ABSTRACTING IaaS: A CLOUD MODEL COMPARISON

Ilari Shafer, Michael Stroucken, Greg Ganger (CMU)

OVERVIEW

- **Monitoring and administration tools should work across different utility computing providers / abstractions**
  - **How fundamentally different are providers?**
  - **What (abstract) model should tools and analyses consider?**
- **Maybe for IaaS (Infrastructure as a Service) platforms**
  - **Here: an initial survey of selected major software/vendors**

COMMON ABSTRACTIONS

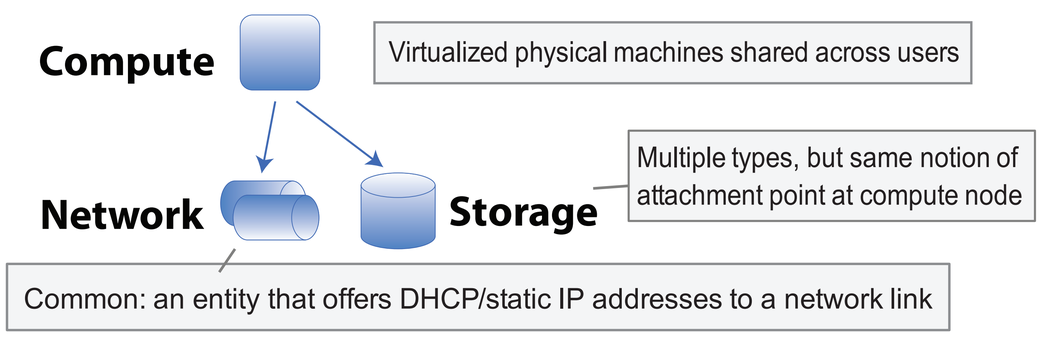

SELECTED PLATFORMS

closed (EC2)+ **open** (Euca)

EC2: largest compute provider
Eucalyptus: EC2 clone by design

open

Focused on scientific workloads
Scripts, Partial EC2 HTTP API

closed SW open APIs

vCloud API and software
Runs on any supported hardware

open

Foundation backed by NASA, Rackspace, Citrix, Dell, ...

Tashi **open**

Focused on data-intensive workloads
Running on OpenCirrus at CMU and Intel

DIFFERENCES

Compute

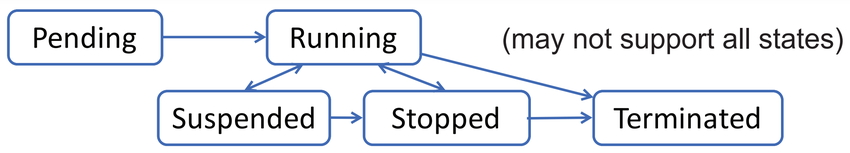

- **vRAM, vCPU (may be fixed based upon "instance type")**
- **Host physical machine (may not be exposed)**
- **Templates/images for VMs, some parameterized customizability**

Compute

- **Few abstraction differences**
- **Different concrete implementations (not abstraction diffs)**
  - **Hypervisors (ESX, KVM, Xen/OVM, Hyper-V, LXC)**
  - **Management data: customization scripts / metadata**
- **Type of hypervisor-level monitoring supported, if any**
- **Existence of leases / expiry dates**

Storage

- **Block Store: EC2 EBS, EC2 Instance Store, vCloud VMDK**
- **Blob Store: object-based, name-addressable (S3, Swift)**
- **Image Store: write-once, only for templates / disk images**

Storage

- **Supported storage abstractions**
- **Exposing storage locality (opaque in OpenStack, vCloud, EC2)**
- **Controlling storage locality (e.g, explicitly local vs on SAN)**
- **Instance startup behavior (e.g, pull?, copy-on-write?)**
- **Presence of snapshot support**

Network

- **Link to network: virtual NIC with IP, MAC**
- **Cloud APIs expose IP of VMs**

Network

- **Level of network abstraction (bridging vs VPC vs vNetwork)**
- **Configuration options (e.g, "elastic IP" in EC2)**
- **L2 networking architecture (vSwitches vs VLANs vs EC2)**

RELATED EFFORTS

- **Cloud standardization**
  - **NIST, IBM Cloud Reference Architectures: high-level**
  - **DMTF CLOUD API and model: low-level, not yet adopted**
- **Cross-cloud libraries/APIs**
  - **JClouds, Apache libcloud, nuvem, Deltacloud**
  - **Designed for least-common-denominator**
  - **Expose compute and object store**

INSIGHTS

- **Core IaaS concepts are shared between platforms**
  - **Possible to define model over many aspects**
- **Still no unified interface for multiple IaaS clouds**
  - **Need to do so to create general tools**
- **Levels of network and storage indirection/abstraction differ**
- **Fundamental differences in underlying network support**
- **Feedback wanted!**

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · Princeton University · UC Berkeley