acmqueue

Moving to the Edge: Overview

The general IT community is just starting to digest how their world is changing with the advent of virtual machines and cloud computing. These new technologies promise to make applications more portable and provide the opportunity for more flexibility and efficiency in either on-premise or outsourced supporting infrastructure. Before taking advantage of these opportunities, however, data-center managers must have a better understanding of service-infrastructure requirements than ever before. The CTO Roundtable on Network Virtualization focuses on how virtualization and clouds impact network service architectures, both in the ability to move legacy applications to more flexible and efficient virtualized environments and in what new functionality may become available. What follows are the key points of interest from that panel. For a more in-depth understanding of what was covered during the panel, please read the full discussion.

—Mache Creeger

The incumbent requirement on IT is to virtualize. It's the only way to benefit from Moore's law. A typical server has incredible capacity—about 120 virtual machines per server—and includes a hypervisor-based virtual switch. Usually, there is more than one server, and the virtual switch has become the last-hop point that touches packets. This allows system administrators to efficiently manage the service platform environment for cost, response time, and capital investment, and to dynamically move workloads on demand from hardware environment A to hardware environment B.

Virtualization decouples the tight integration between software and hardware or applications and operating system, and allows late binding of software by using an abstraction of its supporting infrastructure. As such, it imposes dramatic changes on our notions of infrastructure and network services. Traditionally, application services have been tied to network physical devices, such as intrusion-detection systems, routers, switches, load balancers, and firewalls. With the advent of virtualization, you can no longer reason about a service physically bound to a specific device and that device's relationship to a specific workload. Virtualization breaks the tie between services and the physical devices that support them, and those concepts must be fundamentally reexamined in order to work in a virtualized world.

Abstraction of infrastructure challenges both the notion of where network services reside and old assumptions about physical infrastructure. Binding infrastructure services to specific devices may no longer be relevant to specific workloads. There are two challenges to consider:

- Geographical dependence: the load differences when running an application either locally or at some other location.
- Contextual dependence: the differences in the operational environment when moving a workload from one server to another.

It is important to remember, however, that you will never completely remove the concept of location from networking. It will always be a component of the overall value proposition. Bandwidth consumption will be rarer the farther you go, and latency will be shorter the closer you are. Physical location of networking resources never completely goes away, and the network is never location independent. It will always have a component of geography, location, and physics.

Networking device vendors sell differentiated networking value propositions to their customers,

and most data centers today have highly specialized and configured physical devices supplying infrastructure services to applications. At best, today's cloud architectures typically provide a least-common-denominator service model supporting the most popular devices. Under today's architectures it will be either difficult or impossible to replicate the broad array of specialized physical device services to applications running in the cloud.

Because of its traditional centrality in the data center, the network has been the leverage point. As a result, networks have always been an obvious place to put things such as configuration state. Today, because the semantics are very rich at the edge, the leverage point has moved to the server. Traditional information discovery requiring snooping, learning, and other ad hoc efforts is no longer required, and state information is more easily and directly accessible. The challenge for networking vendors is to define their point of presence at the edge. They need to show what they are doing on the last-hop switch and how they participate in the value chain.

Some network vendors are providing products under the banner of network virtualization that provide a virtual implementation of their physical devices in an effort to preserve a customer's existing investments in legacy infrastructure. Preserving the status quo, however, may prove restrictive in taking advantage of future, more efficient, and more functional network virtualization architectures.

To address the transition of network infrastructure services and avoid being locked into a vertically integrated stack, people who want the dynamics and cost structure of the cloud should do one of the following:

- Wait a little while for better standards before investing in network virtualization. Owners of existing network infrastructure need not worry about the hardware they already have in place. Chances are it provides adequate support for anything they will want to do near term. Future hardware purchases will be driven by a need for more bandwidth and/or less latency and not because some virtualization functions are required.
- Invest today in scale-out commodity networking and make the cloud architecture look like Amazon, Rackspace, or Eucalyptus.

In making any investment, you should minimize location-dependent constraints. You should not be restricted, however, from taking advantage of opportunities that serendipitous locality affords (such as row/rack performance opportunities).

The competition has begun over who will control the network edge inside the server: server, networking, or virtualization vendors. The network edge has disappeared from a dedicated physical networking device and is clearly going inside the server. A server-based architecture will eventually emerge, providing network management edge control that will have an API for edge functionality, as well as an enforcement point. There is debate about how large a role will be played by NICs (network interface cards), I/O virtualization, virtual bridges, etc.

Those who will be left out in the cold during this transition are the folks in IT who have built their careers tuning routers and switches. As the edge moves to the server, where enforcement is significantly improved, new interfaces will emerge. It will not be a world of discover, learn, and snoop, but rather one of know and cause.

People managing servers are not qualified to lead on this change in network services because they don't understand the concept of a single shared network and don't have a clue about networking. Network engineers still have opportunities to add value in a new server-centric virtualized world.

Even though you may structure services to reside at the edge and be VM based, you need a strategy for taking a top-level inclusive view of independent edge-based events. This is especially critical in areas such as security, where you want a global view of a multipoint attack that is below individual thresholds. Top-level correlation will be required to recognize that an attack is taking place.

Although practitioners would like end-to-end virtualization available for service deployment, today server and network virtualization is done in isolation. Given the requirements that overhead virtualization places on application service delivery for a wide variety of SLAs (service-level agreements), whether or not to use virtualization in existing service-delivery architectures is a difficult decision. Q

LOVE IT, HATE IT? LET US KNOW

feedback@queue.acm.org