

[ MUSIC ]

BRUNEL: Welcome to another edition of, Getting the Most out of IBM U2. I'm Kenny Brunel, and I'm your host for today's episode. And today we're going to talk about EDA, or External Database Access. With me in the studio today, I have a couple of guests. I also have a guest calling in from Dallas, Texas. First of all, Helen Beylkin is the development manager for U2. Helen, thanks for coming in today.

BEYLKIN: Thank you for having me, Kenny.

BRUNEL: My next guest in the studio is Dave Peters. Dave is the Project Manager for the IBM U2 data servers, and is going to provide information about EDA.

PETERS: Thank you, Kenny. It's great to be here.

BRUNEL: And finally, dialing in from Dallas, Texas, is Ray Else. Ray is an EDA specialist and has real-life experience with implementing EDA in a production environment. Ray, thanks for dialing in.

ELSE: Hi Kenny. I'm glad I could join the call.

BRUNEL: So first of all, what is EDA? Helen, perhaps you could tell our listeners about EDA: what is it, why might they want it, and what's the best application for one's needs?

BEYLKIN: Sure, Kenny. External Database Access -- or, EDA -- is a technology that enables UniData and UniVerse applications to transparently store their data in an external database such as IBM DB2 or Microsoft SQL Server. And, subsequently access this data without even the users' or application developers' awareness of the fact that the data actually resides in an external database.

Due to EDA, whole power flexibility of the UniVerse, and UniData development environment is now coupled with database independence which is a very important vehicle for integrating U2 applications with other applications and tools in the enterprise environment.

Even prior to EDA, there were multiple ways to replicate U2 data into another database or access a relational table from a basic application. Nevertheless, the unique feature of EDA is an absolutely transparent access to data in external database tables requiring no application rewrite or requiring no knowledge of the foreign database access language such as SQL.

So EDA gives you the freedom to store either some or all of your U2 data in a variety of external databases, while at the same time allowing you to continue using the familiar U2 application development environment.

BRUNEL: Thank you, Helen. Ray, as someone with experience using EDA, perhaps you can give us some examples of why, and how someone might use EDA.

ELSE: Well, as Helen explained, EDA is really slick, sophisticated technology. It can not only convert data from UniVerse and UniData to other databases like DB2, but it can also convert on the fly, U2 reads, rights and select statements into their equivalent SQL statement.

Let me give you a couple scenarios where EDA could come in handy. Let's say you have a UniVerse inventory database, and a Web-based DB2 shopping cart app. EDA could be used to replicate the inventory file data which is in UniData into a remote DB2 database table, so the Web app knows up to the minute inventory figures. EDA could also be used to populate a datamart or data warehouse for BI reporting.

One example is how we've used EDA internally and IBM U2. We actually use it every night to pull data from an existing DB2 view using UniData commands to reconcile SAP payments against our UniData open invoices file. We went live on that application about a year ago. It's been very reliable and we've been very happy with EDA used in that fashion.

Lastly, there's the possibility of using EDA to export a

chunk or maybe even the majority of your data to an external database like DB2 and then continue to run your U2 application against that exported data. This last option is obviously the most challenging.

BRUNEL: Thanks, Ray. Dave, you mentioned that DB2 is a database that's supported by EDA. What other external databases are supported, if any?

PETERS: Well, Kenny, EDA was introduced in UniData 7.1, and at that time, we only support it through the built-in driver that was included with the product. With UniData at 7.2, we introduced an API so that you could write your own interface or driver to any SQL database that you like.

So actually, the answer to your question is really it's unlimited. You can write your own driver to any database that you'd like. We're going to also introduce this feature including the API in UniVerse 11.1.

BRUNEL: Dave, you mentioned that this would be available in UniVerse 11.1; when is that scheduled to release?

PETERS: 11.1 will release late 2009.

BEYLKIN: And Dave, let me clarify for those listeners

who are interested in writing their own EDA driver for other relational databases, using this open EDA driver API. The driver API consists of 16 C language routines that generally resemble our DBC functions such as connect, disconnect, prepare statement, execute statement, face statement and so forth.

BRUNEL: Ray, are there other U2 applications that are required in order for EDA to work; or is it something that you can do right with the data server?

ELSE: You'll want of course the right release of the database server that supports EDA. And then, for example, to access DB2, you would also need DB2 client installed on that same machine. That's a free tool that comes from IBM, and it gives you some great functionality.

For example, with our U2 Business Connect application, we're using EDA to talk through DB2 client using DB2's clients encryption, or transport encryption functionality, so that the data coming from DB2 is in fact encrypted over the lines coming into our application.

BRUNEL: Okay, thank you, Ray. So, now I'm going to really focus on how we get the most out of EDA. And Dave, I'll start with you.

PETERS: Besides your solid working knowledge of UniData or UniVerse, you really need to know the target database or the external database that you want to use. So the key is to really understand that database.

You might also want to take advantage of some of the support services that you can get from either your IT specialists, or your sales rep. They can point you in the direction to get even more resources to help in your project.

BEYLKIN: Right, Dave. And also, as Ray said, EDA not only converts data stored in U2 databases to external database tables, but performs on the fly run-time translation of U2 reads, rights and query statements into external database SQL commands in order to directly pull data from or push data to the external database.

In order to enable this on the fly translation, a metadata mapping between the UniData, UniVerse dictionaries and external database schema is performed.

BRUNEL: Helen, you mentioned EDA enables run-time translation. How is that possible?

BEYLKIN: Yes, Kenny. This translation or mapping is aided with a sophisticated Eclipse-based tool called EDA Schema Manager that is available on your client CD. And the

mapping consists of two major phases. One of them is defining first normal form tables in their primary and foreign key relationships in order to express a pretty complex attribute relationships within U2 files between the single value, multivalued, sub value attributes and their associations.

And the second part of this mapping actually enables converting data attributes, virtual attributes and virtual attributes indexes to their relational equivalents. Once this mapping is defined, EDA transforms uni query and retrieve statements as well as the reads and rights, into the first normal form SQL language. It also continues to support U2 locking and manage transactional process and semantics.

PETERS: And also you don't need to map all of your data to get started with EDA.

BEYLKIN: True, Dave is right. Probably the best mapping feature is its flexibility, because it can translate a map, as many U2 files as you need, from just one file to the whole account. And for each mapped file, have as many or as few attributes explicitly translated to external table columns as you need.

From just [at a z] one attribute, to as many as you wish,

or, more precisely, as many as you need to be exposed to an external database applications and tools. We call it actual attribute visibility. You may have hundreds of attributes in your files, U2 files, but only need to expose a few of them for reporting or other manipulation by the external database applications or tools.

Although all the data in the converted file now resides in the external database, it is more efficient and is even more safe to only explicitly map those attributes that need to be visible to external database applications. And, keep the rest of them in the native and compact multi value and sub value format. So as you can see, some learning and planning is required for a successful and efficient use of EDA.

BRUNEL:        So that sounds like the EDA driver would be a pretty intricate or tightly integrated add-on for the database. So anytime I upgrade something with any application, I'm concerned about the effect it's going to have on performance. So, Ray can you speak to that a bit, as to, should our users be concerned about performance degradation if they are to utilize EDA?

ELSE:         Well, you're right Kenny, that performance is a key consideration for all application developers. From my work with EDA, what I've seen is that the writing performance is typically reasonably good. And if you're

going to read from existing DB2 tables or views that were previously set up, and that performance is good in DB2, then the performance will be good from UniVerse and UniData as well.

Where we have seen performance issues with EDA is typically when people have exported very complex U2 data records that have many multi values and sub values, and perhaps they haven't even bothered to logically associate those fields.

So that what happens over on the DB2 side is to get a lot of child tables, that then have to be joined when you go to read that data, causing performance issues. We've actually addressed this to some degree, adding some new functionality to EDA, but there's trade-offs.

So generally speaking, I can say that you really have to look at, what is it that you're hoping to accomplish with EDA? How were your files laid out? Have you taken the time to maybe create some associations that aren't there? Just, it has to be handled on a case-by-case basis, and we're certainly willing to assist with that analysis.

BRUNEL: Dave, Ray just went over some of the items in regards to implementing EDA. So, we know that it's going to be available, or that it is available currently with UniData 7.2, and will be coming out for UniVerse 11.1 in the future.

For those listeners who are ready to start looking into EDA, what information is available for them out there?

PETERS: Well, Kenny, a good place to start is with the documentation. You can find the EDA manual, which is a separate manual, in the U2 library on the IBM Web site. There's also a good developerWorks article to read called, External Database Access with U2. It'll help you get started a little bit quicker after reading.

Then there's some services that you might want to take advantage of. Lab services can provide you with experienced personnel that will help you get your EDA application implemented faster and would utilize their experts to do so.

Then you also might want to take advantage of premium support. It's a proactive service offering rather than the standard reactive support offering. And through premium support you can get improved implementation times and utilize the best practices of those experts in the support department.

BRUNEL: Thanks, Dave. So it seems like there's a fair amount of resources available for EDA, but how does one go about getting EDA?

PETERS: Well, Kenny, EDA is included with the

Enterprise Edition, but is a chargeable feature with the Workgroup and Server Editions.

BRUNEL: That wraps up today's podcast, and I would like to thank each of my guests for joining me in the studio today. Helen, thank you.

BEYLKIN: Sure, Kenny, and thank you for giving me the opportunity to join this podcast.

BRUNEL: Dave Peters, thanks for joining me today.

PETERS: Thank you, Kenny.

BRUNEL: Ray Else, thanks for dialing in from Dallas, Texas, today.

ELSE: You're welcome, Kenny. I enjoyed it.

BRUNEL: And I would like to thank our listeners for tuning in today, and I hope you found today's podcast useful. Be sure to e-mail us with your feedback, any comments you may have and suggestions for future episodes. Our e-mail address is [U2askus@us.ibm.com](mailto:U2askus@us.ibm.com). You can also click on the feedback link just below this podcast download on the Web page.

Transcripts of today's podcast, which include all relevant links, are available on our Web site by following the U2 podcast link from the main page. Be sure to stay tuned in

for future episodes of how to get the most out of IBM U2.

[END OF SEGMENT]