

Conference Call Minutes, 2010-04-29

ITANA Meeting Minutes - 29 April 2010

----- Attending

Jim Phelps, University of Wisconsin-Madison (chair)
Marina Arsiniev, University of California Irvine
Tom Barton, University of Chicago
Scott Fullerton, University of Wisconsin-Madison
Paul Hobson, Cardiff University, UK
Jim Leous, Penn State University
Steve Olshansky, Internet2
Todd Piket, Minnesota State Colleges and Universities
Erik Lundberg, University of Washington
Ann Kitalong-Will, Internet2 (scribe)

----- Action Item Updates

No updates; Piet posted updates to list.

----- Cloud Computing Discussion: Shel Waggener (UC-Berkeley) Presenting

Shel briefly described work he and Brad Wheeler (Indiana University) have done in cloud computing, specifically that cloud computing is based on two principles:

1. Scale - Can cloud computing effectively service an entire higher ed. institution? a collection of higher ed. institutions? Scale matters to fundamentally disrupt the current paradigm and ability to be as cost-effective as possible.
2. Flexibility and elasticity/dynamic allocation - How adaptable is the environment? Where can we burst/shrink consumption?

Aggregation of demand is necessary when talking about scale to this degree. The proposed solution is not yet clearly defined, and as architects we need to acknowledge and embrace the level of risk in cloud computing so uses of cloud services are managed effectively and to the best advantage to higher ed. institutions. From an architectural perspective, we need to think about how we are solving our campus problems and how they can be better solved jointly and collaboratively.

Additional points:

- In-house management costs on average \$2/GB per month (fully loaded) to manage storage/scale for an individual enterprise.
- Cloud providers can get fully loaded costs of storage down to under \$.30 GB.
- Server layer: average admin can get upwards of 120-140 servers per admin.
- Scale of cloud computing: 1000+ servers per admin.
- Economics shift from capital expenditure model to operating expenditure model.
- Policy and regulatory issues can be roadblocks to adoption, but are inexorable and will be addressed.
- Underlying policies need to be addressed: e.g. who owns the data? Who has access to the data? Who backs up the data? Where are backups stored?

Open Discussion/Points Discussed:

o Many institutions are currently building private clouds; this is an opportunity to extend private clouds into public cloud architecture in a managed fashion rather than the chaos of different units outsourcing on their own, which can result in higher cost long-term. We need to re-think how private clouds should be integrated into public clouds, to consider policies and regulations involved, and share resources between institutions.

o At the low end of the risk scale: architecture design that are supported by CIO and endorsed by campus leadership, includes policies to require some level of utilization, gives a pattern of design against which common architecture can be benchmarked and extended. In parallel, we need to experiment and share results of experimentation regarding how we model utilization of services.

o At the high end of the risk scale: signing large contracts for large capacity with a given provider locks us in individually to a given situation where we still need to acquire services from other cloud providers, but would still need to interface with all cloud providers without control over the API layer. Cost of collaboration slows us down, but without the scale to bring to cloud providers, we are producing a higher-risk model.

o A consortium of institutions, like Internet2, would be well positioned to bring enough scale to cloud providers. Institutions are negotiating with cloud providers individually, which is a higher-cost model to education. A large enough consortium is needed to manage the chaos of different cloud providers to allow for sustainability and scale. Current model means vendor-lock-in: you are able to get your data, but cannot access meta-data; VM environments are proprietary.

o "Uniqueness" is not important at all layers: for example, branding, UI presentation, data management need to be unique; infrastructure does not need to be. There is no differentiation/value-add to the end user community to being unique at the infrastructure layer.

o Standardizing certain layers of the stack is necessary to be able to leverage them in terms of scale and flexibility.

o High-level & Long-term goals (next steps):

- Overall cloud computing strategy needs to be developed.
- IT models are currently push models; we need to shift to a pull model of IT.
- If we don't develop a strategy/model, every unit at an institution will be instituting cloud computing in disparate manners. The challenge will be to understand where we need a level of coordination, and where it's OK to let go of control.

Shel will send additional resources to Jim, for sharing with the group.

Next Meeting - Thursday, May 13, 2010

2 p.m. EDT / 1 p.m. CDT / Noon MDT / 11 a.m. PDT