

Utility Compute Readiness:

We pose key questions for every “Cloud” first; then we build yours and embed the answers.

Do you know what to ask?

Noel Milton Vega
Cloud Architect / Strategist

Deployment Models for Utility Compute

□ Deployment Models for Utility Computing ...

- Identify and describe the 5 Utility-Compute Deployment Models.
- Next, identify several Business-level RISK/REWARD criteria that must be weighed when evaluating a particular Deployment Model for your organization. For each criteria, indicate how great the RISK/REWARD is for each Deployment Model.
- Name a Cloud Deployment Strategy (not Deployment Model) that you, the CEO, can adopt to simultaneously enhance the business REWARDS and mitigate the business RISKS identified above. This strategy can, for example, help you (the CEO/CTO) defend against future technology obsolescence/interoperability Risk in this rapidly emerging landscape. What is it?

□ Service Provider: Utility Compute (Cloud) Economics 101 ...

- Name the two (2) Purchasing/Accounting methods employed by electric utility companies to meet continuously fluctuating power demands. These map directly to Utility Compute designs.
- Is it critically important that “Cloud” CFOs employ the above power company accounting methods; and similarly that the “Cloud” CTO/Architect reflect it in the physical design. Explain how this is accomplished, and why.
- IaaS “Cloud” consumers will be presented with a catalog of VM sizes from which to choose (e.g. *2CPUs @1Ghz; 4GB RAM; 64GB Disk*). These are known as “Performance Tiers”. Identify short and long term consumer usage pattern-changes & trends that will impact the Lifecycle of a “Performance Tier” and, consequently, also impact the ROI of your back-end “Cloud” infrastructure H/W purchases.

☐ Service Provider: Utility Compute (Cloud) Economics 101 (Cont'd) ...

- Name two H/W Purchase Strategies that dramatically increase the Lifecycle and ROI of your “Cloud” datacenter’s H/W portfolio, by increasing back-end reconfiguration agility in response to changing front-end consumer demand.
- Tier1 VMs will be priced at, say, \$.030/hr.; Tier2 VMs will be priced at, say, \$.060/hr.; etc. What method/formula do you use to arrive at these prices? Under what circumstances might you, the “Cloud” CFO, change these prices over time?
- Virtually every mainstream vendor is offering “Cloud-in-a-Box” type solutions. Based on answers to the above questions, explain why attempting to build a Utility Compute platform using these solutions will not work.

Application Replatforming ...

□ Application Re-platforming ...

- Name the 4 basic “Cloud” Services Layers at your disposal today for re-platforming each of your applications into a “Cloud” setting?
- List the Pros, Cons, and Risks associated with each of those 4 “Cloud” Service Layers.
- Name the 4 Application Migration Methods that correspond to each of the 4 “Cloud” Service Layers identified above.
- Each of the 4 Application Migration Methods similarly come with their own set of Pros, Cons, and Risks. List Pros, Cons, and Risks for each of the 4 Application Migration Methods.

Unique Security Considerations ...

□ Unique Security Considerations ...

- In traditional IT, VM-to-Hypervisor affinity is generally measured in months and even years. That is, they don't move between hypervisors nodes much. Moreover, it is common to have all VMs within a hypervisor node belong to the same Organizational Unit (OU).
- In the dynamic multi-tenant "Cloud" setting however, these static relationships no longer exist. Placement of VMs is arbitrary, often resulting in Cross-OU VMs landing on the same hypervisor node; and they are ephemeral, with affinities measured mere hours or days, not months or years.

In the multi-tenant context, can you identify a simple privacy exposure that arises when VMs in the same hypervisor node replace one another? Describe some remedies to this exposure.

Cloud Architectural Considerations ...

□ Cloud Architectural Considerations ...

- The are Three (3) Core Functional Layers in every Utility Compute infrastructure. What are they?
- Explain how maintaining separation-of-purpose between these 3 layers allows a Utility Compute architect to create a design that allows for the incremental, just-in-time addition of cost-effective discreet, commodity hardware.
- Describe the difference between Capacity-Based infrastructure design and Purpose-Based infrastructure design; then indicate which aligns to Utility Compute design and why.
- Provide a semi-detailed design of a private cloud infrastructure. Use whatever orchestration engine you prefer. Requirements: It must mimic a power company's dual purchasing model; must have multi-datacenter support; must allow for just-in-time incremental insertion of commodity hardware (i.e. CPU + Storage); Cannot use any vendor the "Cloud-in-a-Box" solution.

And much more ...

And much more ...

These are just a few considerations that need to be expertly addressed before you can decide what kind of, or even whether a “Utility Compute” (“Cloud”) delivery/consumption model is right for your organization:

- ✓ First, we have answers! And they address every stakeholder concern. Secure an appointment with us for a 1-hour high-level presentation that answers some of the questions posed here.
- ✓ Next, starting with a 6 month-or-longer *customizable* engagement, we can help blueprint-and-build the foundation of your “Cloud” infrastructure, ensuring that these and other critical design considerations are appropriately accounted for.

Contact:

Noel Milton Vega,
Cloud Architect / Strategist