Di **[00:00:25]**: So we were talking about a governance structure in which the executive policy makers and decision makers are able to delegate the technical implementation and the exploration of how to implement this information sharing to a group of experts. So what we'd like to do now is really drill down on some of the architectural decisions and some of the architectural evolutions that CONNECT has been through.

[00:01:02]

Just a reminder that you started this endeavor in 2005, really before a lot of the standards or best practices. You are the pioneers. You're the ones who came into this new land first, so we want to commend you on that. How did you transition to an external identity authentication? What did you have in place in each of your four states, or what did you have to put in place in each of your four partner states to really enable that external authentication?

Mike: It was an evolutionary thing. We went through a couple of different stages. Originally, we did everything basically internally to the state.

[00:02:00]

We had our own user base, our user knowledge, our centralized database of capabilities, aspects, credentialing basically, within our data sets. I think we just kind of extended from there to our first model where we pushed all of the identity concerns at the state level, kind of coding that externally. That kind of got us through Version 1.

As the standards evolved and also as our primary vendor got to know the standards better, because there wasn't really a path to follow, to a great extent, they made a recommendation that we look at possibly a more web-based technology. We were all going to web-based technologies anyway. We were evolving to that kind of thing. That really allowed us to go into a centralized identity provider for things. The vendor, and actually most of the states had some contact with Analysts International, made the recommendation that we kind of go to this web-based thing.

[00:03:03]

It'd be easier to push out to other people. It'd be easier to manage. It would take less work at the local level and at the state level. It seemed to make sense for us at the time. It raised some questions and concerns. How do we actually do that? How do we bring up this website, this internal web-based portal to basically do the identity management? How do we find a home for it? How do we sustain it? How do we look at it over time? Kansas was willing to take that step for us, and it really helped us. It was an evolution of the technology, as well as understanding what was needed and what could be done and how it could be done. It was also a cheaper way to approach it, probably.

Maury: I think a key point here is that each of the states had an understanding of identity management already in place. We had a procedure that I think we've all even evolved ourselves internally since then, because we've learned a lot and even the technology that has grown out of the GFIPM [Global Federated Identity and Privilege Management] experience is allowing us to more mature our internal operations, in fact.

[00:04:06]

But it is a mindset. When you have this sort of top down management of multiple users and multiple agencies seeing different sets of data, it's not, "Well, you get access to this one site with one set of data, and it's all or nothing."

You have to just totally change the way you think. Once you do that, and you've got a technology in place, which was sort of what we already had, then we were able to get to the table as the board and think, "How do we externalize this in a way that makes sense?" and it was a learning as you go. Out of the chute the first day, there was nothing mature sitting there at all. I think we failed a few times even. We tried different approaches.

[00:05:00]

We looked at different technologies. The vendor would come back and say, "Well, here's how we're going to do this security model, we think." Then it came back, "Oh, that didn't quite work like we expected, so we'll go back and try a slightly different technique." To the point we are today, it's a really well implemented way of taking the federation model and putting it against our products in our states.

John: I think one of the technical things that CONNECT found as they evolved was that it was just more efficient to have a portal and have the user authenticate to their identity provider and have the portal then do all the communication to the other resources. Originally what they were doing is after the user authenticated, they would go to a remote resource, and then they'd go to a different remote resource.

[00:06:03]

They weren't going to a portal that had it all together. It wasn't centralized in one portal. What happens is just the messages going back and forth, the response time was not good enough. So this is one of the lessons learned, I think, in setting up this federation, that you do need to look at your message traffic and what's the most optimum way to configure this, and that is something that they discovered. I think that's an important lesson to take away from that particular implementation.

Mike: And that extended as well, not just for the security and the messaging and the timing and communicating with all the states, but also getting the data back, not just the authentication.

Because it really does cascade and keep going on. So for the user experience, you need to minimize security, but it trickles down to getting the data back and getting it all synchronized as well.

John: You can link performance there, too.

Mike: Exactly.

[00:07:00]

Maury: Well, the architecture or the shaping of the way that data is transmitted, NIEM [National Information Exchange Model] has grown since we started.

We've evolved with it. The technology such as LEXS [Logical Entity eXchange Specification] and LEXS Search and Retrieval and applying that, tying that in to the way XACML [eXtensible Access Control Markup Language] now is included in the decision process. It's just –

John: Building blocks.

Maury: Building blocks, and it grows, and we learn and apply. That's sort of where we are now.

Di: I want to make sure we understand this clearly. Mike, can you talk to the architecture of your authentication today? If I'm a Nebraska user, and I go onto the portal, walk me through how you know that I'm really the okay user that I claim to be.

Mike: For Nebraska users, we're basically going through NCJIS [Nebraska Criminal Justice Information System], which is our internal data portal. So one of my users will log in under their normal log in procedures, and if they want to search Alabama, Wyoming, whatever, there is a button that they click on for CONNECT.

[00:08:03]

At that point, we're basically going out and talking to the centralized [CONNECT] portal where we're passing off the authentication, the types of – once we decide what search to make and make that transition – we basically pass off those credentialing back and forth to the centralized site. It goes in, we verify the type of user. Things are all happening in the background of the user. It doesn't matter. But our server is talking to the centralized identity server, passing off the credentials. The server on the other end is going to verify it, pretty much that it's working, and then deals with what type of search it's trying to make.

Is it trying to go to Alabama? Is it trying to go to Wyoming? Is it trying to do whatever it might happen to be? At that point, it measures the credentialing. It goes in sync back and forth with the partner that we're trying to get the data from. Again, all transparent to the user.

John: So you leverage an existing portal that you have in Nebraska for an existing set of users that you vet? And so you're leveraging that as your identity provider. And from that, you are building the credential to go to the centralized portal.

[00:09:14]

Mike: Correct.

And that's the underlying thing. We needed that user database and all of those credentials and those details about the users to operate our own data search capabilities to operate NCJIS.

So we had that built in. We didn't build it externally. It could be built externally by someone who didn't have one and needed that to connect to CONNECT or some other federation. But we had that, and all the states had that as an essential component. We might have had to expand it to say that we needed to know about sworn officers or more details about the users to pass off in certain types of credentials. Yeah, we had that user database in the backend, and that really provides kind of the underlying . . . .

John: So it's fair to say that the existing portal that you had with the centralized data store, that became your identity provider.

[00:10:04]

You didn't start with an active directory that gets consulted to access your central portal. That's going to be very common. As we're looking and evolving here, we're going to see most applications are not – in the government space, at least the space that I've worked with – their applications are not active directory or LDAP [Lightweight Directory Access Protocol] enabled. They basically have their own user store. For example, my DA's [District Attorney's] office has every attorney, every employee, etcetera, all in that one application. I could see that application basically being a data store or an attribute store. When they active-directory enable that application, then all of those attributes will be kept in the active directory. But it's not necessarily true that you have to have a separate directory in order to proceed.

[00:11:01]

You can use an existing directory that may be part of an internal application and generate from it the credentials to participate in a federation.

Mike: And not having just one way to do it. There could be partial implementations of LDAP or active directories that fit certain needs, but not everything.

Di: Did you, as the four state partners in CONNECT, did you have to agree on what those minimum requirements were for user credentials? How did you do that? How did you arrive at that

agreement? Most of your exchanges are in the law enforcement area, right?

Mike: It kind of depends upon data sets. We initially started with driver data: driver license data and driver photos. For the most part, those are open to criminal justice, in most states.

[00:12:01]

In Nebraska, we actually had statutory twists, where the drivers' photos were restricted to law enforcement agencies, DMV [Department of Motor Vehicles] agencies, and then eventually certified officers, but not corrections, not probation, things like that. So from the beginning, we actually had to deal with this notion of limiting users, even though we weren't really thinking totally in terms of credentialing. We had to have a way to deal with that. We kind of took a classical approach: let's just start with law enforcement. But as we get more involved with the nuances of GFIPM and realize that we can use all of these user descriptions for credentialing, that really allows us to do more.

Then when we add on things like corrections and courts, we're able to look at it in a little broader picture, but that really kind of provided some of the initial thought process and some of the original headaches, actually, to tell you the truth.

Maury: Many of our discussions still right now are what are these credentials? What are these attributes of the user that define them?

[00:12:56]

Maury: As we come forward and mature toward the new model, GFIPM 2.0, that allows so much more flexibility, how do we incorporate that back as we credential them locally? Are we making sure we gather those distinguishing points in our system, in our active directory or our identity provider model that allows us to inject those into the way GFIPM sends the information.

John: I think Maury is bringing up a good point: that, in order to understand who the user is, you need a common set of descriptors to describe them, whether they're a sworn law enforcement officer, whether they've had certain training, etcetera. So when he refers to the GFIPM metadata 2.0, he's just talking about a vocabulary of attributes that you attach to the credential when you send it off to access a particular resource.

[00:13:58]

And what you have to do internally is look for that attribute, [for example,] sworn law enforcement officer: what is that in my local personnel system? What is that called internally, because I'm going to take whatever that is – in our county, we have item numbers, and we've got one for officers versus deputies, we can then distinguish someone's role as to whether they're management or line deputy – we would map that to sworn law enforcement officers. Whereas, we have lots of other items in the sheriff's department which are civilian, so those items would ***not*** map to a sworn law enforcement officer.

So there's this exercise that you do to come up with a common credential and a common set of attributes in order to enable this federation to recognize as a resource provider what the attributes are about this particular requestor.

[00:14:58]

When we refer to this GFIPM metadata, that probably doesn't mean much until you put it into some context, but that's what it's about.

Di: One of the biases that we've had in our conversation so far is that these will be human consumers, human requestors. Does the same model work for system-to-system exchanges?

John: Yes, yes. As a matter of fact, CONNECT is leveraging that, in that their portal is presenting the data, but behind the portal the credential is making system to system requests for the data that may be in different distributed areas around the state, if you will, so they are using both. CONNECT was one of the first projects – or it is ***the*** first project – within the Global family of product developments to actually use a service system-to-system interface.

[00:16:00]

Di: So a system is going to be credentialed or authenticated in a very similar way, using this same standard?

John: Actually the message that comes from the portal has a signature from the system that uniquely identifies, "This is the CONNECT portal," and the resource on the other side validates that before it

will accept the message. Then also within that message, it says, "And here's the user who is on that portal right now, and here's their attributes," and then the resource, number one, knows this is an authorized federation member, the CONNECT portal, and then they know the attributes, and they can determine from that what kind of access they will grant or deny. So absolutely, system to system, user with a browser going to a system – both of those use cases are supported.

[00:17:00]

Di: So imagine a situation in which you wanted to engage a new information sharing partner, but this new partner didn't have an identity provider capability, either because they were small enough that it didn't really make sense for them to have that capability or just in terms of their technological maturation, they just hadn't reached that point yet. What are some of the strategies that you would recommend for enabling an identity provider in that kind of a situation? Are there some options that people should consider?

Maury: I think to come to the CONNECT table, we do have an assumption you are going to have some ability to manage your users, some ability to provide authority, some ability to credential them.

[00:18:01]

And it is true that there are many peers of ours across the nation that have not created a capability that is as complex as we may need to implement the attribute approach through GFIPM. I guess what we would say is we have lessons learned. We've got some battle scars in terms of the way we've had to implement each of our states, and some of us are different than others. Mike and I don't do it the same way within our states, but we do have this ability to know who the person is, what their authority, is and how to present that to the system.

We can give that another state. We can say, "You don't have to start from scratch. This is sort of lessons learned how you can go about doing it." Some of the COTS [commercial, off-the-shelf] products, some of the vendors out there now are baking in these possibilities into their new systems.

[00:19:02]

They not only have to take into account the way you manage people, but security models, like the CJIS [the Federal Bureau of Investigation's Criminal Justice Information Services] security model, and then the technologies coming out of Global – the GRA [Global Reference Architecture] and the GFIPM and so forth – so it is going to become much less of a burden on a new partner than it would have been in the past.

Di: So what I hear you saying, Maury, is that the open source and the commercial options out there are much more sophisticated now than they were in 2005.

Maury: Oh, absolutely.

Mike: Well, they exist.

Maury: They exist.

Di: There were none, now there are some.

Maury: And the nomenclature is becoming . . . . It's not like teaching a foreign language when you go talk to someone. You understand you're at the table with the folks that deal with this criminal justice community because that's where our focus is right now: criminal justice.

[00:20:00]

Between the BJA [Bureau of Justice Assistance], the folks at DOJ [Department of Justice] and BJA, have made a spectacular push across our nation to educate. The IJIS Institute, who really steps forward to bring the industry partner community together. It has made such a big difference: because of the nomenclature, that we're all talking the same language. It helps when we want to implement this kind of project. It's not just for us. Whether it's at a micro level or the state level or the county level, just to know that the products that are moving forward are going to recognize these technologies makes it easier.

Mike: I'll just toss out – Maury mentioned the commercial options and other things – there are other ways to build up just an identity provider for somebody if they wanted to do that, or there are ways to centralize those. They're all technically possible. They can all work. They also, though, bring on a level of administrative complexity and administrative need.

[00:21:02]

And so it's more than just the technology, just like everything else.

Maury: It comes back to commitment.

Mike: Exactly.

Maury: You've got to have a willing partner that's going to take on the responsibility and have authority to do so.

John: I would just add that in Los Angeles County, we have over 44 local police agencies, but there's really only about three or four of those local police agencies that have the technical infrastructure and the data centers, etcetera, to securely set up an identity provider. So what about the other 40? How do they play?

What I think needs to happen is either a cloud provider or the county itself or the sheriff needs to set up the identity provider software as a service, and then the small organization can administer the identities and the attributes of their workforce.

[00:22:08]

Some of these workforces are 15 or 20 people. Some of them are smaller, and then you have all kinds of sizes.

I think we have to recognize that 80 to 90 percent of the population out there is served by these smaller police agencies, so they need this service. That's one, I think, tactical and strategic thing that we need to look at in terms of providing – because, as Maury said, you are responsible for vetting and managing your people and reflecting responsible attributes. That doesn't mean you have to be responsible for the technical infrastructure and the support and maintenance of it if you can get that on some kind of a service contract or whatever.

[00:23:00]

Mike: Even larger agencies might have a technical infrastructure or a technical staff, but that doesn't mean they can take on basically a new capability for almost a one off kind of chore, so that kind of shared service would certainly make sense.

John: Specifically within criminal justice and law enforcement, there's an organization, RISS, Regional Information Intelligence Sharing Systems. They offer software as a service, IDP [identity provider] services, for small agencies that can't host their own. The FBI CJIS with their LEO [Law Enforcement Online] portal offers an IDP-type service for small law enforcement agencies that are part of the CJIS criminal justice community. There's two already where they're trusted organizations. You know they have data centers.

[00:24:00]

You know they practice all of the security controls you'd want around protecting that identity information – that are offering those as services to the small agencies that just don't have those budgets or don't have the size to support that themselves internally.

Di: Very good to know. So in terms of the technical architecture of CONNECT, my understanding is that it's the CONNECT portal itself and the mechanism for moving these messages around. You built this in a Microsoft environment? Is that correct?

Mike: Right. It's all .NET.

Di: Could you talk a little bit more about your experience using the .NET platform and trying to develop these capabilities? There have been some easier parts and some harder parts maybe?

[00:25:00]

Mike: I'd say it was more of a default for most of us. We were using Microsoft already. We had a vendor that was familiar with components and with our development staff and went with that. There wasn't, I think, for us a debate about the proper or a preferable technology. It was mainly operating out of a Microsoft environment. The bulk of it was there, and it was what people knew and what we could hire people and sustain it and maintain it. It was kind of driven by prevalence more than some kind of technical checklist that had to be one thing.

Again, the idea of standards is it should be able to be implemented in a variety of environments. So, for us, Microsoft was that standard environment, and I think we all just gravitated to that.

Di: So that was also true for your other two partners, Wyoming and Kansas? They were also .NET shops? Would it be a requirement for any new states? No, and that was the point that you were just making about the standards, right, is that it should be able to speak with each other regardless?

[00:26:00]

Maury: It should be, that's true. We injected across almost all conversations the idea of thinking forward to other states. It has never been just about our four states. We're sort of the guinea pigs to make sure everything's working.

Di: Pioneers.

Maury: Pioneers, I like that word better. But always thinking out to other states that have a different architecture, a different technology base. We've talked to several that are not using Microsoft, and the idea is that we will be able to leverage these standards to make sure that it can just plug right in.

We're just very thankful that Global is working constantly to modernize and evolve these – we keep using that word a lot, but it truly is an evolution of these new ways of doing technology that is making our life easier on the state level and in the way we run our own systems and then as we make CONNECT better.

[00:27:03]

John: Basically as these open standards mature, we define – and "we," I'm talking about the Global products now – we define profiles of exactly what settings to use with those particular open standards so that it doesn't matter whether it's being generated from a Microsoft or a Java platform, an Oracle or an IBM. As long as everybody's following that protocol and that open standard set of settings, we can then exchange our identity credentials, and our service providers can consume and interpret those using any of those various vendors' commercial products. There's plenty of open source products, as well, for the really advanced organization that can support open source tools as part of their production environments.

[00:28:02]

Maury: I look forward to when we do bring on a partner that may not be Microsoft based, even though I feel very confident –

Di: We do, too.

Maury: I feel very confident though in the way Alabama and the other three states are leveraging Microsoft because we have commitments from them, too. They are going to make sure the technology that we have implemented is going to maturate right along the path.

John: With the open standards.

Maury: And it gives us a lot of confidence, secure assurance.

Di: So gentlemen, you've worked on other IT projects as well. Are there any fundamental differences between external authentication and sort of a generic IT project? What's unique about CONNECT that's been particularly rewarding or challenging for you?

[00:29:00]

Maury: Wow, that's wide open. Well, there are some personal things. It's been a joy working with the other states and learning more about them and seeing the way they do business and learning that they don't do things like us. I know Alabama has taken pointers already from Nebraska and Kansas and Wyoming and said, "Let's do that here because we haven't done that." I think they would say the same thing across the board.

Mike: Absolutely.

Di: Cross-pollination?

Maury: Exactly. The cross-pollination makes such a difference. Then being on a national scale, we're able to leverage bringing in experts like John and different points of view or the National Center [for State Courts] or IJIS and different groups to help us create this process all together. To me, that's the most rewarding part about it because we feel like there is an end goal, and there's a difference that we're being able to make now, because of what we're doing.

[00:30:01]

Mike: When you ask how is it different from other IT projects, that's just it – it ***is*** different from other IT projects.

John: Definitely.

Mike: None of us knew what we were getting into, quite honestly. We knew that we needed to authenticate or be able to – we knew we wanted to leverage a single sign on type of system where our users would be able to do that. We don't really have an idea at that time about what that really meant and what that constituted. It's been a huge learning curve at the technical level, at the process level. I like what Maury said. Without the Department of Justice, without Global, without the assistance of BJA, we really would have been floundering – a lot of technical assistance. John has been incredibly useful, a great resource, very useful at that level.

We appreciate everything that went on, but it's been a learning curve for the vendors, for the users, for what it means for policy users. We've had to think differently on all those different levels. We were thankful that we had to be flexible, I think, that if we had a deadline and people were going to pull the plug because we weren't done in four months, we would have been up the creek.

[00:31:01]

So that's good. People have stuck with it, and knowing where it was going and that, as we said, it's evolved over time. So, not a lot of projects necessarily have that luxury of knowing that things are changing, that you have to adapt. That kind of helped us in a lot of ways, I think.

Di: John, what have you seen out there? How is this different than sort of a generic IT project?

John: It's a cultural difference. Our technical programming support people haven't thought like this. They've always said, "I'm a team. I have a project, and ***I*** build – I, the technical team – build user stores for ***my*** application and authorization logic for ***my*** application." Then when you bring to them the idea of an external entity doing authentication, that's a foreign concept. The first thing they say, "Oh, I don't know about this. This is not what I'm used to doing, and why should I trust them?"

[00:32:05]

A lot of questions. It takes – I think the key here at the table – a lot of education and understanding about the business benefits of

moving authentication out of your every single application and into some kind of a central store.

Then you've got the people issues: programming staff don't want to feel ignorant, and most of us are ignorant in this space, and so you're going to have that kind of resistance to change.

Quite frankly, then you'll have your teams telling the manager, "Well, if I do it like I always do it, it'll be done in four months, but if I have to do it this new way you're talking about, it's going to cost you twice as much and take you a year, so which do you want to do, boss?"

[00:33:00]

So you really have to sell the concept and the benefits of moving in this direction all the way from the management right through the programming ranks. And you have to give them those consultant services and give them those guides so they don't feel like they're being challenged in a way that makes them feel ignorant. It's a fundamental human thing.

So this is a very different type of IT project because it's not an IT project: it's an IT development ***methodology.*** I'm changing the fundamental way you develop systems by going in this direction.

Maury: I think we were very fortunate. We had a vendor, Analysts [International], that came in and they were already somewhat plugged in at the Global level.

We were able to work with them very well, but back at our own shops, I'll speak for Alabama specifically, but I think we all had these experiences.

[00:34:01]

The folks back home said the same thing. "I'm not used to that." They had a hard time admitting they didn't know it, and I understand that. People want to do what they know how to do, what they feel most comfortable in. I believe this took a lot of growth. It took a lot of leadership too to say, "Guys, trudge ahead, no matter how hard it is. Learn about it. We'll bring in technical assistance."

Certainly, this is a new area with the GFIPM. I know very specifically when we were doing the NIEM [National Information Exchange Model] implementations early on, nobody wanted to do that. That was just, "No, I can do it so fast if I do this." I'm like, "No, we've got to go through the process. There is a way to make it right." And it works so well now.

[00:35:01]

I think bringing your staff up to speed, giving them the support, that they know that you're going to support them, that even if it's hard, it's going to make a big difference.

Di: Mike, how did you help your internal resources back at home make it over that initial hurdle?

Mike: For us, it was little easier because the vendor that we used primarily for NCJIS was the one that we were using for the CONNECT project. That really helped. It was different staffing, but it was easy to kind of break things in. But my primary developer still had some of the questions about not trusting it and not knowing what it was, even though it was basically his company that was doing the work. He didn't necessarily trust that as an extension. So for us, it was a little easier on the technical side.

The other thing that goes onto trust is that as this all went along, we already had this culture of trusting the NCJIS portal, of trusting sharing data, of trusting the users, that we weren't going to be giving the data out to the wrong folks, that the people that got in should be gotten in, that they had been trained, and we knew what was going on.

[00:36:02]

So I think that really helped us make that bridge to go over the CONNECT step. We hadn't really messed up, yet, so they said, "We've trusted you so far and it's worked, so we'll go with this, and we'll assume that you're going to be making the right decisions." Working with the other states, we knew we had to make the right decisions for everybody, so that helped.

Di: That's great. So what I hear you saying is that you had a similar experience as what Maury was describing in Alabama where maybe that initial step in the evolutionary process which happened inside Nebraska was scary and hard. But once you make it over

that hurdle, the second and the third and the fourth, each one gets a little easier, a little more comfortable?

Mike: The culture changes, but the expectations change as well.

Maury: They do.

Mike: It's the same thing. "You've done this fine. Now, I want to see them. Yeah, bring me the data. I do want to see that data."

Di: "Now, I want Google!"

Mike: "Okay, we can share, but again, bring me the data as well."

Maury: The expectations do start rising up there.

Mike: Absolutely. Absolutely.

[00:37:00]

Maury: But it's great to see that shift in the mindset across the board, and it's the same with Nebraska and with Alabama from the smallest police department up to the largest agency. They want to – their password is going to get them to this data, and they don't have to worry about it. They just know it's going to be there. It is great to see where we've come, and CONNECT is just the next step in that process.

[00:37:23]