

Cloud Computing

Interview with Russ Daniels and Franco Travostino



Dejan Milojević • Hewlett-Packard Labs

In this installment of Trend Wars, I discuss cloud computing with Russ Daniels, Vice President and Chief Technology Officer of Hewlett-Packard's cloud services strategy, and Franco Travostino, a Distinguished Architect at eBay.

Featured here is an excerpt from the in-depth video interview, which ranged from discussion on the challenges of cloud computing to standards and research. The video is available at <http://www2.computer.org/portal/web/computingnow/videos/trendwars>.

Dejan Milojević: What does cloud computing bring? What is its substance?

Russ Daniels: I think of cloud computing as services being delivered where somebody else owns the assets and is responsible for operating it. If you're doing it for yourself, that's not really cloud, that's just scale-out computing. We see that as being highly related to things such as high-performance computing, grid, and so on. At the top, we think about cloud services, and that's where some of the technologies – DCE [Distributed Computing Environments], Corba, and so on come in – which are really much more centered around how to present capabilities in a form that can be accessed programmatically as well as interactively, and tend to be higher level. You can think of them as, “How do I deliver functional capabilities?”

And there we see these things being delivered today over standard Internet protocols. The user experiences frequently being Ajax based, delivered in the browser rather than through the clients, and then what's in between we think of as cloud platforms. This is, “What does this mean? How does a developer understand how to build things to take advantage of that scale-out approach for computing and to deliver their capabilities in the context of that more service-oriented delivery?”

Franco Travostino: One distinction I make between grids and clouds is that the cloud came out of the Web 2.0 mindset. Grid came out of super-computing teams that have a very strong habit of looking at things that are very complex by nature – and they don't fear complexity – whereas clouds come out of a mindset where there has to be simplicity. It has to be something highly participatory; you take it home, and you do it, and you put your application in the cloud. If it takes more than 10 minutes, it's just not right. That's why we saw the explosive growth of Amazon and Google. Or where you go from 2,500 users to 30,000 users, adding several thousand users every hour. That type of highly participatory and highly on-demand response is new and certainly is going to produce an effect in the industry and is going to be certainly quite a change. As well, we will see a lot of companies that are going to produce cloud examples

within their confines without using a provider, but rather building their own infrastructure and doing cross charges as opposed to having a real credit card, and I'm sure that we will see more of that as the industry matures around this concept.

Milojević: Where do you see cloud computing going in five to 10 years?

Travostino: I imagine three types of cloud manifestations. One is the very large-scale offerings by vendors that are optimized to operate with low margins. That's going to be a phenomenon by which the bigger becomes bigger. There will be a few centers blessed with a lot of hardware—a lot of good natural connectivity in a few strategic places. There will be certain niche places. For example, a GPU cloud for people that are interested in some type of high grade parallel math processing—with some type of silicon – types of unique resources.

Then there will be the third type: companies with large IT departments that will create their own cloud, or they will re-brand whatever they have had for many years like they did for grid. I'm talking about Intel, Boeing, and eBay. With regard to what is going to run, if we look at the cloud as being a platform as a service, then the sky is the limit. We can run SOA, Web 2.0, and HPC. We think that anything can run in the cloud at that level of abstraction, so I think that there is a lot of diversity

there, and that's really going to be the beauty for innovators that will lower the bar to entrepreneurs and users that really want to try out and test the market with new features.

Daniels: First, it's critical that we try to have a bit of balance. If you were to take the sum total of workloads running in the world and ask where they will be running in five or 10 years – or, what kind of architectures will they be running on? – the vast majority are going to be running where they're running today. The industry – if you think about it, we're 45 years from the introduction of the mainframe – has gone through computers, microcomputers, client servers, Internet computing of various forms. You think about all of these things, and the only thing that's really gone away, and even that's not really gone, is minicomputers. The distributed computing worlds have grown, but at the same time, there's still a lot of workloads running on mainframes. Similarly, there's a lot of workloads running inside of IT today that will continue to run in the current deployment models for the indefinite future, and I would fully expect that we'll see significant innovation in that form of delivery as well.

A lot of the virtualization and automation technologies that we're building are really focused on improving the economics of that style of delivery. We make it sound like two years from now there's going to be a flip of the switch, and everything will move to this new world. I just don't think that's true. What will be over there are new workloads that have very specific requirements for scaling. They will need to be able to flex capacity so that when they go from a few users to hundreds of thousands or millions of users, they can do it without falling over. There are stories every day about another service that can't

meet the service expectations that their customers have as they move into these new worlds, and so taking architectural approaches that allow that kind of flexing to occur and that adaptation to demand to occur is critical in having the right kinds of platforms that'll allow people to do that successfully. It's the fundamental difference of the architectural approach that we have to keep focused on.

If I have an integrated architectural approach where I can make choices at the highest levels, from my software all the way down to the choices about the types of hardware devices that I'm using, I have the ability to achieve economies in that approach that I can't get in a classic distributed architectural approach that we have in traditional IT. A lot of it really comes down to where do those economics benefit, and the workloads will move where those economics are attractive. In places where you don't really get that huge amount of fluctuation and demand or you don't have other requirements that would make that kind of world attractive, there's not going to be an economic incentive to move those workloads.

Milojčić: Who do you think will use cloud computing the most?

Daniels: We can deliver compute capabilities at lower cost and at higher scale, and that's all highly motivating in the industry. But to me, there's something else which is much more exciting: we have things with persistent availability, we can get to the cloud, and that stuff that we get to in the cloud is pervasive and persistent; it's there all the time. Then, you can start to use that as a way to go after a continuity and consistency of experience that people have as they move through the different circumstances in their lives across a range of devices.

Today, user experience is about a device and an application. The application might be local or remote, but it's still that narrow, siloed experience that we look at. Being able to use the cloud to get after experiences in a broader way is really exciting, and if we do that well, we can make technology useful for a much broader group of people. It's not just consumers, but it's the people that today are nonconsumers, the people that aren't using technology because it's too complex, too expensive, too hard to get to, and that's really exciting. When I think about what that means to the enterprise, it's less about how to use this stuff for internal IT and much more about how to use this to do the things that I want to do to reach those nonconsumers. It's important to the enterprise much more in how they work with their customers or how they work with their partners. Over time, we'll also use it for our internal stuff, but I think those are things that will be less valuable to move into this form of delivery soon.

Travostino: Anybody who can appreciate a variable cost proposition has to benefit from the cloud. Variable costs need to be mitigated against people, processes, technology, and information. It's not just about technology.

We always have to keep in mind four different dimensions: people, processes, information, and technology. I would say that entrepreneurs and small operations have to be the first beneficiaries given that clouds today don't have five-nine [99.999 percent] or seven-nine [99.99999 percent] dependability. Increasingly, we should see more of the Fortune 500 companies. They will be torn between using their own internal cloud within their own IT confines versus a real external cloud by an external provider which they do not have control over. I expect there will

be some dilemmas between CFOs and CIOs because the CFO is going to look at variable cost numbers, and the CIO is going to feel erosion of his or her control in responsibilities. There will be some creative tussles between those two figures in the corporation. That's pretty much where we are heading. Then services. If we overcome trust [issues] and gain confidence in the fact that a cloud can be managed the same way or even better than how we do it internally within a corporation, then we should see a lot of services built upon clouds, just as we have seen for utilities in other industries.

Milojčić: Do you think users will trust cloud computing providers?

Travostino: Well, I'm a technologist at heart, so I will always say that help is on the way. There are technology nuggets that can be thrown at this problem. There are ways, for example, with virtualization that you can protect the application from a compromised operating system. There have been some very exciting and interesting papers in that area of research.

I think that there will be ways to detect if there's an inside job by a disgruntled employee, so there will be a way to minimize that. But, at the same time, the philosophical and psychological impact of a bank or an institution with private data and moving that data outside of the confines of their IT operation is a huge leap forward, and that's going to find a number of inhibitors along the path.

I'm sure there will be some examples of this happening: the bad press that comes with data being leaked, but we see that today without clouds. It's a matter of growing the IT business and applying those technology nuggets that are providing some help – virtualization compounded with extreme ... to the point that you do not have any

loss of confidentiality and private data, even if the operating system is compromised.

Daniels: I think it's the social dimension that drives these questions. The answers are surprising sometimes. In the consumer space, there are companies that provide ways to consolidate all of your financial information from your banks, brokerage firms, 401(k), and all of those things. There's a little gotcha in this, which is for them to do that, you have to give them your login information for each one of those accounts. It turns out that if you use Quicken to do that consolidation, you're doing exactly that same thing, but the user login information is on your computer so you can feel more secure about that. Well, what we know is that the most frequent way of losing sensitive data is the loss of the physical device. We've had bigger problems with confidential data being lost because a notebook was stolen. There's the psychological comfort in having a local device that takes you in, perhaps, the wrong direction. What happens is that the companies that are doing this are relatively small companies. But if it's presented to you from your bank, you think your bank is doing it, and you're willing to do that. You trust the big bank, but you may not realize that they're actually depending upon some small service to do this.

To me, there's a very interesting thing that is going to have to happen for this to succeed. Today, that bank thinks they own my data, so the decisions they make about how to present that information are made based on their own interests and concerns. They may not want to let any access to it, they may not want to add any functionality to make that more secure so that if I want to hand that over to somebody else it's okay. They don't care. They aren't going to go to

About Russ Daniels and Franco Travostino



Russ Daniels is vice president and chief technology officer of Cloud Services Strategy at HP. In this role, he sets the overarching business and technology strategies for HP's approach to cloud computing.



Franco Travostino is a Distinguished Architect at eBay. Travostino is currently focused on automation, virtualization, and grid/cloud computing.

any time and effort to make it easy for me to say, "I only want to share aspects of that information. I'm willing to show the transaction history over here, but I really don't even want to show the account numbers because I don't want anybody to be able to do anything to affect trades or make withdrawals."

Today the banks don't have any incentive to do that because they think it's really me giving access to their information rather than me giving access to somebody to my information. That change of mind is critical to get through these social concerns. Everybody's going to have to recognize that the data belongs to the individual or belongs to the corporation, and they're the ones who can provide – have to be able to provide – the fine-grained access control to that data.

For much more on cloud computing, a video of the entire interview is available at <http://www2.computer.org/portal/web/computingnow/videos/trendwars>. □