages.

Di: Are there any other considerations? For instance, I've heard a consideration of what you just described, Maury, where you've got control, but you've also got responsibility. When you're going to

stand up a distributed PDP model, you've also got the system performance issues that John highlighted. Are there any other considerations that people need to be thinking about when they're designing their architecture?

Mike: One thing is related to auditing, because as people put systems in, you have to know what's being done with your data or with your system. And typically, I know at the state level, we build auditing tools to track everything that goes on. We have to do that. Maury has done the same type of thing.

[00:10:02]

When all of a sudden you're passing data and queries out to other people and you're centralizing part of it, and part of what is going on is away from you – if I'm a Nebraska user, I can track kind of what I'm doing, and I'm passing off to the portal. And we can monitor what's going on in that centralized spot. But if I'm monitoring a query against Alabama, a lot of those details and the results of either passing on or rejecting that query and actually doing a search all really needs to be done at the Alabama level. Working all that out, and what you're going to do if you actually need to get into audit logs? Who's going to monitor that? Who's going to maintain that? That's another level of complexity.

Maury: Since we have two decision points here, there's the idea that there's two sets of audit logs, as well. First of all, the portal itself is going to log everything going through it in terms of who came, at what time, and who they . . . .

[00:11:05]

Di: Everyone who knocked on the door is logged?

Maury: Everybody who knocked on the door. It carries with you that set of credentials. It identifies that person at that moment. And then, back on our PDP, I guess, back at the state level next to our dataset, we're going to log anybody that came to our specific door for that. So it won't be the big, overall, we won't know everything going on in CONNECT, but it'll be a log related to that specific query service that we will maintain and have access to and it'll share the credentials from whatever state or, you know, whatever entry point they came into. We'll have that ability to look at that.

Certainly, back at a higher level, at a policy level for the states, our agreement allows any of us at any given time to share the logs in terms of looking at what the other states did –

[00:12:03]

– if we have any questions. And it's really not a problem being able to deal with each other, to know what's going on.

Di: It's just a way that you designed your accountability to each other?

Maury: Sure.

Di: Is it okay to drill down just for a little bit about that audit logging? Are you using any kind of business analytics tools to try to detect how well the system is enforcing the rules or for any other purpose?

Mike: No, we haven't. But I think we see that we would like to automate and enhance kind of that middle administration piece. Within that centralized portal, there's a need probably to enhance what we do with auditing and with use, so that it's not a manual process. So we can more easily get at that kind of data.

[00:13:02]

And it might be the types of things you're describing. But I think we're still, to an extent, young enough that we can dig into a manual log. But as we expand – both in datasets, number of PDPs, the number of users, the number of partners – enhancing that kind of maintenance component is something we really need to do and hopefully automate at a better level.

Maury: One thing that we did do to divine what goes into the audit log is we went by the federal CJIS [Criminal Justice Information Services] security policy. So hopefully in our states, we're all ready used to looking at those kinds of logs anyway, as we log all the national traffic through NCIC [the National Crime Information Center] and NLETS [the International Justice and Public Safety Network], but very similar to the way we're going to record this now, for any access to CONNECT.

Di: This makes me want to ask you gentlemen, how many users do you have? And what kind of traffic do you have coming through the CONNECT portal?

[00:14:02]

Are you talking about hundreds of users and hundreds of queries every year? What's the quantity?

Maury: Well, it's growing. I think in the border states – Kansas, Nebraska, and Wyoming – I'll speak for you [, Mike,] for a sec – they have a lot more activity just because of the very nature of their location. It's not as many times when someone from Alabama needs to go out looking at someone from Wyoming or Kansas, even Nebraska, yet. And so it's not a tool that is readily necessary for us in Alabama. The traffic's relatively low. I mean, considerably low compared to all the other kinds of queries we do.

I think with the new addition of N-DEx [National Data Exchange] searching directly from our interface, that's going to make a big difference. I think that we're going to shoot up dramatically because we're going to provide that tool to everybody now.

[00:15:02]

They're going to have LEXS [Logical Entity eXchange Specification] and N-DEx in a way that they've not typically had, or not everybody had a LEO [Law Enforcement Online] account. But now it's going to work through CONNECT, so that's going to be very nice. I think for your very real world scenarios of borders, there's a greater demonstrated need right now.

Mike: I wish I had numbers. I don't have numbers of users or queries right at hand. But we're seeing it particularly along the borders, but also people with interstate traffic. I-80 [Interstate 80] is a drug corridor and we know there's a lot of traffic going back and forth from the coast, out of Chicago. So people want forms of identification and as much information as they can find on people. So it's spreading. Certainly within the border counties, as well as our state patrol and the CID [Criminal Identification Division].

Di: When we were talking about identity providers, we were talking a little bit about –

[00:16:01]

– helping your internal technical staff achieve expertise in these new standards, these new technologies. Also kind of helping them

culturally see this new way of doing business. John, I think the language that you used was that it's different from any other IT project because it's literally a new development methodology. It's changing the way we do business as technologists.

When you got to the privacy enforcement point and the privacy decision point, did your staffs have the same needs for learning new skills and new ways of doing business?

Mike: Yes.

Di: Yes! And how did you help them through it?

Mike: That was the tough thing because

[00:16:58]

– you know, we said, we used a vendor. But, independently of that, nobody, nobody – not our day-to-day staff, nor ***anybody*** – really had a good grasp and hadn't done a lot of implementation on this or even attempted development. So it's not like we could sort of help the vendor that we were using to do the development on. They needed to learn it from scratch. And it was a tough thing for them. They admitted to a huge learning curve. It was really, really difficult but . . . .

Maury: I think they took a lot of responsibility on and taught us, in fact. That was kind of a hand-holding exercise as we worked through this process.

Mike: Yeah. One of the strange things: we can talk about PEPs and PDPs, but for the most part, that's not how you, at the time we started, that's not how you were thinking about all these things. We were thinking we need to be able to talk state to state. And we need to be able to have something that works in the middle. And then they needed to figure out how to make it happen with or without a structure of the check name and whatever might happen.

John: There really wasn't even a vocabulary in Microsoft terms until fairly recently: they've come up with what they call "Claims Aware Applications."

[00:18:04]

Three, four years ago, you wouldn't know what that meant. But today what that means is there are claims coming to the application – attributes, etcetera, about the particular request – and the application is now able to look at those claims and do the policy checks and the role checks etcetera, before granting access to a particular resource.

But that wasn't even in the vocabulary of Microsoft; they came up with the "Claims Aware" term. And another term they came up with, which is a little different than the SAML [Security Assertion Markup Language] standard is, "Relying Parties." And people [ask], "What do you mean, 'Relying Parties?' What are we talking about?" That really is, a Relying Party is a claims-aware application that's relying on credentials that it's receiving, not credentials that it's provisioning.

[00:19:00]

And it's [not] doing all the claims collection and all that sort of – it's just validating and ensuring those claims are from a trusted source, and then it's making the decisions about authorization.

So it's very new, very new for everybody. I really think this area, the authorization piece, is probably still another five to ten years before it really reaches critical, mainstream mass. Whereas the federation pieces, we're probably a good five years into ***being*** mainstream, and within another five years, [people will wonder,] "Was there any other way to do business?" So I think we're on some different evolutionary paths between those two technologies.

Di: Do you happen to know from Analysts International, what kinds of resources they found particularly helpful –

[00:20:00]

– in educating themselves and overcoming this knowledge barrier on the XACML architecture, the XACML standards?

Maury: Well, there were discussions with people on the Global technical working committees. I think that's where it was all being maturated, right there. And since this was do it as you go, we're just sort of sometimes shooting in the dark, just thinking, "Let's try this approach."

Mike: There wasn't documentation. You weren't going to go buy a "XACML for Dummies." You know, it wasn't there. So it was a learning curve for everybody. But now there are resources, thankfully, within Global and with some vendors and the states that can hopefully help out other people.

Di: Did you find any useful resources in other sectors? Maybe banking, or anything in the private sector, that are helpful case studies? Anything?

[00:21:00]

Maury: We had discussions about other paradigms, so to speak, throughout our years of meetings.

Mike: It seems like a lot of it, though, went back to the use of XML [eXtensible Markup Language] or development of IEPDs [Information Exchange Package Documents], and things like that, more than on the security side. Because it was a little bit different at least . . . .

Maury: The typical question is, "Well, if a bank can do it, why can't we do it?" We did discuss those kinds of topics.

John: Health services has a big push for the electronic medical record. They have some actual due dates and times. So they, for a long time, have been looking and adopting federated ID, and actually have a workgroup in healthcare XML security and privacy authorization based on the XACML standards.

[00:22:03]

But again, what healthcare is finding is that most of the culture is still not there, so they're doing sort of a tactical exchange of information just using secure email while they come up to speed on how to deploy some of these other IDPs [identity providers] and service providers and coding their policy. So you have some groups that are out in front, but a large group that haven't even stepped or put their toe into this space. But it is another space that, when you go out and look at OASIS [Organization for the Advancement of Structured Information Standards] and some of these sites, you will see publications in healthcare with actual profiles for the SAML, Security Assertion Markup Language, protocol. Some sample XACML policies for getting patient consent, as to whether they consent to have their –

[00:23:00]

– drug history shared with a drug treatment provider or not, and that sort of thing. So there is work going on outside of criminal justice. The other areas I've heard about are in aerospace and in financial. And again, some of these are driven by compliance rules in those industries.

Import-export rules have some very strict controls over what information you can share across international boundaries. And so they've leveraged this architecture: depending on if you're in France, maybe you cannot get the access, but when you're back in the States, you ***can*** get the access to maybe the same piece of information, because of some international policy rules. So those are the other industries that are active in the externalization of authorization policy rules.

[00:24:00]

Di: Understanding that CONNECT has been working with a private service provider for much of the development, do you foresee that the maintenance of the CONNECT pieces, is that going to cause you to think about your technology organizations in a different way? Just as one example: John, I think that you have proposed that it might become a best practice to actually have an identity provider support team, similar to a database administration team today. They've got that expertise, in-house. They do that one thing. They do it for ***all*** of the different applications in the enterprise.

Does the question make sense, about how you might be thinking about your tech organization?

[00:24:59]

John: Let me just clarify a little bit on that. Within each organization in L.A. [Los Angeles] County, the directory service in our sheriff's department, that's all managed by one group in the tech support area. They are the most able to get the idea and get the concept. In terms of the maintenance of the attributes on the people, that'll be distributed out in terms of them updating that central directory. But they're in a position to take that active directory and layer on top of it an identity provider that can take the internal claims and

convert them for inter-agency exchange of claims using the SAML protocol.

Di: John, let me make sure I understand. You're saying that there's a central team who's managing the technical markup and the structure of the IDP itself?

[00:26:04]

John: Of the organization's directory of attributes about their employees. And they also know which application might have supplemental attributes: they are in the position to be able to work out a query to bring in attributes that may be part of a training database, or part of the payroll system, that may be relevant, that they need to go draw [in,] that's not in the core directory.

Di: But they are not maintaining Deputy Roberta Smith's job title and the individual content of each individual employee? That's a very local . . . .

John: That's done by the local management for that particular set of attributes. The payroll systems are done by the personnel department, and when they record that information –

[00:27:00]

– that goes into the directory and then that can be used for different functions. You know the payroll title and you can determine what their role is in the organization, etcetera. Or the training manager's got a training system, and they may set it up so that you can query that training management system to find out if Bill is up on his 28 CFR [Code of Federal Regulations] Part 23 training or not.

So it's a service provider of the actual directory service itself. I'm kind of looking at – those are the folks that can build the other pieces for an identity provider. And on the service provider side, you're going to have to work with the line of business applications and the managers over those applications to determine what the policies are. And then again, the application programmers are going to need some support if they're going to maintain a policy on that –

[00:28:01]

– in terms of working with the department manager and using whatever policy authoring tool you have to record what the rules are around who can access that resource. So you're either going to extend many application teams, but you probably need a core competency group that can provide the consultation and the direction. Just like a database administrator rules, kind of has a core competency to consult with the application teams when you're moving a database into production.

So I kind of see them in a consulting role towards the service provider policy authoring and actually providing the infrastructure for both the policy rules, policy administration points, etcetera. You look around the organization, someone has to do it. That's the one that just comes to my mind. So that's just my personal view on how to introduce this into the IT organization.

[00:29:03]

So I'd be interested what thoughts –

Di: What your long-term vision is for the long-term maintenance, sustaining CONNECT over time? Maybe it will remain with your private service provider? I shouldn't say, "service provider." Your private ***solution*** provider?

Maury: Each of the states sort of brings to the table some different capabilities. We did have a solutions provider come in and help develop a lot of the underpinnings to make this happen. But we still have our own state IT services. In our case, I'm in Alabama, we actually . . . . Because we had a different set of development capabilities in Alabama, and frankly, our architecture was different than Kansas, Nebraska, and Wyoming. We did a lot of the development ourselves internally. So I can't speak for everybody –

[00:30:00]

– about how the long-term vision goes for how they will maintain and sustain. We all have agreed that we will sustain. Part of it, since we all ready have portal capabilities in place, we had them before this existed and they will always exist now. There's a model in place for us to do the management. This doesn't add too much more administration on top. We're building new competencies to allow some of these technologies. But what I'm personally benefiting [from] in Alabama, is now that we're using

them for CONNECT, I'm sort of retrofitting my applications in the state to apply these technologies to our internal applications. We're working on a very significant GFIPM re-engineering, but also the whole idea of NIEM [National Information Exchange Model] coming in.

[00:31:00]

We'd been using the GRA [Global Reference Architecture] for significant amount of our efforts. But across the board now, with these standards that Global has put out, our rule of thumb is you will apply them across the board at ACJIC [Alabama Criminal Justice Information Center] in everything we do, for all new developments. And that's where we're trying to approach this. So I don't think the sustainability will be quite as difficult.

But every time you re-engineer, there's costs, and it is a hurdle. I'm telling you, it's a significant hurdle to take a lot of these systems that we've had in place for years now, even though they were somewhat advanced for their time, and even today they're pretty advanced, still, there's a hurdle and a cost to us to re-engineer them enough to start using these Global standards. We do see the benefit and that's where we're going to go with it.

John: Do you think that you'll have a team that goes and helps consult with the various application support teams?

[00:32:02]

Or do you think all of the application support teams will go to training and figure out how to do that? What kind of tactic are you thinking about?

Maury: The different development teams now, we are requiring them to go to training. We're trying to get technical assistance when possible from the different DOJ [United States Department of Justice] technical providers, which is very good, I mean, there's some great TA [technical assistance] out there. But any opportunity we can find now, we're sending somebody to go learn. You've got to build these competencies up where they're just second nature and they haven't been. So it's just a difficult process when there's not as much money these days.

Di: That's interesting, Maury. What I hear you saying – maybe in contrast to Los Angeles County – is that you see some of these as core competencies for any technologist in your shop?

[00:33:01]

Maury: Mm-hmm.

Di: Mike, how are you looking at it in Nebraska, in terms of your internal Nebraska IT staff?

Mike: I'll kind of bring it into the way we were discussing things earlier, with the centralized approach. At the state level in Nebraska, we'll probably continue with our vendor. We'll probably continue going onto maintain that state level workload and that state level integration and the enhancements, the PDPs that we need to do there at that time, which ties into our user base and everything that way. And I think all the states pretty much agree that we need to take that on for ourselves, getting competency, getting education, getting training, and accepting that as just a cost of sharing the data.

The other thing we have to be concerned about, though, is that when we centralize some of these services, we have to have a way to share that support. We have to know that that needs – we can't let that go away, or everything falls apart.

[00:33:58]

So we have to be sure that we find a mechanism to support that, to enhance it, to do maintenance, to do development, and everything that way. We've been okay so far. We've been fortunate to have grant money. We've also had contributions from the states going in at different levels. But it's going to be a challenge, as Maury said. Money is rolling up in a number of things.

So how we share that, states have been willing to put in money, when they can, how they can, but it becomes tough to ask for a new line item in a budget to sustain that. So hopefully we can minimize some of those things. We talked earlier about auditing, about maintenance, and things like that. We've talked about the benefits of automating some of the development of the PDP and PEP so that we can do that without hard coding everything.

Pay for it once up front and get that integrated into our process so we don't necessarily have to rely on full technical resources all the time to be able to make enhancements, to be able to add data sets, to be able to enforce those policies. That should really help us. And again, there are some commercial tools that are coming out now. We might be able to leverage those.

[00:34:59]

We might have to do our own internal development. But I think we have to build it so that we know that we need to maintain it at a reasonable cost.

Maury: Again for future development, this push is coming down from Justice as well, but the awareness that any new systems in place need to be Global aware. They need to take into account up front. And it's going into RFPs [Request for Proposals] now. It's injecting itself everywhere. Whereas before, it was an argument, "Well, you're just not going to get a broad enough audience for responding to solicitations." Now, it's out there.

I think we've hit this next level of maturity across the whole spectrum that we can feel comfortable with. It certainly wasn't there five years ago, and I don't think it was quite there two years ago even. We're edging into it. But now it's getting there. It's part of the topic of conversation.

[00:36:02]
{deleted}
[00:42:19]

John: I just think it's a shift to an architecture for information sharing. We haven't had an architecture for information sharing before. We've had a lot of ad hoc ways to kind of move information. But now there's actual architecture for doing it that's repeatable, that's open standards-based, that really allows us to share information between these various systems. Bottom line is it's very, very important that we ***do*** share this information. And a lot of things and a lot of problems are going to be identified and solved and problem-solved because of information sharing.

[00:43:02]

I can't tell you how our Board of Supervisors, at numerous Board hearings, are sitting up there saying, "Why is this foster child's

school records – why is he taking the same courses he took before?" Because the education department isn't sharing the foster records with the schools that the county's putting them in. And then, "Why did this person who you knew had particular drug requirements and needs, why weren't those records available for the emergency doctor who's trying to treat them?" There are just a lot of situations where you need information from multiple systems for the particular incident and individual that is being addressed at that particular point in time.

This kind of overused (maybe a little bit) motto is, we're really now in a "need to share" versus this "need to know," meaning you shouldn't have this information –

[00:44:00]

– unless you can tell me ***exactly*** what you need to know and then I'm only going to show you this little bit. But really recognizing the business case for sharing information and for those business purposes how important it is to share the information as opposed to holding it. So I would just kind of conclude with, I think that's the big shift.

[00:44:22]