



## Standards in Cloud Computing Interoperability

*featuring Grace Lewis interviewed by Bill Pollak*

---

**Bill Pollak:** Welcome to the SEI Podcast Series, a production of the Carnegie Mellon Software Engineering Institute. The SEI is a federally-funded research and development center at Carnegie Mellon University in Pittsburgh, Pennsylvania. A transcript of today's podcast is posted on the SEI website at [sei.cmu.edu/podcasts](http://sei.cmu.edu/podcasts). My name is Bill Pollak, and today I am pleased to introduce you to [Grace Lewis](#), deputy lead for the SEI's Advanced Mobile Systems Initiative and technical lead for the Edge-Enabled Tactical Systems Research Team. Her main interests are mobile computing, service-oriented architecture, and cloud computing. Grace has more than 20 years of professional software development experience, mainly in industry. Before joining the SEI, she was chief of systems development for Icesi University where she served as project manager and technical lead for the university-wide administrative systems. Welcome, Grace.

**Grace Lewis:** Thanks, Bill.

**Bill:** In today's podcast, Grace and I will be discussing her recent technical report, [The Role of Standards in Cloud Computing Interoperability](#). Let's begin by having you tell us what interoperability means in the context of cloud computing.

**Grace:** That's an interesting question, because the cloud computing community uses the term "interoperability" to refer to the ability to easily move workloads and data from one cloud provider to another or between public and private clouds. However, if you think about the definition, that really means portability. And, so, in this report, to be consistent, we use interoperability in the way the community uses it. Towards the end of the report, there was actually a discussion on what it would mean for a true interoperability, which is the ability to exchange information between systems, for example.

**Bill:** I see. So, please tell us a little bit about the research described in your report. First of all, why do we need standards in cloud computing interoperability anyway?



---

**Grace:** The biggest fear is really vendor lock-in. People want to have the freedom to move from one cloud provider to another in case the relationship between them isn't working, service-level agreements aren't being met, other providers have better prices, or even if their provider goes out of business, which is not unusual in today's world. If there aren't standards, then moving between providers could be very difficult.

**Bill:** And, how are these standards developed?

**Grace:** Well, like most standards, they are being developed by standards bodies, industry groups, or even nonprofit organizations. The problem right now, which is discussed in [the report](#), is that there are way too many standards. And, one of the things that I even question in the report is if we really need new standards, or whether we can just leverage existing ones.

**Bill:** What are some of the ways that this research could be leveraged by organizations in the DoD and industry?

**Grace:** So, the report starts out with a long list of these standards-related efforts that I just talked about. They're very different types of projects. Some focus on standardizing parts of a cloud-computing solution, such as workloads, authentication, and data access. Others focus on standardizing how the parts should probably work together as a solution. And, then what the report does, it talks about how some of these standards support what I would call "the four, typical, cloud-computing, interoperability use cases," which are (1) workload migration, (2) data migration, (3) user authentication, and (4) workload management.

There is then a discussion about how, from a use-case perspective, the main use cases that really benefit from standardization are the first two: workload migration and data migration. Because those are the ones that would really mitigate vendor lock-in concerns. From a more technical standpoint, what this means is that it requires standardization of things such as virtual machine images or APIs [application programming interfaces] for the cloud storage. Then the report talks about the three main service models, which are [infrastructure-as-a-service](#), [platform-as-a-service](#), and [software-as-a-service](#). And, then [the report discusses] how standardization relates to each of these different models.

One of the conclusions in the report is that the one that really most benefits from standardization is infrastructure as a service because the main building blocks are workloads that are represented as VM [virtual machine] images and storage units that are just data, whether type data or broad data. And that ties back to the first two use cases that I said before, which were workload migration and data migration.

Now, platform-as-a-service and software-as-a-service, they benefit less because the use cases are very, very different. When an organization decides to adopt platform-as-a-service or software-as-a-service, they're not just thinking about extending their basic IT resources. They're interested in



---

the value-added features that these organizations provide—whether it’s libraries and platforms in the case of platform-as-a-service, or application software in the case of software-as-a-service—and organizations select them precisely based on these value-added features. So, what they end up doing is that they end up having a commitment similar to the commitment that you make when you buy any type of software product. And, expecting platform-as-a-service providers and software-as-a-service providers to standardize these features would be kind of equivalent to asking an ERP [enterprise resource planning] software vendor to standardize all of their features. And, this is not going to happen because it’s not in their best interests.

**Bill:** I see. So, tell us what the future direction is of this research.

**Grace:** So, one of the topics of discussion in the report is, “What area should we standardize first?” especially thinking about all the standards that are out there. Right now, the way that we work, mainly in industry, it’s a very manually-negotiated process between cloud consumers and providers. Where the cloud resources are located and the negotiation is done, is all at what I would call design time, and then comes the provisioning process. In the report, this is called the first generation of cloud systems.

We’re starting to support scenarios in which the location negotiation is done at design time, but the provisioning happens either at design time or run time, and the actual instantiation happens at run time. And this is what I call a second generation. An example of something related to this would be a cloud-bursting strategy, in which developers would design a system for an average load, but then the system automatically can balance out to a cloud provider when it reaches its full capability.

The more far-reaching scenarios, or “third generation of cloud-based systems” as I call them, all the operations that I talked about before—location negotiation, provisioning, and instantiation—would happen at run time. And reaching this third generation of cloud-based systems is probably where most of the research is going to focus. It’s not easy because it requires cloud consumers, cloud providers, and software-vendor groups to work together to define standardized, self-descriptive, machine-readable representations of things such as basic resource characteristics, such as size, platform, and API [application programming interface]; more advanced resource characteristics like pricing and quality attributes values; negotiation protocols and processes; and billing protocols and processes.”

So, going back to the question about where to start, I think for now standardization really should focus on the basic use cases of user authentication, workload migration, data migration, and work load management. And, then use those as a starting point for these more dynamic use cases in which negotiation, location, and all that stuff happens at runtime. And of course, there would have to be a real business need for this more dynamic scenario. I would say read the rest of the report to find out more about this question.



---

**Bill:** Grace, thanks for joining us today. You can download all of our technical reports and notes, including Grace's report, [The Role of Standards in Cloud-Computing Interoperability](http://sei.cmu.edu/library/reportspapers.cfm) at [sei.cmu.edu/library/reportspapers.cfm](http://sei.cmu.edu/library/reportspapers.cfm).

This podcast is also available on the SEI website at [sei.cmu.edu/podcasts](http://sei.cmu.edu/podcasts), and on [Carnegie Mellon University's iTunes U site](#). As always, if you have any questions, please don't hesitate to email us at [info@sei.cmu.edu](mailto:info@sei.cmu.edu). Thank you.