Carl Franklin and Richard Campbell interview experts to bring you insights into .NET technology and the state of software development. More than just a dry interview show, we have fun! Original Music! Prizes! Check out what you've been missing!

Carl Franklin

Richard Campbell

Text Transcript of Show #599
(Transcription services provided by PWOP Productions)

Andrew Brust Gets All BI On Us
October 5, 2010
Our Sponsor

Lawrence Ryan: .NET Rocks! Episode #599 with guest Andrew Brust, recorded live Sunday, June 27, 2010.

[Music]

Lawrence Ryan: This episode is brought to you by Telerik, by Haystack, and by Franklins.NET - Training Developers to Work Smarter, and now here's Carl and Richard.

Carl Franklin: Carl and Richard here. You're listening to the .NET Rocks! Live Weekend and we're here with Andrew Brust. Hi, Andrew.

Andrew Brust: Good afternoon. How are you guys?

Carl Franklin: Good, good, good. Is the weather is nice and sunny in New York as it is two hours north in New London, Connecticut.

Andrew Brust: It's impressively humid outside. It's also very, very busy because I live about five blocks from the birthplace of the Gay Pride Movement and today is Gay Pride Day so there's a big party outside.

Richard Campbell: Nice.

Andrew Brust: Yeah. Lots of rainbow flags.

Carl Franklin: Half-naked men walking abreast.

Andrew Brust: Walking abreast? Actually, I haven't seen any of that. It's been very...

Carl Franklin: No breasts allowed? Sorry.

Richard Campbell: I've always wanted to be one of two men walking abreast.

Andrew Brust: It's quite a commerce occasion at this point.

Carl Franklin: Yeah.

Andrew Brust: I mean you see Budweiser banners applauding Gay Pride.

Carl Franklin: You know, I'm like Mark Miller. I want to know where's the parade for guys who want to do it with two girls at the same time. Where's my parade?

Andrew Brust: Probably at the other side of the Hudson River I'm guessing.

Carl Franklin: Can we have a flaming heterosexual parade? Is that possible?

Andrew Brust: We could try.

Carl Franklin: All right. I'm done.

Richard Campbell: Okay. Well, thanks. Thanks so much for calling in, Andrew, it's been really great.

Carl Franklin: Yeah, it's been great.

Richard Campbell: What have you been working on these days, my friend?

Andrew Brust: I have been working on a collection of stuff. My latest passion is to give myself a makeover as a wise guy.

Richard Campbell: Aha.

Carl Franklin: Wow. Because you're not one I would take as a wise guy, for a wise guy. But I guess maybe living in New York, you know, it's sort of seeping into your soul a little bit, the wise guy thing?

Andrew Brust: I could hold it back for only so long. But yeah, I've been, for over a year now, writing a column called Redmond Review which I've been using as a bit of a soapbox on various issues affecting, you know, now at the API level but affecting Microsoft and Microsoft partners and customers and the whole ecosystem.

Richard Campbell: Sure.

Andrew Brust: And trying to blog a lot more actively too. I have this secret plan to ask Visual Studio Mag and Redmond Developer News to have their websites carry my blogs, also that I would be forced to post more often.

Richard Campbell: Right.

Andrew Brust: And that has worked out pretty well. In this week, I actually went down to the city council and testified at a hearing on Open Government Data.

Carl Franklin: Open Government Data.

Richard Campbell: What does that mean?

Carl Franklin: OGD for short?

Andrew Brust: Well, I suppose so. Actually, Microsoft has a technology called OGDI for short, which is the Open Government Data Initiative, and that did enter into my testimony but I pushed that out until the end so that I could focus on more non-

partisan issues if you will. But yeah, the city councilor is moving to create a law that would have all city agencies and you may not know but New York City has 30 or 40 Mayoral agencies within its government.

Richard Campbell: Wow.

Carl Franklin: Yeah.

Andrew Brust: I think we're the fourth largest, or third largest municipal or state and local government where our city government is bigger than most states as is the budget. So the idea is that basically every agency needs to publish all of their sort of non-secure, non-privacy related data in raw form. That's proposed legislation and I think it's a great idea but it's something that the council has been working on for over a year so...

Carl Franklin: So it's imperative that these agencies share data and because of that they need to define Open Data protocols? Is it about protocols or is it about accessibility?

Andrew Brust: Well, it's mostly about accessibility although protocol enters into it and politics enter into it so Google was down there testifying in person. Microsoft was not although supposedly they submitted a testimony in writing. But Google is very happy to host everything for free which makes sense because then they can crawl it and leverage it. It's mostly not about format and not about protocol. It's mostly just about the mandate that the agencies must share their data so that really citizens can have a more transparent government and entrepreneurs can have access to the data and create products that add value to it.

Richard Campbell: I mean obviously there's some work here on the government's side. They've got to decide what is private and what is sensitive and shouldn't be published. I mean, you're talking about every council meeting, every kind of meeting and every document produced by government having to be reviewed before it can be published into the system.

Carl Franklin: It sounds like a job for OData.

Richard Campbell: And that's what the Open Government Data Initiative or OGDI is actually based on. It's based on OData, it's based on Azure.

Carl Franklin: Beautiful.

Richard Campbell: What that means, of course, is that the data comes out in AtomPub-based XML format using a collection of open web standards including REST and HTTP. What's nice about it, and that a lot of other governments apparently are not

really thinking about is that it's not just about publishing aesthetic feeds but it provides an API right away so that data can actually be queried and the data can even be updated which makes sense. You can imagine the service where people could report crimes online. Imagine, for example, you had an app on your phone that allowed you with a tap of a button to report a crime happening right at that moment and perhaps the GPS in your phone would alert the authorities to your precise location, maybe the camera could even enter into it.

Carl Franklin: Well, I don't know if that would work because then somebody has to vet what's the crime and what's a guy with plumbers crack.

Andrew Brust: That's absolutely true. There would have to be...

Carl Franklin: Which is a crime, but it's not...

Andrew Brust: There would have to be some algorithm like noise reduction.

Carl Franklin: Yeah. I can't tell with an algorithm.

Andrew Brust: That's true over the phone too, right? It's just that you don't quite get the same scale.

Carl Franklin: That's right.

Andrew Brust: So there's lots of data that I think is not controversial. I mean, potholes, crime statistics at least are not all kinds of economic data.

Richard Campbell: It still feels like a booby trap, you know. If you're too casual publishing the stuff, you're going to publish something that people are going to decide was sensitive and you're going to be in trouble. If you take the time to review it, then you'll take a long time to publish the stuff and you're going to be in trouble. Like I feel for the government in the sense that there's no good way to pull this off, you're going to have problems.

Andrew Brust: Well, that's the story of being in-charge of infrastructure. Right?

Richard Campbell: Yeah.

Andrew Brust: Nobody wants to be the person that brought in the new phone system.

Richard Campbell: Because if it works perfectly, nobody can tell. If it probably doesn't work perfectly, then it's all your fault.

Andrew Brust: That's right. The perfect success is if nobody says anything.

Richard Campbell: Right.

Carl Franklin: I know SQL Server makes it easy for you to publish data through OData, but what about other databases?   What about Oracle databases? What about MySQL? What about non-relational databases? Is that a nightmare waiting to happen?

Andrew Brust: If we're talking about OData specifically, there are a couple of ways to think about it. I mean, first of all, if you're on the Windows platform in terms of exposing the data, then anything that has an enterprise framework provider to it --and so far that's mostly in the Microsoft sphere...

Carl Franklin: Uh-hmm.

Andrew Brust: But there have been loud rumors that Oracle is coming soon. But anything with an EF provider would be exposable. Beyond that, what the guys have actually done with OData is they've decoupled it from Windows and from enterprise framework and they basically said here's an open protocol that you can implement natively, and actually IBM went ahead and did that with their DataGrid technology. They have a DataGrid, you know Cloud DataGrid product, and basically decided that they would implement their REST interface using OData as the protocol. My understanding from talking to folks on the team is that the OData team itself didn't even find out about this until they saw a video about it on YouTube.

Richard Campbell: Ha.

Andrew Brust: That's how independently at Microsoft you can do this. Now having said that, that's the only example I know of a non-Microsoft platform doing that but it does at least prove the concept.

Carl Franklin: This portion of .NET Rocks! is brought to you by our good friends at Telerik. We've been blown away by the uptake in the quick adaption of Silverlight. It's no secret though. The platform didn't provide for consistent integration with the Web Analytics Services. Well, not anymore. As you might have already heard, Microsoft announced the Silverlight Analytics Framework which solves the abovementioned problem. But what's also interesting is that Telerik already provides support for the framework. Telerik is the first UI components vendor to offer handlers for the Silverlight Analytics Framework. Using RadControls for Silverlight, you can immediately benefit from the advantages of the platform and start tracking the statistics of your applications. You can read details and download the handlers at telerik.com/silverlight.   And hey, don't

forget to thank Telerik for supporting .NET Rocks! on their Facebook fan page, facebook.com/telerik.

[Music]

Carl Franklin: I feel sorry for the .NET developer that has to charge with going into some, I don't know, dv2 or some kind of -- maybe even a non-relational database but certainly a non-Microsoft database or a non-touchable database and having to implement OData. If you have access to tables and what-not, do you think it would be a fairly straightforward thing to do if you're going to do it yourself?

Andrew Brust: Well, I think you have to define the premise. If you've got somebody who's a real XML nerd, then yeah, I think it's a fairly straightforward thing to do.

Carl Franklin: Yeah. You really have to be that guy.

Andrew Brust: Yeah, yeah. But you don't have to be a .NET developer. The point is the OData interface for IBM's DataGrid was not implemented in .NET and it need not be because all this is about is pushing the data out in the right AtomPub, extended AtomPub format, and interpreting the URLs with the query adaptions strong at the end of them the right way.

Carl Franklin: Right.

Andrew Brust: And that's what's so nice about it. It's that it really is decoupled from a particular platform.

Carl Franklin: Pete Brown says from Skype, "The killer app for Open Government Data, there's now a public OData feed of all the red light cameras in DC and speed cameras.

Andrew Brust: There's a bunch of District of Columbia data on OGDI. There are a couple of federal agencies, US federal agencies. It's a little bit random but the City of Edmonton in Canada has decided and has implemented their own Open Data service on a GDI. So there are real live examples out there.

Carl Franklin: OData is awesome. I mean, I can't believe how widely it's been adapted in such a short period of time.

Andrew Brust: Well, I mean my observation is that it has been widely adapted within Microsoft and I'm not saying that to be coy. I mean, we all know those two work closely with Microsoft. No, it's not a monolithic organization.   All the different product

teams within the company tends to work shall we say independently and the idea that so many of them have standardize on this, to me that's a world in miniature right there that is kind of united around something. To the outside world, it just kind of looks like, wow, okay, great. Microsoft standardized on its own technology. A big whoop, but it is a big whoop. Reporting Services in 2008 R2 can render all of its reports and not just the text data but even the chart data can be rendered as OData, all SharePoint list in SharePoint 2010, as you know the table storage in Azure. If you didn't know, the SQL Azure now, when you push your database up there, there's literally one checkbox you can check off and it will publish the entire database as a set of OData feeds.

Carl Franklin: So cool.

Richard Campbell: There   are   external companies...

Carl Franklin: Yeah.

Richard Campbell: External to Microsoft that are supporting OData.

Carl Franklin: Netflix.

Richard Campbell: I mean, obviously Netflix is doing it but I think they've got a fair bit of Microsoft technology involved as well, but I mean I've been looking at pieces around how IBM's REST engine is working on the OData protocol.

Andrew Brust: Right.

Richard Campbell: So there are other groups out there that are saying this is not a bad thing. Isn't Microsoft submitting OData to the internet taskforce or trying to get it and made as a generalized standard, not necessarily company specific?

Andrew Brust: If they're doing that, I wasn't aware of it. I would expect that they will at some point, but that I think they wanted to do more with the OData standard and more innovating before they put it through the bureaucratic process of submitting it to a standard organization. My guess is that that's the ultimate goal.

Richard Campbell: Right.

Andrew Brust: But as soon as you do that, everything slows down. So you really want to get...

Richard Campbell: You want to get as mature and in use as possible.

Carl Franklin: And as pointed out by Fritz in the IRC chat room, Dallas is really the killer feature of OData.

Andrew Brust: Uh-hmm.

Carl Franklin: The killer app.

Andrew Brust: What I love also about the OGDI, the Open Government Data Initiative, is that it wasn't even developed by a team in Redmond. It was actually the field developer and platform evangelism group that is dedicated to the US public sector. They're the guys who build it, and they build it not as a product but really as a whole code base starter kit for building your own -- for governments to build their own Open Data platform and that code is completely open source.

Richard Campbell: Interesting.

Andrew Brust: There's a lot of -- I mean, you know, people outside the Microsoft world may agree that's with cynicism or with skepticism but I think there's a lot of selflessness on this. I think this is, you know, if you look at OData, first of all I think the team is very earnest about making it truly open and not having a commercial agenda to it and just really getting it so that we have something on the internet that can be fully expressive of all relational data and accommodating of all relational data operations, and then if you look at OGDI it's just done to help governments get online faster.

Carl Franklin: You know we really haven't talked about Dallas and all the goodness that lie therein on the show before, Richard. Maybe it's a good time to talk about Project Dallas. Andrew, what's been your experience with it. Maybe you could just lead off with telling everybody what it's all about.

Andrew Brust: In effect, it's a way to leverage OData and get your data feeds out there and do it in the Azure infrastructure.

Carl Franklin: Yeah. And there are so many feeds that are available within this group of data feeds that will blow your mind that you can just start consuming right now.

Andrew Brust: That's right. The idea is to have this be really, you know, I hate to use Microsoft competition's term to define their own offering, but this really is sort of an app store for data feeds. Right?

Carl Franklin: Yeah.

Andrew Brust: The idea is if you've got their -- first of all, if you want to make data public, fine. If you want to charge for it, here's the platform to do that in a

reasonable way without you having to build a bunch of infrastructures yourself. So it's the same sort of idea that smaller concerns can get online quickly and can get in front of a really large consuming audience.

Carl Franklin: Here's one right here that just went up. 2006 and 2007 Crime in the United States. data.gov provides extraction of offense, arrest, and clearance data as well as law enforcement staffing information from the FBI's Uniform Crime Reporting, UCR, program.

Andrew Brust: Yeah. There's a bunch of United Nations data up there too.

Carl Franklin: Right. UNESCO statistics, info USA business analytics, US, UK, and Canada.

Andrew Brust: Right.

Carl Franklin: And some of these are free and some are not. Is that the story?

Andrew Brust: I think at the moment, because Dallas is still a codename and the project is still in its formative stages, I had thought actually that for a limited time at least everything is free but eventually there will be commercial subscription concept up there as well.

Carl Franklin: What a great way to do mashups of...

Andrew Brust: Absolutely.

Richard Campbell: Yeah. The next generation of mashups. I totally agree.

Carl Franklin: Yeah, that was your... I think you said that about OData earlier in the day.

Richard Campbell: Yeah. Hey, Andrew, C# Fritz on the IRC chat is saying "As a software architect on a large multi-tenant web-based application, my primary concern with OData services is security."

Carl Franklin: Hmm. Indeed.

Richard Campbell: "Can you speak more about how to handle authentication in OData?"

Andrew Brust: I was thinking that that question might come up even sooner, but here is it now.

Richard Campbell: How did it take so long?

Andrew Brust: Right.

Carl Franklin: Because it is an Open Data protocol and you can do CRUD operations.

Andrew Brust: That's right.

Carl Franklin: I'm thinking it's IIS.

Andrew Brust: You're thinking. Yeah. I mean, it's sort of the same story that SOAP web services where when they first came out which is...

Richard Campbell: Right.

Andrew Brust: Right. There's no built-in security but it's just HTTP, man, so use SSL and use HTTPS or use your own schema on the server and you can authenticate people. I know that the team is working on making that story a lot better and I think that's one big reason it hasn't been push out to any standard organization yet. I mean that, you know, the team is well aware that not all Cloud data services will be public Cloud data services.

Richard Campbell: Right. Yeah. Crime stat is going to be freely available, but there are plenty of other things that aren't going to be. So how do we -- it's not just securing it, it's also how do we have an effective building engine around it, how do we make sure that it's used properly. Those are fairly tough questions to answer.

Andrew Brust: Absolutely. But I think the solution that comes out will be robust but it's definitely, you know, it's iterative and the first and second iterations involve the premise that the data is open. Most of it though aren't going to accommodate updates and so forth. I mean, you can if that's what you wanted to do but it's sort of an all for nothing premise.

Carl Franklin: Right.

Andrew Brust: Not a roll-based premise for that right now.

Carl Franklin: You know it's pretty interesting how RSS as a protocol was really utilized by journalists and newspapers, magazines, blogs, that kind of stuff where you've got these short stories. But OData really takes that idea and encapsulates all kinds of data and now you're seeing what we thought we're going to see with RSS which is the widespread adaption of a XML-based protocol for exposing sets of data that can be queried over REST.

Andrew Brust: Right, and we're seeing it instead based on AtomPub.

Carl Franklin: On Atom, yeah.

Richard Campbell: It's interesting.

Andrew Brust: Yeah.

Carl Franklin: Because Atom is more robust.

Andrew Brust: And I remember kind of when RSS went from 1.0 to 2.0 and a bunch of sites also had Atom, I was scratching my head saying why even bother.

Carl Franklin: Right.

Andrew Brust: And then it turns out the dark horse kind of won. You know, David Winer might not be so please with that but there it is.

Carl Franklin: Yeah.

Andrew Brust: Really, you know, the power of standard is in everybody agreeing upon it and using it.

Carl Franklin: Right.

Richard Campbell: And you're exactly right. RSS went that way too. RSS blew up on its own. When it came out so many people implemented it, it effectively became the standard. It had nothing to do with its quality.

Andrew Brust: Right.

Richard Campbell: It had nothing to do with its feature set. It had to do really with its simplicity which comes back to its name.

Carl Franklin: Jersey Shore says, "OData is the new cables."

Richard Campbell: Tables?

Carl Franklin: Cables in quotes. Do you mean like hooking a cable from one thing to another, sort of like a virtual cable or are you referring to something specific? I'm not sure.

Andrew Brust: I don't know. Yeah, maybe it's a hooking stuff together reference.

Carl Franklin: Daryl Miller says, "In theory, authentication and authorization should be orthogonal issue to the OData service. Use whatever security your underlying protocol supplies."

Andrew Brust: Right.

Carl Franklin: "In this case it's HTTP, in the case of Microsoft stack it's IIS."

Andrew Brust: Right. Which is why I said this is exactly the same thing as the first generation of SOAP Web Service. Right?

Richard Campbell: Yeah.

Carl Franklin: Yeah.

Andrew Brust: Sort of like this shrug off like, well, that's not the point. I mean, security is handled by the protocol we're running on. But then came the WS-Star standards for making things secure.

Richard Campbell: And we all know how well those went.

Carl Franklin: He means cable...

Andrew Brust: They were there because they needed to be there. I mean, arguably that's what made REST so popular when SOAP got so complicated.

Carl Franklin: Yeah.

Andrew Brust: But for enterprise use, they were so absolutely necessary.

Carl Franklin: He means cable as a journalism term.

Richard Campbell: Like telegrams?

Carl Franklin: Yeah, like telegrams.

Andrew Brust: Oh, oh, oh. Wow.

Carl Franklin: I missed that one.

Richard Campbell: I don't know. It seems like a stretch to me, but okay, I won't disagree.

Carl Franklin: Well, it's certainly, you know, the metaphor is there as a glue, as a way to hook up data from one -- real data. I mean, this isn't just publish a list of items or something.

Richard Campbell: Yeah.

Carl Franklin: This is a queryable data source.

Richard Campbell: Chops of data, yeah.

Andrew Brust: Absolutely. And this is where all my little worlds started coming together. You guys know I've sort of been in the developer world for a long time and the BI world for a goodly amount of time but not as long...

Richard Campbell: Sure.

Andrew Brust: And now there's thing called PowerPivot; and it's tying all these things together.

Carl Franklin: Yup.

Andrew Brust: Because PowerPivot is the sort of rich client consumer of OData feeds including Open Government Data Initiative fees and Dallas feeds. Basically they let you bring all this data into a model. We can talk about what a model really is, but right now we'll just call it a model and then eventually push it into Excel and do real analytics on it. So not only is the data out there, but you can actually start pulling stuff in and breaking crime statistics down by zip code, or neighborhood, or council districts, or economic strata, or what-have-you. It becomes very, very simple to do that.

Carl Franklin: This portion of .NET Rocks! is brought to you by the Haystack Code Generator for .NET, Code Generation on steroids. Want more control over your Code Gen? You want your code generator to give you Silverlight 4.0, WPF, and ASP.NET CRUD screens? The Haystack Code Generator for .NET will generate entity, data, and business rule classes for all your SQL Server and Oracle tables, views, and store procedures. Haystack generates ASP.NET, WPF, and Silverlight user controls, View Model classes, and WCF Service Layer classes for true and tier applications. Check out codehaystack.com, download the user manual, and watch the videos from more information on this great product. They host a live webcast every two weeks. You can sign up at pdsa.com/webcast and see how Haystack will shorten your development cycle.

[Music]

Richard Campbell: It makes sense with your BI background looking at OData as sort of a new form of ETL. You're changing this...

Carl Franklin: What was the acronym you just use?

Richard Campbell: ETL?

Carl Franklin: ETL?

Richard Campbell: Yeah. It's...

Andrew Brust: Oh, ETL. I thought you said VPL, Visible Panty Lines. You completely drove me...

Richard Campbell: Visible Panty Lines?

Andrew Brust: That's given where this conversation began...

Richard Campbell: What are you talking about?

Carl Franklin: Are you thinking about that parade conversation we were having earlier?

Andrew Brust: That's my point, yeah.

Richard Campbell: ETL, my friend. Extract, Transform, and Load. Any BI guy knows that.

Carl Franklin: We're not BI guys.

Richard Campbell: He is. Only you.

Carl Franklin: But the listeners aren't. Explain your acronyms.

Richard Campbell: There you go.

Andrew Brust: Extract, Transform, and Load. What really do you mean?

Richard Campbell: What I mean is how much work did we do as BI folks when you want to set up a data warehouse? This was all about bringing desperate data sources together, extracting them from the source, transforming them so that you could combine those data together and then loading them into a common sets so that you can analyze them.

Carl Franklin: Oh, yeah.

Richard Campbell: And OData turns this on its head.

Carl Franklin: That's essentially what we've been doing for the last 10 years as .NET developers anyway.

Richard Campbell: It's trying to organize data so that we can actually combine it in useful ways.

Carl Franklin: Yup.

Richard Campbell: So the idea that we put in OData feed out there so that anybody, you skip the E and T part pretty much, just go to the L, go load up any data you want.

Andrew Brust: Right. But that's an issue. It's that there's not much transform available and it's all like PowerPivot right now.

Richard Campbell: Right.

Andrew Brust: I mean there is in that you can create your own formulas and a brand new formula

language called DAC, and you could transform formulaically I suppose but there's not much opportunity for doing serious data cleaning on that.

Richard Campbell: This is where you're going to get into whammies because you're going to go grab an OData feed on crime stats and you're going to want to try and match it to some census data on income levels and there's no like set of GIS data, there's no like set of time ranges. You know, the crime stats are given date and the census data is from a given year. All of those things become challenging. This is where you get into this transform side of the problem. I'm trying to line things up so that they meet up in a useful way.

Andrew Brust: True.

Carl Franklin: Yeah. It's almost like you need to agree on data types all over again.

Richard Campbell: Right.

Carl Franklin: Isn't there a standard organization out there that's trying to define standard types for all kinds of data like this?

Andrew Brust: That would be a good thing.

Carl Franklin: It's not Biztalk.org?

Andrew Brust: I'm not aware of it, but I mean you mention the census. Right? The census have this concept of, oh, I think it's called the census name place. So when you hear census data about a town or a village or something like that, that's what it's called. But that entity is not necessarily code terminus with an actual incorporated village, or municipality, or zip code, or postal address. All these things are kind of overlapping but not quite align.

Richard Campbell: Yeah.

Carl Franklin: I think that was one of the problems that Microsoft had with the metadata problem of WinFS. Do you remember they were trying to find, come up with metadata specific to different types of files and different pieces of data that you could search on and all that kind of stuff. Am I dreaming that happened?

Andrew Brust: I don't know how you got this back to longhorn from PowerPivot and the census.

Richard Campbell: Wow.

Carl Franklin: Well, no. But it's all about associating a metadata or having types that are standardized...

Andrew Brust: Right.

Carl Franklin: So that you can take globs of data from all sorts of desperate things and hopefully jam some metadata into them or around them that is a standard recognized type.

Andrew Brust: And by the way it's census designated place, not in census name place. I apologize. I think there's even the CDP as an acronym for it. But yes, you're right. Standard taxonomies are somewhat a contradiction in terms with a lot of public data right now so it's an issue and probably a good opportunity for entrepreneurial software companies out there who can introduce some solutions to help clean it up in kind of an automated way.

Carl Franklin: And Daryl Miller says from Twitter, "High traffic websites survive because of caching. With OData's highly variable query capability, how will OData services handle high load?"

Andrew Brust: Wow. I think the loads are going to be pretty -- I mean, the ability to handle high loads are going to be pretty well correlated to the infrastructure. I don't know that OData is going to be the issue one way or the other. If you're hosting it on your single web server of your DSL line, you're probably going to have some issues if you're using Azure infrastructure or for that matter Amazon Web Services. There's nothing about OData that marries you to Microsoft stack really, then you're going to have better results if I understood the question correctly. If it didn't, please smack me.

Richard Campbell: Yeah. Do you think that Daryl Miller's question is about the caching and the querying capabilities? This is ultimately an implementation issue. I don't think anybody is necessarily committed to a given implementation per se. So it's just a standard. This is the same question you would have against SOAP, or RSS, or anything.

Andrew Brust: Right. It's an XML format and a wire protocol.

Richard Campbell: Right.

Andrew Brust: That's all it is which is both good and bad, but we got to start at that low level and then try and get some consensus around it. So apparently Google has something rather similar which is unsurprisingly called G Data, and I think they're both actually based on the same idea of taking AtomPub and extending it so that you can have in effect a user-defined schema that aligns up with the columns and tables and databases, or sometimes it's less structured data but you still structure it that way. So it would be interesting to see if it will emerge or

live on unpeaceful coexistence or peaceful coexistence. Google is definitely proactive in terms of hosting government data and services because actually Google Maps is integrated. Like the first time anybody has every done this, it's integrated with New York City transit authorities subway map and scheduling data.

Richard Campbell: Right.

Andrew Brust: So they definitely got a head start.

Richard Campbell: Well, and they disintegrate to be going on the IRC channel around the whole caching side of this so now you get into data variability as to how often is the New York subway schedule going to change, can we actually propagate cache items out there well, and who does the querying. Do you just pull the whole set down and query it locally or do we provide services on the backend to limit the amount of data you're hauling around?

Andrew Brust: It's probably going to be all of the above.

Richard Campbell: Yeah. I think it's all implementation at this point. You can choose.

Andrew Brust: It it's an iPhone app, I bet a lot of the data comes down. If it's something more browser-base, then it will probably be a lot more dynamic.

Richard Campbell: Yeah.

Andrew Brust: But again, the good thing about OData, it's not just the DOM, it's not just the flat file DOM. It is XML serialization of live data coming out of the data source.

Richard Campbell: Right.

Carl Franklin: So you have an Android phone, Andrew?

Andrew Brust: I do, I do.

Carl Franklin: Andrew Android.

Andrew Brust: I have the original droid actually as my phone.

Richard Campbell: Wow. The Motorola one.

Carl Franklin: And an iPad.

Andrew Brust: I do, both.

Carl Franklin: You know I actually took a fake, not fake but non-deserved offense to a caller who I thought was saying to Charles Petzold "You know, how come you don't have a Windows Phone? You're like the Windows guy." You know, I thought he was calling him out because he had an iPhone or something, but really what he meant was how come Microsoft hasn't given you a Windows Phone 7, and I misunderstood because I'm always getting crap from people who say "You're a Microsoft guy. Why do you have an iPhone?" And I say "Because it's a good phone, because I want a good phone." You know, loyalty goes so far. I'm loyal to Microsoft for the things that I like, like Visual Studio and .NET.

Richard Campbell: Well, how many iPhones are held by Microsoft employees?

Carl Franklin: They love them.

Richard Campbell: Yeah. We're walking around campus a few weeks ago and they've all got iPhones.

Carl Franklin: Right.

Andrew Brust: Despite the Wall Street Journal article that basically claims that any Microsoft employee who had one, in addition to not getting reimbursed, also took a bunch of political sort of heads for it which I think was nonsense.

Richard Campbell: Yeah. I don't think it's true at all because there are some pretty darn senior people inside of Microsoft with iPhones.

Andrew Brust: Yeah.

Carl Franklin: So is Microsoft capable of pulling off the consumer device too that everybody thinks and hopes that they will?

Andrew Brust: They're definitely capable of it. They're also capable of dropping the ball so the question is which way will it go. I mean, it's going to be tough because Microsoft is, and it always comes back to Android. I'll explain what I mean in a second, but Microsoft is sticking with the motto of working with OEMs which means it's not going to control...

Richard Campbell: The phone.

Andrew Brust: The hardware to a micro-managerial level. They definitely have a very serious set of minimum requirements. It's a pretty high bar so I think the phones will call less together. They'll have a lot in common and the overall quality will be high, but it's still not obviously the same degree of control as the Apple has. As we have now learned even in Apple's case, even when you do have full control, I guess it can happen like antennas that don't work if

you hold your phone in your left hand. My biggest concern is how good the touchscreens are going to be because I think Apple has really set the bar there and I think the best we can do is meet it. If we can exceed it, that would be wonderful but I can't just imagine that. So with all these different OEMs, you know, are all of them going to meet that standard or not? You know, I worry that some of them won't.

Richard Campbell: I can't imagine -- I think I said this earlier today, I can't imagine more than four different Win Phone 7s. Just because the spec is so high, it's going to be incredibly hard to build a phone that has unique features.

Andrew Brust: That's right and they're not allowing -- not only is the spec high, but they're not allowing OEMs to put their own sort of customize UI on top of them.

Richard Campbell: Right.

Andrew Brust: Like the Sense UI from HTC or Auto BLUR from Motorola.

Richard Campbell: Right.

Andrew Brust: And as I said it keeps coming back to Android because Android does use the exact same OEM model that Microsoft did. The only difference, and this will sound kind of goof ball, but the only difference is Google doesn't actually charge for you to use their operating system.

Richard Campbell: Yes.

Andrew Brust: Other than that, it's basically the same model. Take it, do what you will, and let's see what happens. Android, by and large, has worked out really, really well.

Richard Campbell: Although the app store is chaos and actually taking a given app and running it on a lot of different phones is hard. Because there are so many different Android phones, sometimes stuff just blows up.

Andrew Brust: And there's at least kind of three very prevalent builds of Android out there, right?

Richard Campbell: Right.

Andrew Brust: There's 2.1, there's 2.2 which is just eking its way out and I guess there's still 1.6 and some of the older and cheaper phones. What I think is really neat about Android is you have this real like flagship phones like the Droid, like the Nexus 1.0, like the incredible Sprint. Then like we had a babysitter come over last night so we could go to this birthday party I was telling you about and she had some $50

LG phone from Verizon that I didn't even know existed with the slider keyboard and everything else and running Android and I use it for a few minutes. It seems really solid so you have some really low-end phones too in terms of price points.

Richard Campbell: Right.

Andrew Brust: And they've actually been able to make the OS kind of scale those both ends of the market which I think is like impressive.

Richard Campbell: Uh-hmm.

[Music]

Carl Franklin: .NET Rocks! is recorded and produced by PWOP Productions, providing professional audio, audio mastering, video, post production, and podcasting services, online at www.pwop.com. .NET Rocks! is a production of Franklins.NET, training developers to work smarter and offering custom onsite classes in Microsoft development technology with expert developers, online at www.franklins.net. For more .NET Rocks! episodes and to subscribe to the podcast feeds, go to our website at www.dotnetrocks.com.