

BriefingsDirect Insights Analysts Examine HP-Oracle Exadata Release, Extreme BI, Virtualization and Cloud Computing Trends

Edited transcript of BriefingsDirect Analyst Insights Edition podcast, Vol. 30, on Exadata, extreme BI and cloud computing, recorded Sept. 26, 2008 from Oracle OpenWorld in San Francisco.

Listen to the podcast. Download the podcast. Find it on iTunes/iPod. Learn more. Sponsors: [Active Endpoints](#), Hewlett-Packard.

Dana Gardner: Hello, and welcome to the latest BriefingsDirect Analyst Insights Edition, Vol. 30. This periodic discussion and dissection of IT infrastructure related news and events, with a panel of industry analysts and guests, comes to you with the help of our sponsors, charter sponsor Active Endpoints, maker of the ActiveVOS visual orchestration system, and also Hewlett-Packard via the HP Live! Podcast Series.

I'm your Host and moderator Dana Gardner, principal analyst at Interarbor Solutions. Our panel this week consists of Joe McKendrick, an independent analyst and prolific blogger on SOA and BI topics. Hi, Joe!

Joe McKendrick: Hi, Dana, glad to be here.

Gardner: We are also joined by Brad Shimmin, a principal analyst of Current Analysis. Hey, Brad!

Brad Shimmin: Hi, Dana, thanks for having me.

Gardner: Jim Kobielus joins us. He's a senior analyst at Forrester Research. Hello, Jim!

Jim Kobielus: Hey, Dana, and Hi, everybody!

Gardner: And Dave Linthicum, independent consultant, joins us this week. Thanks for coming, Dave.

Dave Linthicum: Thanks, Dana, thanks for having me back.

Gardner: We are going to be talking about the news of the week of Sept. 22, 2008. We'll be looking at the HP-Oracle announcements and other news made here at Oracle OpenWorld. We'll be talking about cloud computing and the notion of "on-premises" or "private clouds," and how data portability might actually work among and between different clouds -- both "public," if you will, and "private." We'll look at recent virtualization news from VMware, HP, Red Hat and Citrix.

An Exadata 'Shocker'

Let's start our show this week with Jim Kobielus. Jim, you and I are both here at the Oracle OpenWorld. We had an unusual announcement around optimization between hardware and software from Oracle, which has traditionally been a software-only company.

Oracle and HP introduced two Exadata products. I wonder if you could fill in our audience on what Oracle did this week.

Kobielus: Yes, this week Oracle announced the release, in partnership with HP, of a very high-end data warehousing appliance. They may not use the word "appliance," but that's in fact what it is. It's a configured and optimized bundle of hardware and software, with database storage, so it meets very high-end data warehousing requirements. It's called the HP Oracle Database Machine.

It encompasses and includes the HP Oracle Exadata Storage Server, which is a grid storage level server. One of the key differentiators on that storage approach is that it puts the query processing in the storage subsystem. As a result, it can greatly speed up the processing of very complex analytics. What Oracle and HP have essentially done is take a page from the Netezza book, because that is, of course, the feature of the Netezza performance system. [Netezza] has an appliance that they accelerate your data models through by putting this whole processing back close to storage. But

the HP Oracle release does much more than simply taking the page out of that Netezza book.

What they did essentially is they also shot across Teradata's bow, because this is Oracle's petabytes-scale, data warehousing-solution platform. The HP Oracle Database Machine that they demonstrated at the show definitely screams. And it can scale -- Oracle says that it can scale with almost no limits, and that remains to be seen.

But it can definitely go to a much higher scale in terms of capacity than the current Oracle Optimized Warehouses that they have already begun to ship with HP and a variety of other hardware partners. The Oracle Optimized Warehouses max out at several hundred terabytes. And, as I said, the HP Oracle Database Machine can go well beyond that.

This is a shot both across Teradata's bow and across Netezza's. And Oracle Chairman and CEO Larry Ellison from the stage directly honed in on both of those competitors by name.

It was classic Ellison, a very well put together presentation. But quite frankly when you begin to analyze the various claims he made, they don't all hold up. Or rather, he is presenting a lot of the Oracle-specific differentiation. Yet they were very impressive.

Gardner: This is a significant departure for Oracle and HP on several different levels. On one hand, we have a combined hardware-software product from two different vendors. We also have a new parallelization process, with their architectural design, where the database and the storage are very close. The processing can take advantage of massively parallel processing. We also have the fat pipes in the form of InfiniBand connections.

So we have an architectural departure. We have a hardware-software departure, and we also have this interesting alliance between HP and Oracle, making and selling a product together. How does this all strike you, Brad Shimmin?

Shimmin: Well, I was shocked, shocked, absolutely simply shocked. This is because historically Oracle has strayed so far away from the

appliance market. It's been surprising to me on a number of occasions when I have spoken with them about acquisitions. Actually they made one acquisition earlier this year, they acquired a company that had an appliance that was very successful, and they chose to simply kill it outright because they "did not want to play in the space."

With that said, I am glad to see this happening. I really am, and I think that HP is a good partner with them because I don't feel that the two companies really bumped into one another in terms of Oracle's core constituency. So I think it's a good play all around, and I am glad to see Oracle finally getting into this. Now if only they would release parts of their middleware as an appliance, I would be very happy.

Kobielus: And, in fact, Brad, Larry Ellison indicated that they seem to have some plans for that. They really resisted details -- but they seem to have some plans to "appliance-ize," if that's the word, more and more the Oracle Fusion Middleware stack.

Shimmin: There are other quite prominent middleware stack players that are moving to appliances, as well. I can't mention their names but the use of appliances seems to be of great interest to more vendors.

Gardner: I have also been picking up on this interest in the appliance business. IBM has been into this with DataPower SOA Appliances for a while now, but IBM has not really extended use of appliances out as widely as I was expecting. I have also heard that TIBCO may be building an appliance for complex event processing (CEP). So, yes, I think we are going to see more of this.

Brad, I want to go back for one second to the HP-Oracle relationship. It almost seems now that Oracle has anointed HP at some level as a preferred hardware supplier on storage, if not also other aspects of hardware. What does that mean for EMC and some of the other storage hardware providers? They are no longer on an independent or third-party-friendly level with Oracle, right?

Shimmin: Absolutely. I think that all of those relationships will come under strain from this. There is no question about that. And it seems to me that this makes Oracle look a lot more like both IBM and Sun

Microsystems and EMC, in terms of having some sort of competency in hardware. So I think there are going to be a lot of far-ranging ripples from this relationship that will change the way the market functions.

Gardner: And how is all of this going to come to market? If you want to buy the Exadata warehouse you actually have to go through the Oracle sales force. Oracle is going to support it, and sell it, price it, and then HP is going to service the hardware. So in essence, HP is the supplier to Oracle, and Oracle is the principal vendor. Does that mean anything anybody out there?

Kobielus: It means that Oracle is taking much more of a marketing lead on the HP Oracle Database Machine than they have with any of the Oracle Optimized Warehouses. So Oracle is very much staking its data warehousing go-to-market strategy on this new product, and on this partnership with HP. That said, HP is providing all of the technical support on the new products. So it's not like Oracle is really becoming a hardware vendor, rather they are going to become very much a software vendor, but has staked their future on delivering the software on this one particular hardware partner's platform.

Gardner: Actually it allows Oracle to operate at a solutions level and so take quite a bit more of the margin across that total data warehouse solution, right? And that undercuts the data array providers significantly.

Okay, so let's talk about what we do this thing. We heard that the 1 terabyte-sized data sets and higher start to hit performance issues. And that then prevents companies from adding more queries on their warehouses, and also reduces the amount of additional data that they want to put into their warehouses.

So we could have hit a wall somewhere around 1 terabyte databases. This approach, this architecture in the Exadata hardware-software optimization claims to blow that away, that it can deal with the largest sets, of 10 terabytes and up, with very high performance. What does this mean for business intelligence (BI) analytics? What does this mean for bringing more types of data and content into the warehouse? What are the business outcome benefits?

Joe McKendrick, what are your thoughts on the BI perspective on this market development?

McKendrick: Well, it certainly moves the business intelligence arena forward. Looking at what the rationale is for having an appliance in this market -- versus what's has been happening for the previous decade with data warehousing -- it really says a lot about what's needed in the market. Data warehouses, when you get into the multiple terabyte range, are simply too complex and have high cost of ownership. That's made BI a fairly expensive proposition for companies going this route, and the cost is tied into the maintenance, the updates for the warehouse software, the organizational effort, and the input required to make a large data warehouse go.

Now there is a trend emerging. I am sure Oracle has an eye on this as well. It's toward open source. We are seeing more open source in data warehouses too. This is open source at the warehouse level itself, at the database level itself. [Sun Microsystem's] MySQL for example has been pointing in this direction, PostgreSQL as well. [And there's Ingress.]

Gardner: Well, that's another distinct issue. Now with Oracle and HP cooperating, why shouldn't we expect Sun to come out with something quite similar, but with MySQL as the database, and their [Sparc/UltraSparc] processing, and their rack, cooling and InfiniBand, and of course, their storage?

McKendrick: I wouldn't be surprised, I wouldn't be surprised one bit if we see some kind of response from Sun fairly soon because Sun still makes its money from hardware.

Gardner: Right, now if Sun does that then IBM will certainly come out with something around DB2. We should expect that, right?

McKendrick: Yes, yes, definitely. And I think there is an emphasis on simplifying data warehousing, making data warehousing simple for the masses. Microsoft, love them or hate them, has been doing a lot of work in this area by increasing the simplicity of its data warehouse and making it available at more of a commodity level for the small to medium size business space.

I think we're going to see more in the open source data warehousing space, and Oracle is looking at that as well.

Gardner: Let's go to Jim Kobiellus. Jim, [using Exadata] we can start taking 10 terabyte data sets and delivering analytics in near real-time and deliver query results out to various business applications on huge scale. We can also start looking at this as cloud infrastructure -- where we are going to be providing data as a service, BI as a service even. And then we have Sun, IBM, and perhaps Hitachi, and all these other guys that are jumping in with their own data warehouse appliances, and they start beating each other up on price, and the price comes down in the market. Are we then entering an era of affordable extreme BI?

Kobiellus: For sure. Well, extreme BI, that's really BI and data warehousing with very large data sets, with very demanding real-time loading scenarios, with very extensive concurrent usage and so forth. We are already in that era. If you look at what's going on -- and actual deployment, enterprise deployments -- like 10 terabytes, those are in fact much of the data warehousing solutions in the market. That's the joy of data warehousing, enterprise departments are between 5 and 15 terabytes. They are being handled quite well through a lot of symmetrical multi-processing (SMP). So these are around in the market.

Now our data warehousing and BI environments are in the hundreds of terabytes, and up to the petabytes range and beyond. A lot of these are in the cloud already. I can't name names yet. There are a few things right now that I can't show. But well-known Web 2.0 service providers are already above the petabytes scale in terms of the amount of data that's persisted, and in terms of their needs and their ability to do continuous concurrent loads into those humongous data warehouses in real-time. In these extreme data warehousing environments you may have millions upon millions of queries hitting that data warehouse all the time.

Gardner: But aren't we with Exadata taking this from the high-end, roll-your-own, computer-science gee-whiz level down to much more of an off-the-shelf, forklift upgrade level? Aren't we now getting to

extreme BI at much more a commodity, or at least something that's much more germane across many more types organization?

Kobielus: Oh, yeah, for sure. It's getting down into the affordable way to eventually bringing cloud data warehouses down into the range of the main market, as well as for large enterprises. ... The one thing -- one of the other important outcomes from my point of view this week at Oracle OpenWorld -- was the fact that Oracle, now in conjunction with Amazon's Elastic Compute Cloud, has an Oracle cloud -- the existing Amazon cloud can take Oracle database licenses.

They can move those licenses to a cloud, hosted by Amazon EC2. Using tools that Oracle is providing they can move their data to back it up or move databases entirely to be persistent in the cloud, in Amazon's S3 service. So this is very much a lead in. I strongly expect that the other enterprise database vendors over time, maybe in a year or two, we'll also offer similar deployments and flexibility for their data warehousing customers.

Gardner: Okay, let's go to Dave Linthicum on that. Dave, you're familiar with moving data around the cloud. It sounds like people will start getting comfortable with this from a risk and from a reliability/privacy/control issue level. And then it's a no-brainer to start moving fairly massive data sets, or for backups or extend-enterprise sharing or federation of data -- what have you, into cloud infrastructures.

How important from your perspective is what Oracle announced in conjunction with Amazon this week?

Linthicum: I think it's very important. I think that the economics -- that it's much cheaper to do cloud computing than on-premises stuff -- and you can prove that at each and every time, or run into an issue around cultural, and kind of total protection issues within the enterprises ... I think those are falling down as time goes on. Go back in a time machine five years ago, and start talking about running major enterprise applications delivered as SaaS, they would have laughed at you.

Today, everyone is using Salesforce.com, and just a bunch of other

SaaS delivered applications. So enterprises are getting their minds around cloud computing, understanding the concept of it. So moving information into the cloud is not really much of a leap. You already have customer information existing on Salesforce.com, or other SaaS providers out there.

I think that this is one step in the direction that, in essence, we're going back in time a bit, moving back into the time-share space. A lot of things are going to be pushed back out into the universe through economy's scale, and also through the value of communities. It just can be a much more cheap and cost effective way of doing it. I think it's going to be a huge push in the next two years.

Gardner: Does it seem reasonable that Oracle would test the waters on this, in terms of market acceptance, with Amazon? Once people get a little more familiar and comfortable with it, then Oracle comes out with its own cloud offerings?

Linthicum: Absolutely, I think that Oracle is going to have a cloud offering, IBM is going to have a cloud offering, Sun is going to have a cloud offering, and it's going to be the big talk in the big industry over the next two or three years. I think they are just going to get out there and fight it out.

I think you are going to have number of startups, too. They are going to have huge cloud offerings as well. They are going to compete with the big guys. And they can -- because it's very simple to put up infrastructure. It's fairly cost effective, and you can get out there and start battling it out with them. Quite frankly, I think, maybe the more agile, smaller companies may win that war.

Gardner: In other recent news, Brad Shimmin, we have heard quite a bit of virtualization, and cloud compute discussions from VMware, from Citrix, and from HP. We saw some acquisition from Red Hat that brings them into the hypervisor space. Maybe you can help our listeners understand a little bit better the relationship between virtualization, management and platform vendors, and how this whole notion of private or enterprise clouds works.

Shimmin: It depends on the perspective we have, right? It depends

on if we are talking about virtualizing the datacenter, virtualizing the desktop, or moving facets of the datacenter to the cloud. If you are trying to understand how, as we were just talking about, these smaller players are able to use things like Amazon EC2 to get into the market -- or if we are talking about moving the desktop to data center cloud -- what I want to understand as a customer is just what the SLAs and protections are from these providers, whether it's IBM or Amazon. And, by the way, another one we need to mention is Cisco, which will be using the WebEx platform as a SaaS platform and SaaS solution for the enterprise.

The point is that as a customer you don't just want to know what the [performance reliability figures] are, you want to know what sort of wrapper these vendors are putting around their solutions for things like security and policy management enforcements. It's not just the fact that they will be able to secure the data, but it's about being able to control and manage the data, and have visibility into the data; whether it's something that's sitting in some sort of virtualized instance in your own datacenter, or whether it's something that's sitting in some federated system that might be shared between Cisco and Amazon.

Gardner: I found it interesting that these vendors are basically tripping over themselves and rushing out to the market, way before these private clouds have even established themselves. Yet the vendors are declaring that they have the infrastructure and the approach to do it. It sort of reminds me of a platform, or even operating system, land grab -- that getting there first and establishing some of the effective standards and coming up with industry-common implementations gives them an opportunity to at some level or format create the de facto portability means.

This is a layer above virtualization. And virtualization is there to bring all the legacy stuff into play, but what do you do with the new applications? What do you do with the new services? Dave Linthicum, what are your thoughts on a meta-operating system in the cloud? Are we in a kind of a race to be first to that?

Linthicum: I think that's a ways off at this point. I think people are going to put aspects of the infrastructure up in the cloud first. And I

think that the platform-as-a-service (PaaS) and the ability to provide a development infrastructure, storage infrastructure, some deployment infrastructure, and things like that -- that is all going to be a bit of mix and match. I think people are going to do little tactical projects to kind of dip their toe in the water to see if it is viable.

However as we go forward, I think that's the destination. If you look at how everything is going, I think everything is going to be pushed up into the cloud. People are basically going to have virtual platforms in the cloud, and that's how they are going to drive it. Just from a cost standpoint, everything we just discussed, the advantages are going to be for those who get there first.

I think that very much like the early adopters of the Web, back in the 1990s, this is going to be the same kind of a land grab, and the same kind of land rush that's going to occur. Ultimately you are going to find 60 percent to 80 percent of the business processes over the next 10 years are going to be outsourced.

Gardner: What about this issue of data, these massive data sets, and bringing some of that up into a cloud? Is it going to be just standards-based interoperability for my data set and your data set to play well with each other? To what level are we hung up by different cloud implementations, and therefore perhaps also different data implementations? Does that need to be solved?

Linthicum: Yes, I think it does. I think that you are going to find that integration does occur in the clouds, just like it does within the enterprise, from enterprise to enterprise. The reality is that people have information up there with different semantics, different data formats, and all of that stuff has to be transferred one to another.

I think that the idea of integration in the cloud, which I have been involved with personally over the last 10 years, may actually start to be used. I think that people are going to have to do transformation around control, filtering, all of these things as information moves between these partitions out in the universe. Ultimately integration is going to be easier. We know lot more than we did 15 years ago when I wrote the book, Enterprise Application Integration. But I think that it's still going to be needed and a necessary thing. So maybe integration

in the cloud companies should start pushing forward.

Kobielus: I hear what you're saying. I think that's an important point to put forward, which is that these clouds, these data warehousing clouds -- data warehouses that are external to the firewalls, the multi-tenant environments -- are multi-domain, multi-entity data warehouses with strict separations between the various domains, which are often searching with particular customers. But like a supply chain application in the cloud, it is the probably the best place to put all that data so that companies and suppliers and the customers all have access to common pooled data in a common externally hosted environment.

What that raises then is that the data warehouses in the cloud, really become data federation in the cloud. So all these different data sets, the divergent schemas and so forth, need to be normalized to a common semantic layer in the cloud provided by that cloud vendor. So then you are into the data federation vendors that had a huge footprint in the enterprise, those guys then need to provide their capabilities in the cloud for these types of supply chain and B2B applications.

I am talking with a couple companies, like Composite Software and some others, where they have well established data federation to manage virtualization layers. Those guys need to get cracking to put a lot of that into a cloud environment to enable this level of data integration and federation in that cloud environment going forward, and make it scalable.

Gardner: Well I think we will have to leave it there. We have been discussing announcements from Oracle OpenWorld, other news in the virtualization space, and how these relate to the future of "extreme BI," as well as what cloud infrastructures might look like from a variety of vendors in the future. I want to thank our panel for joining us for BriefingsDirect Analyst Insights Edition, Vol. 30.

I also want to thank our charter sponsor for supporting our podcast, Active Endpoints, maker of the ActiveVOS visual orchestration system, and Hewlett-Packard via the HP Live! Podcast Series. This is Dana Gardner, principal analyst at Interarbor Solutions, thanks for

listening, and come back next time.

Listen to the podcast. Download the podcast. Find it on iTunes/iPod. Learn more. Sponsors: Active Endpoints, Hewlett-Packard.

Transcript of BriefingsDirect podcast on Exadata, extreme BI and cloud computing. Copyright Interarbor Solutions, LLC, 2005-2008. All rights reserved.