Di **[00:00:10]**: So we're talking with Tony Lucich about Orange County's centralized identity provider and also a particular application of that with the juvenile court, which is the JUICE project. I thought maybe where we could take the conversation now is: help us understand what kinds of challenges you faced helping your own IT [information technology] staff in the county and in the court understand the benefits of external authentication and authorization and how you help them over that hump.

Tony: Interesting, as you would say, "the hump."

[00:00:51]

The challenge is that as we move identity into, in our case, an identity appliance, and the staff no longer is tasked with adding people to Active Directory or even being able to log into Active Directory, which they felt was their core business process, and the managers who traditionally said, "Well, I just call Fred, and he puts me in, and I get all the access I need," and they now have to go through a workflow, you're changing something that's pretty core to the culture.

So I think getting over the hump is the challenge of it. Certainly initially, they all were excited about any new technology, but then they realized it was going to be a change in their ability to finger control the way things happened. They felt that this was an obstacle to their ability to accomplish what their manager might ask for. By design, part of that is true. The manager wants access to a given system, he also has to go through the request and approval process for that entitlement of an access.

[00:02:01]

But at the same time, that's taken awhile.

So it's exploring the value proposition so that they understand that there are regulations around it. We should have been doing this industry best practice for a long time. Understand that in the cases of access to – let's just say healthcare records or court and criminal data – there's CJIS [Criminal Justice Information Services] regulations about it, there's rules why CLETS [California Law Enforcement Telecommunications System] data needs to be controlled. So it's a matter of stepping up to the new regulations.

For me, it was education and then getting them on board. I won't say they're all on board yet, but the grumbling has died down. And they now have been through enough demos that they understand how the system works. They've all requested access to things, got it granted, they've seen how it works. There's still a little fear of a black box that controls identity, but that'll probably exist.

[00:03:02]

You've probably seen the same thing.

John: Absolutely. Yeah, it's taken about a year, I would say, from introduction of the concept of identity management being more centralized and not being a component of your application directly for them, through a number of training sessions, presentations, and thinking about the benefits of that identity management piece, about a year.

Tony: And in the case of many of our applications, I actually brought in a third party security assessment vendor so that it was purely independent. And I said, "Great. Take that application. Go through it. Figure out how to breach its identity." In most of our legacy apps, that was doable. Then we'd go back to the team and say, "Okay, wouldn't it be nice that that's not your responsibility anymore? You can focus on the business logic around whatever your application is tuned to, and that that identity is elsewhere?"

[00:04:01]

Again, that was helpful in them seeing that maybe they didn't have that in their skill set, that they were great at the SQL and the information around how that business logic flowed, but that identity is a very specialized – particularly in this anytime-anywhere environment.

John: So I think what Tony did is something that I haven't been brave enough to do, is actually have a security assessment of their present way and recognizing how many of the users that have access to their system shouldn't have access anymore because those particular users have retired, or they've moved to another department, etcetera, and kind of recognizing that management of that identity at the application is not the place to do that. It really should be managed in a centralized manner, and it should be managed by the people – it should be updated and approved by the people that are managing the individuals.

[00:05:01]

Tony: The other thing we did was change our security focus to a thing called GRC: governance, risks, and compliance. So for all of our systems that we have in house, we have an entry and basically a risk table, which maps our total county risk profile. In the risk table, it looks at the system, the owner of the system, obviously, and then what kind of data classification. Are we talking about regulated data, and to what extent? Is it CLETS data, if it's regulated under CJIS? Is it healthcare data that is county confidential, or is it healthcare data that is basically HIPAA [Health Insurance Portability and Accountability Act] covered?

These generalized guidelines that say, "Oh, everything in healthcare is HIPAA covered," become false from the start, but as soon as there's a slight grey area, then people's enforcement across the board becomes not uniform, because they say, "Well, that piece of data wasn't HIPAA covered, so maybe that one isn't."

[00:06:00]

So although you train people on how to be sensitive to what PII [personally identifiable information] and EPHI [electronic protected health information] and all that means, it doesn't mean it applies to each system. What we did was we mapped the GRC for that. Again, what kind of data is it classified? We also map when was the last penetration test on it? What's the population of users, internal and external? And what are the controls that are associated with that system to enforce that authorization?

Once we map all that out, it was very easy to talk to the owners of the resources and say, "Look, there's a little bit of opportunity here. This hasn't had a security assessment on it since it was created or in the last five years. Did you know that these regulations apply to that particular service that you're offering? And, oh, by the way, have you reviewed the application in the context of those new regulations?"

Again, using the GRC matrix was another way to get the owners of the system and the data to kind of step up. And they were all willing to, although funding is always a challenge, to step up to look at it.

[00:07:03]

Di: So going through that process, which might be called a privacy impact assessment in some of the Global products, or a risk assessment, but in Orange County lingo you're calling it the GRC? What did that enable you to do in terms of stratification of different levels of identify proofing? Are you able then to conclude that analysis by mapping some of these data resources to particular levels of proofing? You mentioned a token or a grid card.

John: A grid card, two factor authentication.

Tony: So things that it enabled us to do were to look at a data classification policy for the county, because previously we didn't have one. Many organizations think that they understand their data classification, but in our policy we said, "Look, if this document doesn't say right here in the corner that this is public, this is county confidential."

[00:08:11]

So you establish what's the default on a document? Many organizations think that just because it's there, I guess I can share it. It doesn't say confidential anywhere. It doesn't say private.

So what we were looking at was getting that clarity about the data classification, and what are your obligations relative to that, and not only the user's obligations, but the system owner's obligations, so we introduced terms such as the data owner, the data custodian, and made very clear about what your roles are in that policy. Again, that wouldn't have come out if we hadn't done the GRC.

John: Just in terms of the federal government, they have NIST [National Institute of Standards and Technology] standards, levels 1 to 4, for both identity proofing and for authentication strength.

[00:09:03]

So you kind of assess the overall risk level using their scale, and then you match the appropriate credentialing in your identity proofing before issuing a credential to someone to a particular level as well as the other NIST levels, which are basically how strong your technical authentication is. Then basically, they can restrict access to resources until you're at a particular identity proofing level and a particular authentication level. If they all match, then you get access.

Tony: If you're trying to get to your personal e-mail on an Outlook web server, the risk is very low. If, on the other hand, you're trying to get to a system that is going to allow you to control the infrastructure or the network, now I need to dial up the controls on you. Now, you're onto a real token, not to just logging into the portal. Again, using that data classification was very important in terms of figuring out how the policies inside the engine had to go for granting access.

[00:10:08]

But again, this is all part of that cultural thing that I think we talked about a minute ago that takes a while for people to become comfortable with. "Why are we changing from the way we've always done business?" My argument was, "We're changing because the regulations have changed. We're changing because you used to do business in your office, and if you want to do business in your office, that's fine. But you're now saying you want to do business on your iPad at home, so ***you're*** changing the game. It's not just me, so let's talk about this in terms of you want to change, I want to enable your change. What can we do to come up with a new set of controls?"

Di: It seems to me that there might have been some potential information sharing partners in the county who at one time felt that there was no way they could share information.

[00:10:57]

But now what you're able to demonstrate to them is a level of sophistication and a level of technological enforcement where you can really overcome their resistance. It's not a question of whether we're going to share or not. It's a question of how will we enable appropriate information sharing with the appropriate controls? One classic example is maybe some school officials who believe that FERPA [Family Educational Rights and Privacy Act] absolutely prohibits them from sharing any information with anybody. But now, you are able to bring to them answers that enable them to share information. Is that fair? Going through the GRC and having the data classification policy and having this identity proofing?

John: If reviewing FERPA policy shows that you have a valid business purpose as an external agency for needing access to particular records at the school, then you can refer that back to the schools and say, "And we can give you what you need to know in terms of the user for this particular business purpose."

[00:12:14]

Now we can convey that so that you can only provide records to people that have an authorized business use for that.

Tony: But I think your example is really going to come back to a different question. And that is previously the statement was, "There is no way to technically, securely provide that data to just the users who are authorized." So I come back with the answer that says, "There is now a way that you can do that." We still have to discuss FERPA's subtleties, and we still have to find an advocate to help us get through it the first time, but once we get through it the first time, I think we'll be okay because others will see success and say, "Okay, great, move on."

[00:12:57]

So it's a little different strategy for how to get there, I think, because really we need to have an advocate. They need to know that it works, and that's really what's so exciting about the JUICE [Juvenile Information Content Exchange] program with the Orange County court system is that we've got that strong executive leadership really reviewing all of the documentation about who can share what. We've got this pattern of people that we're already sharing, and now we're just going to speed it up and make it easier for them. So this is an opportunity to show its success, and really the courts are really leading that effort of showing the success for sharing of the criminal data.

Di: So Tony, as you were helping your county IT staff get on board with this new way of doing business, were there any technical resources that you found particularly helpful? Maybe they were technical resources about some of the standards or some implementation or case studies. How did you help them see how the nuts and bolts of this thing were going to work?

[00:14:02]

Tony: This 2007 document that we talked about.

Di: The [Global] Technical Privacy Framework.

Tony: It has been very helpful. [This] was an opportunity to meet with someone who was very influential in its creation, so that was fun. So that was a resource that was helpful: there is a framework out there. The vendor that we're working with, OpenIAM, thoroughly understands the problem and the solutions, and have been doing identity forever. They've been a great resource. On the court side, [Orange County Superior Court Chief Technology Officer] Snorri [Ogata] has got a great development team, Danny and his team, and they've been great, because they found things like the Oracle entitlement server didn't work, so they had to create their own pieces of certain parts of this puzzle and go back and get the GFIPM [Global Federated Identity and Privilege Management] pieces in place.

So really, it's all about, at the ground level, getting the right team in place, and those have been very helpful resources.

[00:15:02]

I would like to say there were more of them out there, but really this is a – very early adopters. We're the pin here, and it hasn't gotten widened yet. There are still only a few people doing this successfully.

John: This is not mainstream yet.

Di: Now that you're operational with your OCID [Orange County Identity] identity provider, what impact has it had on the structure of your IT staff in the county? For instance, what did it do to your help desk? What did it do to the way you structure your IT staff for supporting the identity provider, as opposed to being embedded in the business units, maybe? Can you talk through some of the organizational changes?

Tony: The business units have been more able to focus on the business, as opposed to the buzz around the business.

[00:16:00]

The businesses, many of them have their own IT help desk, and they were dealing with password resets, obviously. They were dealing with changes of staff, somebody coming in there and having to provision.

Basically, [OCID] automates all of the provisioning and de-provisioning. You join an organization, and it has the ability to create your e-mail, set up your accounts, give you access, tell you what systems your prior person had access to that you might need access to. There's a lot of getting the "noise" out of it, and now you can get back to doing your job.

From the IT perspective, there's still that sense of loss because some of them haven't gotten past it yet that that's not a system you control. Their concern has always been that they can't log into the identity solution and make changes. And the bottom line is: that's by design. As the security officer for the county, there are three of us who have a role that we can log in, and we can force something through the workflow to speed up.

[00:17:00]

But none of us can go in there and go directly to the SQL database and make changes. The IT folks are still a little bit upset that they don't have direct access into the SQL tables where they could give John access to all the CLETS data in the county. As I say, I tell them every day, "I don't have access to that. I can put in a request, and I can move it along a little faster, but that's the process."

Di: There is a rumor out there that probably somewhere in the neighborhood of 30 percent of the development cost of a new application involves that user provisioning and the authorization logic. Now that some of that is gone out of the new applications that you're developing in Orange County, are you starting to see some return on investment in that way, that you're able to develop applications more quickly? Or maybe it's too new for that?

[00:18:03]

Tony: I think it's too new for that. We have a couple of projects in place where we are going to be using [OCID] instead of internal application authentication. But the challenge for me is none of that funding is going to come back to my project. They're going to take the budget they put in two years ago for doing it, and they're going to say, "Oh, great! We don't have to do that work," but, by the way, they still have that much funding. It's not a win-win in that regard, but what is a win for the county is the fact that it's standardized, that we can enable-disable the granular rights across the board, and the supportability: if that's a third of the program, the changes and change control, one third of that is going to be standardized now, so that's easier.

Di: There is also this idea then that – as compared to new applications, which are going to be designed with OCID in mind and relying on that – there are legacy applications in Orange County. Talk to us about how you have integrated this new identity provider with your legacy applications.

[00:19:28]

Tony: One of the nice features about the solution architecture that we've come up with and the open source vendor we're dealing with is that OpenIAM supports seven or so different connectors. So that means if you have a legacy application and it used Active Directory as its basis, great: you just take the connector and point to the identity solution's instance of Active Directory, so it has its own data, and it publishes an Active Directory and an LDAP [Lightweight Directory Access Protocol].

[00:20:00]

So legacy applications that used an AD [Active Directory] or used an LDAP, {snap} they're working tomorrow. It's a very simple transition, right?

Di: Wow.

Tony: So that's a real power, because that's how applications in the last couple of years have been: they've been really targeted at LDAPs and Active Directory.

Then you go back a little further in the legacy, and we've got things that had SQL tables, and they had their internal user table that had passwords and credentials and phone numbers, etcetera. So what happens is OpenIAM supports a SQL connector, and it points to their table, and it now controls the table. So again, legacy, no problem. They've got web services.

So we were looking for, on our criteria selection, we were looking for vendors that had a wide range of connectors, easily interfaced to, because nothing like having identity where the users can go in and update their profile, but all that happens is you've got an accurate phone directory.

[00:21:00]

No, we wanted more than that. We wanted the ability to integrate to all those back end systems, and that meant having the connectors.

Di: Tony, have you encountered a legacy application that you have not been able to integrate with Orange County ID?

Tony: I've got one right now, and the issue is that we don't have the source code for it. It was in-house written. All of those things they say don't ever do: don't write it in-house, don't –

So basically, it works how it works. However, what we did is we added a feature to OCID so that you put your credentials into OCID for that app, and it auto launches. You get the experience of a single sign on, but it's not totally controlling the back end. So it doesn't actually – it can't go into the app because we don't know how the app works. So we've got a group looking at it to decide can they figure out – where's the table? How are the elements working? In that case, we're going to end up rewriting it.

[00:22:00]

The business logic on this – and again, this is a twelve-year-old application – so the business logic on it really needs workflow, it needs identity, and a few other reports. I can just have OpenIAM – the solution – has a workflow built into it, it has an enterprise service bus built into it, it has SQL tables in it, so basically I can use those elements and rebuild this app faster than I can figure out how to retro-engineer it.

Di: So would that fall under the category of a wrapper? That Tony's going to be creating a wrapper for this 12-year-old, in-house, no-documentation application?

Tony: The solution we have today is a wrapper. The use of the app safe is a wrapper around it. Where we're going to end up going is I'm going to rewrite it using the modules that exist and are well documented now of the identity solution, because, again, it provides all of the web service tools and pieces that you would want anyway.

[00:23:01]

Tony: We have some 25-year-old systems still, major systems in [Los Angeles] county that run on IBM's IMS database. IMS is a hierarchical database that preceded relational databases, preceded Oracle as a company even being formed, okay? We have a number of those large systems. But IBM has developed a set of tools and adapters that you can get to those IMS databases and expose the information with web services. So, in the case of, let's say, Orange County, the portal and a particular data source: the portal could call the web service, which is calling the IMS database, which is bringing it back, and then you can put in the policy in the middle of there and the roles, etcetera.

[00:24:01]

I guess I would call it another interface into the application that then you can adapt this whole infrastructure as one of the pathways. You could still have the traditional pathway for people that you have now or you can transition them and then turn off their traditional access.

Di: So if you are a technology leader out there, and you're really concerned about the cost of entry into external authentication, what I hear you gentlemen saying is that there are a lot of tools out there to enable the integration of legacy applications with these identity providers, and that they should not be afraid.

Tony: I would think it's a natural fear because I had it myself when we first started the project.

[00:25:00]

I made a list of the apps in the county that had the highest usage and we'd get a lot of bang for the buck on, and I'm like, "Oh, my gosh, what can we do with these?" But knowing what I know today – and again, I don't know that that assures anybody –

Di: It does.

Tony: But the reality is looking at that list today, I go, "Check on SAML [Security Assertion Markup Language], check on AD, check on LDAP, check on web service" – each of them down the table translates to a very simple integration time, low cost, low money.

Whereas before I would have said, “Oh, I don’t even know how to integrate to that.”

So it’s going through that quick application assessment, dialogue that we are doing with each of our agencies, and then saying, “Okay, which of these tools or connectors or adapters do we say is appropriate?”

John: You can have an adapter that’s going to front-end your user ID, passwords. You can have an adapter that’s going to go to the database in the case that it’s a database of internal users.

[00:26:01]

You can have an adapter that actually goes to particular data, and the only challenge on that adapter, in terms of the data exchange, is ensuring that the policies that the core application applies to who and what information you can have access to is replicated then in your authorization engine that you’ve externalized, so that would be an analysis piece that you would need to include.

Tony: For example, you might have an underlying application that doesn’t have a screen timeout. It didn’t have one. Well, then, what happens in the policy for OCID for an end portal is we say, “If that’s the application you’re launching” – because it’s a child application that shows up a separate window in the browser – “then apply a timeout to it.” We can add – in a wrapper through the portal – features that may not have existed in the underlying.

[00:26:59]

Di: That brings you back to the point that the work that you have done makes these data sources more secure than they were previously.

Tony: Right. Plus the fact that we run regular third-party security assessments against the identity solution, which you can’t do when you’ve got that many separated applications. That would cost a lot more. But if it’s centralized, I can bring somebody in on a regular basis, which I do, to run that third party. Back to your keys to the kingdom, this is a very safe key.

John: I was curious: so you do hire security assessment folks to come in periodically.

Tony: I have three master contracts with Foundstone, Accuvant, and McAfee, and they come in and do assessments for this kind of work. Which is part of my obligation as an identity provider: to make sure that's a secure front end.

Di: Periodic evaluation that the system is working as designed.

[00:27:59]

Tony: Although we have tight change control on the code releases to make that there's no hiccups introduced or changes of the code, I still periodically want to make sure that something didn't change.

[00:28:11]
{deleted}
[00:28:16]

Di: Is there anything else that you'd like to bring forward about helping your technological staff through this change, this evolution?

John: The change control process is what it's going to take – it's really marketing to your people and providing them good tools and education so that they can be enabled to take this on. Because whoever is managing this centralized IDP, that's a new role for someone in IT that they didn't have before.

[00:28:56]

In the converse, the person who had that role on each of those applications is looking to that group and trying to understand, "All right, so you're going to authenticate, and you're going to pass me a credential, and instead of me putting up a log-in screen now, I'm going to take this credential and go to the appropriate resource."

Tony: The only thing I would add is about service levels, relative to identity as a service. We've structured a whole set of service levels that the vendor must meet, and we've run through what I'll call use cases for types of calls, when they happen, and then everything maps back to the SLA [service level agreement] table to make sure that the service levels are there. Because identity as a service has to be met. This is not one of those that you can say, "Oh, I'll get to it next week." No, if there's a problem, it needs to be there. The systems have to be redundant. And part of getting the technical folks on board has been that they need to see that.

[00:30:00]

We've done desktop exercises with them, so they see how those SLAs are met, so they can feel comfortable acting as advocates for us with our other users or resource owners. Because if they're hosting the database or the criminal justice system, shall we say, server, they need to be confident that identity is going to be as supported as their server. So they have to become part of our advocate group.

Di: And you do that through very explicit expressions of what those service levels are going to be, and then demonstrations of those –

Tony: Actual desktop exercise of how that SLA is going to be met. What happens? Who does what? Who places the call? Now, we're living it. We haven't had any calls, but they know what to do if it was to have a problem.

John: It's a 24-by-7 operation, once you go down this path.

Di: That makes sense.

[00:30:56]