



RUNAS RADIO



<http://www.runasradio.com>



Richard
Campbell

RunAs Radio is a weekly Internet Audio Talk Show for IT Professionals working with Microsoft products. The full range of IT topics is covered from a Microsoft-centric viewpoint.



Greg
Hughes

Text Transcript of Show #143
(Transcription services provided by [PWOP Productions](#))



Kevin Kline on the State of SQL Server and More!
January 13, 2010



[Music]

Brandon Wenn: From runasradio.com, you're listening to RunAs Radio, the Internet audio talk show for IT professionals with Richard Campbell and Greg Hughes. This is Brandon Wenn, announcing show #143, with guest Kevin Kline, recorded Friday, January 8, 2010. RunAs Radio is produced each week by PWOP Productions, providing professional media and podcasting services online at pwop.com. You can follow the boys on Twitter at twitter.com/runasradio.

Richard Campbell: Thank you, Brandon. This is Richard Campbell. You're listening to RunAs Radio. With me as always, my co-host, Greg Hughes.

Greg Hughes: Hey everybody, Richard, once again.

Richard Campbell: I'd to speedy opening this morning. What can I tell you fast?

Greg Hughes: Speedy, speedy Gonzalez.

Richard Campbell: Feeling speedy today, absolutely. Well, without further ado, hey, let's just jump straight into the show. There are lots to talk about here. I've got Kevin Kline with us today. Kevin is the Technical Strategy Manager for SQL Server Solutions at Quest Software. He is also a Microsoft SQL MVP, and the founding board member of PASS, that's the Professional Association of SQL Server folks.

Greg Hughes: Great conferences they put on.

Richard Campbell: Absolutely. And the author of several books including "SQL in a Nutshell" from O'Reilly. Kevin is a top rated speaker in industry trade shows, have been active in the IT industry since 1986, which is just making himself sound old, and a good friend of mine too. Kevin, it's great to have you on the show.

Kevin Kline: Thanks so much, Richard. It's a pleasure to be here.

Richard Campbell: We always end up talking about SQL Server stuff anyway. What's on your mind these days?

Kevin Kline: Well, you know, there are so many things happening with SQL Server that are fun and exciting. You know, one of the things if you paid much attention you've probably seen that there's been a bit of acceleration in the release schedule so we had a really long delay between 2000 and SQL 2005, but we've only had maybe 24 months between SQL Server 2008 and the CTP before SQL Server 2008 R2 come out. You know, that's interesting and that produces a lot of opportunities for us to learn new

Kevin Kline on the State of SQL Server and More! January 13, 2010

things and to see what's coming ahead from Microsoft.

Richard Campbell: I realize that the 2000, 2005 thing was not SQL team's fault. I remember then they were ready earlier and it was the Whidbey, the new UI that was really holding them back. You know, it's funny how those things start to happen, that they get these delays and stuff to thresh out. Now I worry that we're too fast. The folks are like, "Oh my God, another version already? Are you kidding me?"

Greg Hughes: Right.

Kevin Kline: You know, that is a great point and from an enterprise IT perspective, which is my background, really big multi-hundreds of SQL Server shops, this is a huge concern and something that I bring up to my colleagues at Microsoft through the MVP channels and just through personal relationships that are bring up quite often because one of the biggest issues for an enterprise IT shop is just controlling change and change management, and so 24 months is rapid enough for IT shops to start considering maybe we're just going to deploy an update every other release.

Richard Campbell: Right. Let's skip a version.

Kevin Kline: Exactly.

Richard Campbell: You say change and immediately I think of the mantra "Change is good, you go first."

Kevin Kline: Perfectly said.

Richard Campbell: Yeah.

Kevin Kline: So many IT shops they want to see how everybody else fares first.

Richard Campbell: Yeah. I can't afford to be a pioneer, I don't need arrows in my back.

Greg Hughes: Right.

Richard Campbell: Let me know how that works out for you. If it's okay for you, maybe I'll come along as well. Is there stuff in R2 that really matters, you think?

Kevin Kline: I ask myself that same question really and that's something I'm always really keen on because I used to be a very hard core Oracle person, I use Oracle for many years. As the author of SQL in a Nutshell, I really keep track of the pulse, keep my finger on my pulse of both DB2 and My SQL and so I always want to know where they're going and for quite a while with Microsoft and SQL Server I have difficulty perceiving their strategy overall with the new features and new releases because adding more



features sometimes doesn't help stuff and doesn't necessarily make the product cohesively better. It does advances here and there but not cohesively better, and I saw a cohesive strategy with Oracle because they were buying applications and although they definitely had been advancing the database you could see that they had kind of a strategy that included let's get all the ISVs and open big application vendors on Oracle and that's going to drive a lot of database adoption.

Greg Hughes: Right.

Kevin Kline: So what's driving database adoption for SQL Server besides a great price, it actually really crystallize for me with SQL Server 2008 R2. I was a little bit surprised and there was my big aha moment. When they introduced it at the past conference, they showed this new feature called PowerPivot and it's an add-in for Excel and SharePoint and whatnot. And so what they have been doing in release after release after release, of course the database engine itself has been getting better but they have been adding almost an order of magnitude more BI features so each time analysis services, we get much more powerful. A release or two ago, they added Reporting Services and that became very potent way to get your data out. One of the things they saw with the integration with Reporting Services and SharePoint was huge proliferation of power users starting, you know, making use of data but in a very uncontrolled way. PowerPivot is a feature now that's just very, very good cohesion between the relational engine passing that data to the BI engine and that BI engine then handing it off to users, power users who normally would be uncontrolled and now we can control it, we can get them huge amounts of data, we can make sure that the data is up to date, that it's not somebody with an access database in some dark shadowy corner of the organization. We can actually keep tabs on them and kind of make them part of the team. So what they have done, and I apologize for the long discussion there, but what they've done is they really brought BI right there to the desktop. They had talk about it kind of in passing before but the aha moment for me when I saw the demo was that this is really what they have been promising in terms of business intelligence for years, that we can actually enable the best information -- or what's the old term, the right information to the right person at the right time.

Richard Campbell: Right.

Kevin Kline: It really does that.

Richard Campbell: Well, Microsoft always seems to struggle with the client side of BI for some reason. They had some opportunities there and I guess they're going to land on Excel and that's sort of like the "ad-hociness" of that, that you can just sort of pull some data, it's not too formal, you could do some

Kevin Kline on the State of SQL Server and More!

January 13, 2010

reporting. But I agree with you, the biggest thing for me is are these data actually secured, are they auditable like do we know who's grabbing data and how.

Kevin Kline: Exactly. You know, and one of the experiences I had, and probably both of you have seen it at some point or another in your career, is somebody, you know, somewhere at their queue business, just a smart user, they start clicking away, they get a password to this database or that and they pulled their own data down, they have their own little Excel spreadsheet or access database, and the manager in the cube down the hall sees who successful that information is and empowering that person and like, "Hey, can I use that too?" Pretty soon there's not just one or two or three copies, there's like 12,000 copies of it around the company.

Richard Campbell: Right.

Kevin Kline: And it's completely out of control.

Richard Campbell: Well, and it's also that's when you discover that he left off a major chunk of the query and they're working from wrong data.

Kevin Kline: Exactly, or half of the people who are using it are using his second iteration and then the other half are using the 15th and so there's a huge disparity.

Richard Campbell: Yeah and he has now become his own little IT department.

Kevin Kline: Exactly, exactly.

Richard Campbell: Oh man.

Kevin Kline: It's just like the change in that stuff we were talking a moment ago. It introduces just uncontrolled change and you can want spending as much time supporting that kind of stuff as the real applications that your bonus plan is based on.

Richard Campbell: Yeah.

Greg Hughes: You talk about SQL Server and adaption with this new, you know, the more frequent release cycle now. My experience has been watching development teams sort of looking forward to versions of SQL Server and the release are about to come and say "Well, what can I do now that it works really well?" SQL Server quite often is really a foundational piece of an application and a lot of the SQL Server or other requirements for building a software application from scratch sometimes are based on what SQL Server has to offer.

Kevin Kline: Exactly, exactly.



Greg Hughes: When the truth is SQL Server 2008 gave us native transparent encryption that really changes the way that people are building secure applications to leverage SQL Server as a store. What are some of the things in the R2 release that maybe sort of fit along on those lines?

Kevin Kline: Well, there are a few things that are kind of the flip side of that coin which I find kind of interesting. But one thing they've added, which they've talked a lot about and which is going the other direction, is something that's called Master Data Services. At first when I look at this I thought, "Wow, this is kind of odd" because Master Data Services is the kind of thing that very few enterprises use. You know, it's not something like every developer would say, hey, I need to build this as a foundation into my product, this new application we're building. What it does, Master Data Services, is that it basically creates kind of -- it's almost as much process and automation of kind of best practices around process for saying what data in this company is kind of the official, the master copy of this data. And so working to declare this particular database over here in its version of customers, that's going to be the master version and everybody else has to derive their data from this master version.

Greg Hughes: Got you.

Kevin Kline: They're about customers. You know, this other database over here, this is the one that, you know, the SAP HR data module, that's where we are declaring that to be our master HR data and this is a kind of platform that I see very few shops using, but I think that this is interesting to a certain percentage of Microsoft customers to -- you know, or this kind of top tier Fortune 50 kind of companies, very large government organizations who have been going out their data to many, many different people and I found that one group of people call Hewlett Packard, HP, and then another group calls them Hewlett Packard, another group HP Ink and so they have this terrible problem with data quality.

Greg Hughes: Right.

Kevin Kline: So that's one side of the equation, not quite the developer's side of things. With the developer side of things, I'm a little bit surprised with there was a really, really strong and forceful effort at SQL Server 2005 release to really strongly integrate Visual Studio and SQL Server together and they were even release on the same schedules and whatnot.

Greg Hughes: Right.

Kevin Kline: You know, I was really excited and hopeful that we would see lots of things like much more integration with change management systems like Team Foundation services, Visual SourceSafe,

things like that and that didn't happen at the time but like you're saying we are seeing more robust features added that the developers can hook into and take advantage of.

Greg Hughes: Yeah.

Kevin Kline: But my story there is, I hate to use harsh words, but my story there is a little bit one of disappointment because there are some really, really amazing new features maybe a release ago. Let's say like Service Broker, very, very powerful means of doing asynchronous messaging and transactions and it's something that developers could really make powerful use of it. Some of these things kind of come out, I don't know if half baked is the term, but you could take it so much further and make it a really strong launching point for developers to build all kinds of new services on. But I haven't seen them take some of these really, really cool things that developers could utilize to the next level.

Richard Campbell: I do find the developers generally don't know that there's anything in SQL Server other than SQL Server. They don't know about Integration Services, they don't know about Service Broker, they just don't know that all these stuff are in there, like there's no excuse. In the latest version of SSIS, there's no excuse for you to write looping code to parse text files ever again.

Kevin Kline: Great point, great point.

Richard Campbell: But they don't know that. We've done shows on the other side, on .NET Rocks! with guys where we're saying like "Look at how good this is. You should not be writing this code by hand anymore."

Kevin Kline: Right and you bring up a great point too there, Richard, in that I think that Microsoft kind of recognizes this in the way they've launched SQL Azure. It doesn't have a lot of those services and also you can see pretty clearly with its launch that really it's targeted at the developer. I mean, they launched it at PDC. They didn't launch it at PASS or at TechEd for a broad and general audience. This is for developers and it really just have the kind of things that you guys think of when you think of a database, just those kinds of barebones sort of things.

Richard Campbell: Absolutely. You know, the other angle on this that I find interesting is that I agree with you that they went to all that trouble to put the SQL Management Studio as part of the same tools at the Studio and then it seem like they kind of drop the ball, then along came DataDude which is the version of Studio specifically for database folks that does allow us to check in the data structures, schema and so forth, as part of the source code and actually make us a full-fledged member of the development team, but



I don't see real DBAs using it because it's Studio and Studio is for developers.

Kevin Kline: Exactly. You 100% hit the nail on the head.

Greg Hughes: Why buy Studio if I don't need to buy it.

Richard Campbell: Well, and I think part of this is a licensing problem that folks are just why am I going to go down that path, and Microsoft has now reorganize all the versions because that was what was happening, it was that you had to buy the database edition. It's like you have to buy the test edition and nobody wears one hat to have one version like that so they've reorg all the versions in 2010. I know I've been coaching developers. You know, the one developer who can write decent SQL and actually get his way through stored procedure, that sort of that database developer guy as oppose to the DBA who works with developers, that's the guy who ends up running this version of Studio and he's generating scripts that then get fed to the DBA to evaluate to see if it actually have been applied to the database. That's sort of the way I've been coaching people to use this tool.

Kevin Kline: And that's what I see too. So if you have a large team and let's say you've got 5, 10, or more developers, there's usually that one really smart guy with the database...

Richard Campbell: Yeah.

Kevin Kline: And he's the one who gets the Team Foundation, what is it, Database Professional Services?

Richard Campbell: Yeah.

Kevin Kline: DataDude. And he's the one guy who gets that and then everybody else are still with Visual Studio.

Richard Campbell: Yeah and they're just sort of gets left there.

Kevin Kline: Yup, yup and you're right, there are so many fantastic features in that product too. It's a shame that they're not kind of release and distributed in a way that the DBA themselves could take advantage of it or even expose in way that a regular developer on the team could use it as well. You know, you bring up an interesting point. I hope you don't mind I segue towards that, but one thing that has been kind of interesting thing for me to watch is around licensing and pricing. One of the things that, you know, if you harken back just a few years, you know we all probably remember the days back when there was one phone company and then it was broken up into several phone companies, Pocket Bell

Kevin Kline on the State of SQL Server and More!

January 13, 2010

and Bell's House and all that, and if we use the old game board from Parker Brothers or Milton Bradley monopoly, where there's just one company that runs the show, and then in Economics 101 we have something else called Oligopoly in which there's just a handful of companies that kind of rule the market if you will. You know, we remember that from the old telecom days, the big three from Detroit, just a few movie studios for example make 80% of the movies in the USA.

Richard Campbell: Right.

Kevin Kline: We really have that situation now very clearly. Although it's not the kind of thing that a lot of people think about, it's very clearly happening right now with major database platforms. So at Oracle, Oracle's big conference earlier in 2009 in the summertime, they announced, oh gosh, 35%, 40% price increases and sure enough in the fall of that year, I believe it was at the past conference, Microsoft announced, hey, here's our price increases. DB2 from IBM has always been priced about the same as Oracle so there was a period of time when you could count on SQL Server being priced quite differently from Oracle but I think that we're seeing we're now entering a period in which it's very competitive with the Oracle in terms of its market share now and so I think we're going to see things acting in a much more kind of oligopoly sort of economic situation which if they raise their price, we'll raise our price.

Richard Campbell: Sure, yeah. I totally agree because it's really only three players that matter, right?

Kevin Kline: Exactly and add in the fact that Oracle has bought Sun...

Richard Campbell: Yup.

Kevin Kline: Or is attempting to so you let that stand and My SQL now who used to be kind of the Southwest Airlines of the price leader in terms of low-end cost...

Greg Hughes: Yeah.

Kevin Kline: You know, they're going to be the Oracle.

Richard Campbell: Yeah.

Kevin Kline: And so it's another situation in which you can say, hey, if you never paid attention in those kinds of economic transfer of large patterns, if you wanted to buy right now a four-CPU high-end enterprise edition copy of SQL Server it's going to cost over 100,000 bucks.

Richard Campbell: Wow.



Greg Hughes: That's a lot of money.

Kevin Kline: Yeah, it's a lot of money and of course that's list price so if you're a big company you can negotiate and there's enterprise agreement and SA's, Software Agreements and whatnot.

Richard Campbell: Right.

Kevin Kline: You know, people think of SQL Server as a price leader and I think it is definitely a price performance leader, but I don't think you can say it's just at least solely on cost now because those prices are definitely going up.

Richard Campbell: I think the other side of this is that SQL Server is so much more performant than the average app needs.

Kevin Kline: Yes and that's one thing I joked/talked about when I was running one of my last benchmark papers, it was I have a little test lab and I wrote a paper and said, "Oh, you know, this and such offers are big performance boost over the old way of doing it," and I had real trouble even on my little two-CPU, two-core Dell box here with actually getting it to fail in pushing it too hard. It was performing so well that it took a huge amount of work for me to find out a way to find out a way to, you know, the bottleneck and make it throttle back.

Richard Campbell: To actually bury it. That's fine.

Kevin Kline: Exactly, yeah. Yup.

Richard Campbell: You know, the other direction that I found interesting in the marketplace now is the No SQL movement.

Kevin Kline: Yes.

Richard Campbell: It's this idea of I don't want a database.

Kevin Kline: And I think we're going to see more of that too and it goes back to what we're talking about just a few minutes ago even with the tooling with Visual Studio and DataDude and things like that. At the end of the day, a developer says "I just want to make this application work."

Richard Campbell: Yeah.

Kevin Kline: That's what I want and I don't want to have to learn a new language. If I don't know SQL, I don't want to have to bury my head in that and learn how to do a good SQL statement. So what we're seeing is there's quite a few -- you know, it's a Greenfield right now so there are a lot of new things popping up but there are a few thought leaders. Hey Dupe is one. There are a couple of other No SQL

databases out there and they're attracting a lot of attention because they are a means by which you can access your data. Typically, these are especially good with unstructured kinds of data. You know, things like maybe XML documents or PDF or whatever you might want to throw into them, and they're blazingly fast and they're non relational so that means they don't require a normalize set of tables or primary keys or anything, any of those sorts of things.

Richard Campbell: Right.

Kevin Kline: But really I must be honest, they're an obstacle for a developer to make their application run. One of the things that it is interesting when you look at this non-SQL database is they use a kind of technology called Column Storage. So traditionally with the Relational Database, they use what they call Table Storage. We do things in rows, we define a table.

Richard Campbell: Right.

Kevin Kline: And what is happening actually with the R&D efforts at places like Microsoft, and Oracle, and IBM with their DB2 tool is they're actually working on Column Store technologies themselves so they want to be able to get to that position where they have this blindingly, blazingly fast column look-ups for this kind of unstructured data as well. It was something that one of the keynote speakers from Microsoft spoke at PASS last year. His name is David DeWitt, and so the No SQL movement I think when I first heard about that, I thought "Oh, this is just kind of a group of angry developers that are like no matter what you're selling, we're not going to take it anymore." But now that I start to think about it and kind of analyze what's happening behind the scenes, I think there's going to be some legs to this, there's going to be some longevity because it does deliver on what it promises which is really fast and it just works, you don't have to learn a lot of new stuff and it's the kind of thing that the other major database platforms are starting to pay heed to. Kind of logically, the really great thing about No SQL databases is that they're easy to construct in a way that's memory resident based in memory.

Richard Campbell: Right.

Kevin Kline: And much more so than a relational database. What we're looking at today with the current state of the hardware is the biggest bottleneck is going to disk.

Richard Campbell: Yeah.

Kevin Kline: So even if we add flash drives, that's still our biggest bottleneck. So if we can get more and more of the database to work in memory, we can achieve enormous scales of improvement and



that's why it's interesting to Microsoft and Oracle and DB2.

Richard Campbell: But we're also risking the data. You know, I remember before relational databases -- I'm dating myself now, when we were using hierarchical databases and network data store and stuff where we had to walk the tree ourselves and the relational data that we were really focused on is reliability. I remember doing the test where you wrote a record and said the record is written, you pull the plug out of the machine then you power it back up and it would show that yes, when I say the data is written, it's bloody written. It's absolutely reliable and these things are doing that.

Kevin Kline: That's right and that's the key to any of these new database systems. They have to be cost effective. They have to be reliable and they have to be secure. In other words, that they keep the wrong people out and they left the right people in.

Richard Campbell: Right.

Kevin Kline: That's the challenge both for the No SQL databases. That's also the challenge for any of the Cloud database vendors too. You know, some of these different databases in the Cloud claim to be safe and reliable and secure and whatnot, but they're not transactionally consistent.

Richard Campbell: Right.

Kevin Kline: You might lose a transaction here and there.

Richard Campbell: Yeah.

Kevin Kline: As well as we have famously seen major outages of data because it wasn't backed up properly.

Richard Campbell: Uh-hmm.

Kevin Kline: Some of those same challenges exist on a No SQL side. It's still very amateurish.

Greg Hughes: Sure.

Richard Campbell: And as soon as you're working across the internet, you've got a reliability problem.

Kevin Kline: Absolutely, yes.

Richard Campbell: You know, there's just a basic weakness there but I also think that -- I do a lot of work on high performance systems, we routinely sacrifice reliability for performance. If you stick a caching layer between your application and your database, you've just broken your reliability of data writing. You're already doing that. I'm realizing now,

Kevin Kline on the State of SQL Server and More!

January 13, 2010

as we start looking at the No SQL movement, that what they're talking about is stuff we're already doing. We just call it something else. We call it cache.

Kevin Kline: Right, that's right.

Richard Campbell: You know, we've already got a transactional integrity problem. We've just been ignoring -- we're willing to do that for that performance improvement.

Kevin Kline: Right.

Greg Hughes: Once you have the data in memory, then work with the data in memory. There's no reason to keep going back to the disk.

Richard Campbell: Right. It's certainly not synchronously anyway.

Kevin Kline: Uh-hmm.

Richard Campbell: It's an interesting set of problem and I think from an IT perspective trying not to just be the guy who always says no, there is the "Hey, you know we have a service level of agreement here that says these transactions are reliable. Are we still working with that? Are we compliant to this if we go this way?" Here's another side to this that I think is starting to surface as well but actually I think a little bit older than the No SQL movement. It's really the whole ORM mindset, the Object Relational Mapping mindset the developers are getting into, or some of them at least are getting into today where they're not building the database first. They're building the object first and saying "Just store this. Don't tell me how you did it, just keep it. When I ask for it back, give it back." And so you're literally having databases being generated off of the back of the object grasp of an application.

Kevin Kline: In the long run I think that's probably the way to go. Our computers and processors and things like that are getting so powerful now that it's the kind of situation where it seems like it would be possible that we could let the machines do more of that.

Richard Campbell: Yeah.

Kevin Kline: And in that way we could spend more of our time trying to solve the business problems which I think is something that sometimes gets lost in the mix of all of this and we spend our time and our brain power thinking about elegant and beautiful solutions.

Richard Campbell: Yeah.

Kevin Kline: You know, there's nothing that drives me crazier than to see two DBA type of guys arguing about whether they're data model at third



normal form is the best one compared to this other guy's data model for the same database in second normal form or fourth normal form.

Richard Campbell: Yeah. Well, and you're right. We got that pounded into us in the '90s and we just can't seem to let it go. So this idea that here's this database that's largely generated from a tool and immediately the DBA may start screaming, then I think I'm not going to report directly to that data. I've got an ELT load to a reporting database anyway so what do I really care what the structure of the incoming data is. It's just that super faster the app, it's totally self managing, I don't need to put any of my resources into creating anything for it, it just happens so I'm focused now on the transformation process that pulls that into a form that reports well.

Kevin Kline: Yeah, that's a beautiful thing.

Richard Campbell: It's a different way of thinking and it goes against the grain. The hair on the back of my neck is literally standing up but I can't find anything really wrong with it.

Kevin Kline: Right, right. You know, I think the old tradition is that we control and we know every step of the process.

Richard Campbell: Well, because we're being held accountable for it.

Kevin Kline: Exactly. But at least from the DBA side, the thing that's funny to me is typically they would not, a DBA would typically not see ground to an application that they knew was generated by -- automatically generated by a system.

Richard Campbell: They'd be horrified.

Kevin Kline: But they would happily go ahead and install a third party application that was written by a company that has no idea what a join is.

Richard Campbell: Right.

Kevin Kline: And I see that all the time in our customers as well. You know, you go to their shop "Oh, we're having a problem with performance on this particular application," and we get into the code and we see that it was really written for Oracle but they're running it on SQL Server.

Richard Campbell: Right.

Kevin Kline: So it takes advantage of Oracle's processors and preferred algorithms and think things like what. So in a sense we see many, many customers who do that, but just like you're saying, you know, the hair on the back of your neck would stand up if you suggested some alternative that really is the same thing.

Richard Campbell: Yes.

Kevin Kline: You know, you see that in that control and that performance but it just goes against the grain for them.

Richard Campbell: Yeah. What I see the same thing going back to that No SQL movement is this whole "Guys this isn't that new." We've already been violating acid rules within this application using caching layers. So just because we're -- and just because we're admitting it doesn't mean it isn't already happening.

Kevin Kline: Oh, and so very true, so very true and developers have been looking for short cuts for a long time. I mean, one of the most common questions I get when I speak at events like TechEd or one of the big conferences, people come around and they'll say "Hey, what I really want to know is what is the syntax for slash faster?"

Richard Campbell: Where's the thing, the red button?

Kevin Kline: Yeah. Where is the little switch I can put, and it says just make SQL Server go faster and stop screwing around with me. So they're always looking for that shortcut and like you said also, none of this is new in the sense of the theoretic set that these ideas have been around since the '70s.

Richard Campbell: Yeah.

Kevin Kline: Probably even before then, you know. Everybody who's new to the field thinks that virtualization is the greatest newest thing, but we saw that on old IBM mainframes at the MVS decades ago.

Richard Campbell: Yeah, that's true and I also really have a sense and I guess it have just been in this long enough that we've swung too far over to the relational database. We store everything in relational database even stuff that isn't related.

Greg Hughes: Even if we don't need to.

Kevin Kline: Yes, exactly. Just like you said, you know, certain things were padded into our heads in the '90s and this is one of the things about IT that is a bit of an annoyance to me. It's that it's very kind of fad-driven, you know. What is the greatest, newest technology? Oh, let's all -- I mean we all go there in a herd, we all make the same mistakes together and try to figure it out, and then 10 years, 8 years later something else comes along and everybody runs over to that and pick it up. So definitely that's something that I've seen as well.

Greg Hughes: Fad-driven and you know, an IT has a tendency to watch what everybody else is



doing and to say, well, if everybody also is doing it then that must be what I need to do so I'll do that, as oppose to getting in and really understanding what is it that you're deploying and what is it that you're leveraging in order to get the job done.

Kevin Kline: A great example of that for me is outsourcing. You know, so many people jump into outsourcing because that's what everybody else was doing and then they wind up saying "Gosh, this is a disaster," and they pick it up and try to rebuild their team here in the States and so forth and if they had thought about the whole situation from not what everybody else is doing and then I better do what they're doing, but what does the business need and what are the best options for our organization. In a sense, my joke is I like to blame the airline travel magazines because the CIO of the company happens to be sitting in a seat on an airplane flight and he got a little bored and he flipped open whatever the American airlines or the Delta Magazine is and there's an article in there about some particular management fad or some kind of technology fad and there are CIOs all over the country right now who are beating up their staff saying "We got to get on this Cloud thing." They've got no idea why.

Greg Hughes: Yeah.

Kevin Kline: It's just that was in the magazine when they were on the airplane.

Greg Hughes: Right, exactly. What's the one message that you want people who are thinking SQL Server, whether they're new to it or it's old hat, that you wanted them to have right now. Maybe we can use that just to tie up the show. We're pretty much out of time.

Kevin Kline: If I had one message to someone who is new to databases, maybe they had a successful start-up and they're "What are we going to build our next generation platform on? What are we going to build our first generation platform on?" My message to them about SQL Server would be that SQL Server has really arrived in terms of both performance and manageability, and with 2008 R2 we now have the opportunity to get that information to the right people at the right time and the right format in a way that we've never been able to do in the past. So it's very, very powerful and yet still very usable for everyone throughout the company.

Richard Campbell: Kevin Kline, thanks so much for coming on the show.

Greg Hughes: Thanks Kevin.

Kevin Kline: Gentlemen, thank you so much. I appreciate it very much.