OpenCloud: A Value-Added Cloud for Internet2

Introduction

The Cloud is rapidly changing the face of the Internet, building on two disruptive technologies: multi-tenant virtualized clusters and Software Defined Networks (SDN). Multi-tenant virtualized clusters enable scale-out designs with flexible resource use and good cost/performance. SDN makes it possible to both manage complexity and customize the network. These technologies are being deployed throughout the Internet; they are not limited to the data center. Network operators are migrating away from purpose-built hardware appliances and moving towards infrastructure that exploits virtualized commodity servers and SDN at the very edge of the Internet.

The emergence of the cloud is impacting all IT sectors, including the Research and Education (R&E) community. This puts Internet2 at the center of the discussion of how Universities and community anchors can best take advantage of cloud technology. This paper puts forward a vision for a *Value-Added Cloud* called OpenCloud, which we propose Internet2 build in collaboration with its member organizations.

Value-Add Strategy

When developing a cloud built-out strategy, it is natural to focus on deploying the commodity servers needed to offer Infrastructure-as-a-Service (IaaS), but one needs to look no further than Amazon Web Services (AWS) to understand the real challenge. The AWS console shows 32 available cloud services, only one of which (EC2) is about infrastructure. The rest are services that Amazon has built on top of their infrastructure, many of which leverage still other services in their ecosystem.

The story at Google is similar. Google App Engine is effectively a front-end to a collection of Google services, including BigTable—a NoSQL database built on the Google File System and the Chubby Locking Service. Recently, Google has layered a fault-tolerant database called Spanner on top of BigTable. The main takeaway from both examples is that the key to success is (1) offering services on top of the infrastructure, and (2) leveraging existing services to lower the barrier to providing new services.

In addition to cloud providers that operate their own data centers, there are a plethora of over-to-the-top service providers that leverage commodity IaaS to build value-added services. Some augment the commodity infrastructure with dedicated resources of their own (e.g., Netflix stores video content in S3, but has built a purpose-built CDN to improve delivery speeds), while others provide all their value in software running on commodity infrastructure (e.g., Dropbox and Box provide value-added storage services on top of S3, Snapchat supports a social network on top of the Google Cloud Platform, and Pertino provides private virtual networks running on top of EC2).

These trends suggest a three-part strategy for Internet2. First, Internet2 should treat existing clouds as providing useful commodity services (rather than try to replicate those services), and focus instead on adding value to these services by leveraging its two strategic advantages—a global footprint and relationships with Universities and community anchors. Doing so gives Internet2 an opportunity to provide value along six dimensions:

- *Responsiveness* – deliver low-latency access to end-users.
- *Bandwidth* – avoid upstream bottlenecks and transfer costs.
- *Predictability* – offer compute, memory, and bandwidth guarantees.
- *Privacy* – keep sensitive data on the customer premises.

- *Trust* – spread risk over multiple commodity providers.
- C*ustomization* – tailor access for targeted user communities.

Second, Internet2 should create an open service framework—similar to the ones that empower Amazon and Google—that lowers the barrier to deploying new services and supporting the rapid customization of existing services for targeted user communities. Doing so is essential to catalyzing and sustaining its ability to provide value-added services.

Third, Internet2 should cultivate all possible sources of value-added services, including community initiatives that target R&E use cases (e.g., scientific discovery and course delivery), experimental services developed by the research community (e.g., fault-tolerant NoSQL databases), and commercial over-the-top-services that can provide even more value if given access to Internet2's footprint. Access to Internet2's infrastructure might be in the form of simple infrastructure (i.e., empty VMs), or building block services (e.g., caches, proxies) that the over-the-top providers can leverage to enhance the service they currently deliver to end-users.

Example Value-Added Service

An example value added storage service, called Syndicate, illustrates the possibilities. As shown in Figure 1, we view Internet2 as an integral part of a multi-tier/multi-provider cloud. It includes commodity cloud providers as the top-most tier (on the left in Figure 1), and additional tiers distributed across Internet2 (the three right-most tiers in Figure 1).

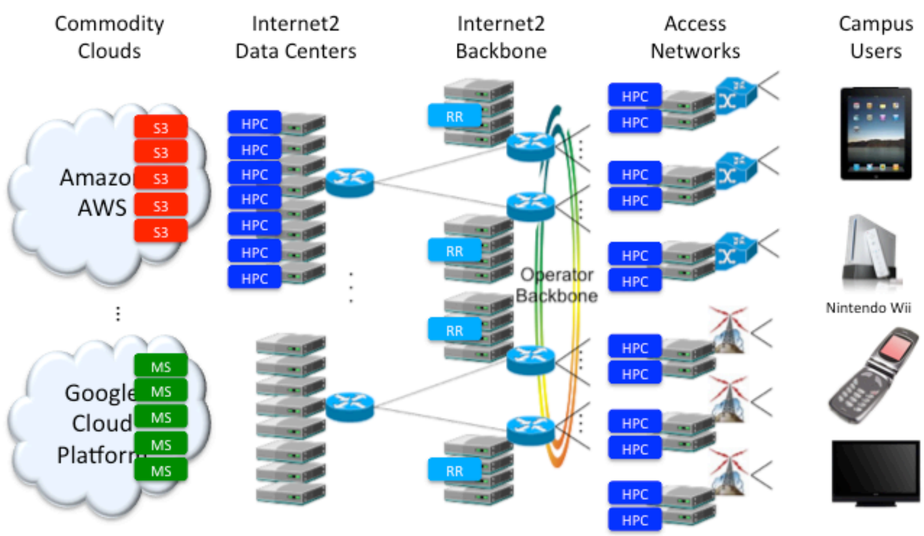

Figure 1. Four services distributed across an Internet2 cloud and multiple commodity clouds.

On top of this infrastructure, Syndicate is constructed from the composition of four scalable services, two of which are deployed on the Internet2's infrastructure (denoted HPC and RR), two of which are deployed on existing provider infrastructure (denoted S3 and MS). The individual services are:

- HPC denotes HyperCache, a scalable caching service based on Akamai's licensed CDN product. HPC provides value by increasing aggregate read bandwidth and reducing end-user latency. HPC is distributed across Internet2's network to form a caching hierarchy, with cache misses at one level resulting in fetches at the next level.
- RR denotes RequestRouter, a scalable request redirection service based on Akamai's licensed CDN product. RR provides value by ensuring user requests are satisfied by the best

available cache taking into account proximity, load, and content type. RR scales trivially, but is replicated throughout Internet2's backbone primarily for high-availability.

- S3 is Amazon's scalable storage service. Syndicate treats S3 as a commodity block device that provides data durability. Figure 1 shows S3 as an example, but a user can elect to use any combination of S3, Glacier, Dropbox, Box, and others as its backing store, where using more than one provides value by mitigating the risk of using any single cloud storage service.
- MS denotes Metadata Service, a new consistency management service implemented on top of Google App Engine and using its NoSQL database to scale performance. MS is necessary to manage cache consistency in the face of write operations, whereas HPC, as is the case with all general-purpose CDNs, assumes cached data is read-only.

We make four observations about this example. First, just because Syndicate currently uses Amazon and Google services does not prevent Internet2 from offering similar services of its own. Leveraging existing commodity services reduces the time-to-market for a service like Syndicate, but if at some point cost analysis dictates that it would be cheaper to build than to buy, Internet2 could build its own durable cloud storage service.

Second, using HPC and RR to augment data stored in S3 is no different than using HPC and RR to augment Akamai's global CDN, Google's Global Cache, Amazon's CloudFront, or Netflix's OpenConnect to deliver over-the-top video content. The idea is that by leveraging Internet2's points-of-presence at the edge of the network, it is possible to add value to a global service. This is the heart of what it means to build a value-added cloud.

Third, while we identify four services as being composed to form Syndicate, there are potentially many others that also play a role, but we consider them part of the underlying service framework and so do not typically name them in a composition. For example, Syndicate also depends on Shibboleth for identity management and OpenStack's Nova (Compute-as-a-Service) and Neutron (Network-as-a-Service) for provisioning infrastructure. This illustrates a core organizing principle of our approach, which we call *Everything-as-a-Service (XaaS),* where providing a uniform interface to a rich combination of building block services lowers the barrier to implementing the next service.

Fourth, in the same way that Syndicate adds value to S3, it is also possible to further customize Syndicate for specific user communities. For example, various scientific communities might use Syndicate to access large scientific data sets (e.g., GenBank) as part of a discipline-specific workflow. Syndicate can also be customized to support community-specific security requirements (e.g., HIPPA), or data consistency requirements. In general, the approach is to provide programmatic interfaces that make it possible to create customized platforms from a combination of existing services.

OpenCloud

OpenCloud is a prototype value-added carrier cloud being built jointly by the Open Networking Lab and Internet2. It runs on Internet2's wide-area backbone, with clusters of commodity servers deployed at five geo-distributed data centers (two east coast US, two west coast US, and one in Europe), ten Internet2 routing centers in the US and Asia, and a growing collection of University campuses across the US.

OpenCloud runs a new cloud operating system, called XOS, on top of this physical infrastructure. XOS adopts *Everything-as-a-Service (XaaS)* as its organizing principle. It leverages OpenStack for low-level VM management, but goes beyond OpenStack to make services a first-class abstraction. In effect, XOS defines a *Service Management Toolkit* that simplifies the process of creating, operationalizing, managing, and composing services, with the goal of making it possible to rapidly deploy new services and program new cloud platforms. An open source network operating system, called *ONOS*, plays an important role in XOS by leveraging Internet2's underlying SDN infrastructure to support service composition.

Finally, we have implemented a collection of demonstration services using XOS, most notably the Syndicate, HPC, and RR services described above. XOS has also been extended to include EC2 and both public and private OpenStack-based clouds as a source of commodity VMs. It is also being used to demonstrate how commercial Telcos can leverage cloud technology to build their access networks.