

Position summary: *eOS* – the dawn of the resource economy

John Wilkes, Patrick Goldsack, G. (John) Janakiraman,
Lance Russell, Sharad Singhal, and Andrew Thomas

Hewlett-Packard Laboratories, Palo Alto, CA 94304

{john_wilkes, patrick_goldsack, john_janakiraman, lance_russell, sharad_singhal, andrew_thomas}@hpl.hp.com

We believe that achieving the benefits of a resource economy, which supports the execution of services wherever and whenever is most convenient, cost-effective, and trustworthy, represents the next big computer systems research opportunity. That is, the emphasis in the operating research community should move away from extracting a few more percentage points of speed from individual computing resources, and focus instead on how to size, provision, and manage those resources to serve the needs of a rapidly diversifying set of services. HP Laboratories are embarking on a major endeavor to pursue this, and are actively seeking research partners to collaborate with us.

The vision

Scene: a Corporate IT director's office, the day before a company board meeting. The COO, Chris, knocks on the door, and comes in without waiting.

Chris: Jean: what's all this about us getting billed for computing resources in Singapore? How am I going to explain that? We don't have a facility there! What's going on here?

Jean: Calm down! It's ok—really.

Chris: Not good enough. You know I have to give a bullet-proof answer tomorrow ... so you don't have much time.

Jean: Ok, ok. Do you remember the end of the last month? We had the R&D guys needing to do their protein shape calculations to meet an FDA deadline ...

Chris: Yes—but weren't they using some cheap compute cycles in Prague that you'd found for them?

Jean: ... and then the marketing team wanted a new computer-generated-graphics commercial in time for Comdex that included film of our Malaysian manufacturing plant as a backdrop ...

Chris: Yes—but you said that wouldn't be a problem ...

Jean: ... and the financial results that you are holding, by the look of it, needed some decision analysis that we don't usually do, in collaboration with our Japanese partners ...

Chris: Yes—but you told me ...

Jean: Please—if I could finish?

Chris: Sorry. It's been a bit hectic today.

Jean: No way could the Prague facility keep both the chemists and the video team happy: both the computation needs and storage space requirements were way over the top. And the Malaysian plant has nothing, really, so it looked like it would cost us a fortune. ... All because those geeks couldn't get their timing right.

Chris: [Sigh.] I must have told them a dozen times ...

Jean: But then our *eOS* system discovered that we are co-buying storage space for the Malaysian lot-failure analysis data with our Hong Kong subsidiaries in Singapore; and it

checked out the effects of migrating the data back to Prague and the computations to Singapore.

Chris: But wouldn't that have been a huge hassle to get right? Moving all that data?

Jean: Not at all—I didn't even find out until 2 days after it had happened!

Chris: What do you mean? You let it make a decision like that?

Jean: Sure! It even reported that the average response time for our OLTP jobs were 30% better than usual—they usually get hammered by the decision support people at the end of the month. Probably because of the time-zone effects. If you look, I think you'll find we even saved money - we used to have a team of people doing this stuff, trying to keep one step ahead of the next wave of demands. They could never keep up, so the customers were always unhappy—and they were people who we couldn't really afford to have spending their time on that when there were more important things they could do for us, like rolling out new services.

Chris: But what about the users while things were being changed?

Jean: They hadn't even noticed! My biggest headache is our accounting systems: they make the resource location visible at your level—but nobody else cares.

Chris: You didn't even have to come up with this solution yourself?

Jean: Nope. I didn't do a thing. *eOS did it all!*

What we're up to

The proliferation of computers and the Internet into all aspects of commerce and society is well under way. Many of the fundamental technical issues to do with the components of modern computing systems are either solved, or well in hand. It is our position that the next wave of innovation—and hence research opportunities—lies in the field of aggregating pools of computing resources in support of the explosion in scale, complexity, and diversity of computing services. The *eOS* program at HP Labs is aimed at removing the technical barriers to this happening.

eOS is not a single artifact: it is better thought of as a set of related research activities that are striving towards achieving the vision described above over the next few years. It is akin to other research efforts (e.g., Oceano, Grid, OceanStore) in large scale systems in its scope. The research focus in *eOS* is to discover methods to abstract and virtualize computing and storage resources and make them available upon demand at a global scale.

A fuller version of this paper is obtainable from the URL http://www.hpl.hp.com/personal/John_Wilkes/papers