



*Announcer* You're listening to the Sun Microsystems Podcast Network. Welcome to another edition of *Innovating@Sun* with your host, Hal Stern. Today's topic, *Data Warehouse Appliance* powered by Sun and Greenplum. And now, here's Hal Stern.

**Hal :**

**Hello and welcome to another edition of *Innovating@Sun*. I'm your host, Hal Stern, Vice President of Global Systems Engineering at Sun Microsystems, and my guest today is Luke Lonergan, who is the Chief Technology Officer of Greenplum. So, Luke, welcome to the show and, I don't know, tell us about what you do over at Greenplum.**

**Luke :**

Hi, Hal, thanks for inviting me. I'm the CTO and Co-founder of Greenplum. We're a software company that has created a shared-nothing MPD database built on Open Source. And we're a great partner of Sun and vice versa and what I do is I spend time working on future developments. I do a little bit of performance work and I also work a great deal on product development with customers to get our database to be the world's best at business intelligence and data warehousing.

**Hal:**

**So, you threw out a couple of phrases there - data warehousing, business intelligence, database systems, shared-nothing architectures – that I think have been around in the industry for the better part of ten or 15 years. What are you doing that's different? And in particular, what are you using of Sun's that you think is giving you an advantage here?**

**Luke:**

Sure. I think we actually – one great benefit we have is actually timing. So, our team has built some of the world's largest super computing systems in the past, using Linux and Open Systems. And when we looked at what was going on in the massive explosion of data and data warehousing, my thought was that we could do the data warehouse job – the big data management and analytics job – using Open Source and Open Systems technology. So, we just happened to – about a year-and-a-half ago – to find Sun on the up-and-coming with some really great, new technology; and in particular, some of the Thumper box, otherwise known as the X4500 – I like Thumper better – has just all the right characteristics to do this kind of work. So, what we've done is, we've taken our Open Systems, Open Source approach to massive data analytics and database, and put it together with this new Thumper box, and we scale with clusters of those, to create some of the world's largest databases.

**Hal:**

**So, let's dig into that a little bit more. So, inside Thumper you have enormous cross-section of bandwidth from the disks right into memory, clearly, you know, in terms of being able to get your bits on and off the CPU and to do something with them, it gives you an advantage. What happens as you go up the stacks? Are you taking advantage of features in Solaris? Are you taking advantage of some other Open Source database technologies?**

**Luke:**

Yeah, this is really a three-legged stool. Maybe there are more legs, but the key things that I found compelling about Thumper and Solaris in this case – we have some terrific Linux expertise, so I was kind of looking at Solaris from a side point of view – but, when the fusion of Solaris and VFS, plus this hybrid storage server came out, we saw that as an opportunity to really leverage the consolidation of all the things that would normally go into a SAN storage infrastructure into one 4U rack-mount box. And so, we’re actually one of the very first adopters in very large enterprise deployments of the new ZFS file system. And that really revolutionizes the way that you can deploy large numbers of disk drives. And these systems – we currently have a 400-terabyte data warehouse deployed at one site which is comprised of 40 of these Thumpers. And in that 40-Thumper collective, we have something on the order of 2,000 disk drives. So, guess what part fails the most? Disk drives – they fail all the time. And underneath us, VFS is doing predictive self-healing and correcting for any check-sum errors or any media errors, without any adverse effect on the performance of the data warehouse. So that – the combination of the ZFS and the Thumper, which is this hybrid storage and server in one box, plus our parallel database, which basically spans across multiple of these boxes – if you need to grow out of one, you just add more – that combination of three parts really makes up the data warehouse appliance.

**Hal:**

**So, you’re really building a grid of database servers, if I can be somewhat blunt about it.**

Luke:

Yeah, no, absolutely. And I think what’s different here is that, you know, we built grid systems for high performance computing for years. Our team have built systems with thousands and thousands of individual servers, but this is really taking that kind of technology from a decade ago and bringing it right into the departments of major corporations. And I think bringing this type of technology in a package that really is easy for people who have SQL relational data models, but want the massive scale and performance that have been coming through these innovations. And Sun’s been right at the core of that with this new hybrid storage server and the ZFS file system.

**Hal:**

**So, I want to pop up one more layer, ‘cause you talked about a SQL database, and I think that historically, one of the things that’s been hard – and certainly, one of the things I’m hearing from customers now that they feel might get harder, is dealing with the level of parallelism, dealing with this notion of grid computing outside of your father’s high performance compute, the large, scientific applications, you know, large, easily parallelized Fortra [?] applications. So, let’s start a little bit on the database side – what’s that look like?**

Luke:

Yeah, no, that’s actually really interesting. I think there are a couple of interesting developments lately that have captured people’s imagination, and one of them is MapReduce from Google. So, there are more and more people that are really confronting this data growth problem, and people are using data to create information and revenue out of information - more and more. So, kind of this revolution of big data hitting every aspect of the businesses that we’re running has forced people to reevaluate how they do their processing. Yet, everyone really loves the SQL model and SQL has been evolving over time. So, the nice thing is that SQL intrinsically has the property that the data is accessible and tractable, with parallel approaches. So, what we’ve done is effectively created, you know, for lack of a better or simpler term, a MapReduce engine, a dataflow engine that scales across thousands of CPUs, and it does so transparently, through the SQL. So – and we built all this, incidentally, and we’ll get there in this discussion, through using Open Source technology. So, we built our technology inside of the Postgres database, and the main reason for that is, number one, it was a great technical foundation, and number two, it’s not some new, wacky database nobody’s ever heard of. You know, if you pop open a Google search

window and type anything about Postgres, you're going to get lots of people's information about how to use it. So, everything really looks just like a Postgres database with Greenplum database. And the key thing is that the parallelism is automatic and transparent, but that power is accessible to people who are either writing SQL queries; they're doing ELT, you know, transformation of data; or new applications that need things like the MapReduce kind of dataflow. Those are all accessible through our parallel engine.

**Hal:**

**So, starting with some Open Source software, particularly the PostgreSQL database, you're able to go and add features to drive parallelism of this, so someone who's well-versed in SQL, who understands the shape and size and the schema of the data is able to take advantage of what you've done in parallelizing this work across multiple data servers, and then to go up another logical layer to then go create information out of this, to go take the abstraction of rows and rows of a SQL database and turn it into something that will actually go drive the business intelligence, you're using some more software that's available through Open Source, in particular, the MapReduce libraries, to go create a programming environment.**

**Luke:**

Yeah. No, I think the real key here is that people really are, at this point, driven to produce business results quickly. And that's why we did this appliance with Sun. Sun actually sells this system as an appliance; you can get a 100-terabyte system that will roll in in a week, you know, turn on in a day, and you'll be loading data and querying in parallel across all the CPUs, all the storage. You know, we're doing ten-gigabytes-per-second of SQL bandwidth on that one-and-a-quarter-rack system. And that's basically just a built-in thing or attribute of this appliance. And the main thing there is that it's not business value, so people doing the analytics need to get away from having to worry about, how do I partition my data? How do I organize this and tune it and put indexes on it – most of our customers don't even run with indexes. So, the main thing is about performance and simplicity. We have a customer, I mentioned before, who's got a 400-terabyte system and they just loaded a trillion rows into it. You know, the numbers are becoming staggering; and to them, after they were done, it was really no big deal, everything's running so fast that they don't have to think about it. So, doing this kind of work is becoming increasingly important and what we're doing is, we're making it increasingly simple.

**Hal:**

**And again, I think what you've hit on here is the secret sauce, if you will, was in building on some readily-available ideas of others, captured in Open Source software packages, and then producing a product that you sell that commercializes those things and exposes a set of interfaces out to someone who may not be National Labs caliber parallel programmer.**

**Luke:**

Yeah, no, and absolutely, the terrific thing about this is, again, kind of timing-wise, we've been contributors to the PostgreSQL community for many years; we've been building this parallel database for over four as Greenplum; and the really wonderful thing is that what we've been able to do is use that core of PostgreSQL and build into it only the things that really add value to it. And that's one of the powers of the Open Source community. First, we start with this great worldwide community – there were two million people worldwide that download the latest version of Postgres – and that helps us because there's a good foundation of knowledge out there about Postgres. And so, we're really building on that, and we've added to it. So, we didn't have to reinvent all of the things that have been in development for the last 15 years with Postgres – it was already there. So, we don't write drivers, but we get the benefit of all of them - you know, ODBC, JDBC, Pearl – we have a customer who recently – they loved the fact that you can build parallel aggregates and imbedded procedural functions in Pearl. So, they're a big Pearl shop; they love to write things in Pearl, and so they drop these routines directly inside the database. And we're

working on advanced interfaces for things like R – the statistical analysis or statistical computing package; so, they run directly inside the database engine. All these kinds of things, really, are leveraging the power of existing Open Source projects like Solaris, like PostgreSQL, like R, and making those available in a form that's really just a plug-and-play. You get the system, you plug it into the wall and it's ready for you to go. And once you plug it in, you're not only plugged in to run your queries, you're plugged into a community that's doing that kind of work with you.

**Hal:**

**So, Luke, what you just discussed, I think, answered the question that I get asked a lot, which is how do you make money from Open Source software – and you just gave a great answer to it, which is: You go and you start with a set of communities of developers, a set of communities of advocates – that's how you know that there's a market for it and certainly, a community of developers for it - and then, you go and you add commercial value on top of it. And clearly, I'd say if you look at all the features and the system-level integration that's been done by Greenplum, it's an impressive little package - and emphasis on "little" – it is a physically small form factor for an awful lot of compute resource.**

**Luke:**

Yeah, I know and actually, that's where the Tier One systems vendor that Sun is comes into play. We have an advantage that's somewhere between three and six times with this appliance of floor space, power consumption. We hit performance numbers that are in excess – we, basically, are the fastest query engine around and we also load data faster and all that kind of thing; but the key thing is that where it takes some of our competitors – in one case - 33 racks of gear to do that 400-terabyte database, for us, it takes seven. So, the difference in floor space, cooling and power consumption are tremendous. And that's where Sun's been innovating, right? So, what we leverage there is all of that work on the backend side of things to make very efficient units. And in this case, it's this hybrid storage/server. But, the architecture that we had to develop for the software layer to take advantage of that was the novel piece; and so, leveraging to your point, the Open Source community, we found that there were people who already liked the Postgres interface, they already liked using the software, and by combining the innovation with the existing community, what we're trying to do is bring something that is kind of readymade, has a ready audience. The way we make money is we basically monetize that innovation; so, we sell the appliance as a premade package that gets you all that value, and we take payment for having that convenience of having it all put together and continuing innovation that follows it.

**Hal:**

**Great. So, final thoughts – what do you want people to think about? What do you want our audience, the developers, Sun customers, future Greenplum customers, to go think about?**

**Luke:**

Well, I'd say it's time to innovate, it's time to give it a try. Using large data has become a lot easier, so try and break out of the mindset that it's difficult to handle tens of terabytes of data, or even hundreds. It's actually become easier to do that. So, it's time to try some experiments and start to get in the fray. You know, so, the first rule of life is to show up; in this case, there's a lot going on, and I'd say just grab ahold and give it a try. Sun's a great partner; I think you'll find that you can get a system from them to test out very quickly. And of course, we're gonna keep innovating. So, good luck to everyone doing it – I think it's going to be a great decade.

**Hal:**

**Great. Well, Luke, I want to thank you for being our guest here. So, I've been speaking with Luke Lonergan, Chief Technology Officer and Co-founder of Greenplum, and you've been listening to another episode of Innovating@Sun. I'm your host, Hal Stern.**

*Announcer      You've been listening to Innovating@Sun. Join us next time for the latest in innovation from Sun Microsystems. Only on the Sun Microsystems Podcast Network.*