Moving CCAP To The Cloud

Alon Bernstein
Cisco Systems

Abstract

Network function virtualization (NfV) is gaining traction as a viable method for implementing network appliances on generic compute resources (e.g. the same intel processors used in laptops) instead of custom-built networking hardware.

What would be the benefits of fully virtualizing a Converged Cable Access Platform (CCAP) as an NfV appliance? The presentation will outline the new ways in which a virtual CCAP can solve challenges of scaling, performance, availability, qualification, test and other operational issues in a cost effective way.

Moving CCAP To The Cloud

This paper does not prescribe a specific functional division between the various components that make a virtual CCAP solution. The paper will outline the benefits of moving both the control plane and data plane of a CCAP solution to the cloud (a.k.a. NfV) and the overall network architecture surrounding it.
CCAP owns the physical access to the plant and that part is not ready for virtualization yet, therefor this paper assumes that the virtual CCAP relies on an external solution for HFC access.

Network Function Virtualization Overview

Up to the mid 80's there were not many dedicated network devices. Most routers where implemented on general-purpose servers that had multiple interface cards. However, those general-purpose servers were fairly expensive and as demands for networking increased several companies began building cheaper devices for the sole purpose of routing/switching Internet traffic. Custom built network devices ruled the earth for a couple of decades but as general purpose CPUs became more powerful and better integrated with Ethernet input/output (IO) it was demonstrated that they can be used for forwarding millions of packets per second making them a viable alternative to custom built network devices – in a way completing a circle from server to appliance and back to a server implementation again.
Anything that can run natively on a CPU (a.k.a "bare-metal") can run on top of a hypervisor, opening the door to place these networking devices in the cloud alongside other cloud applications. As a result of these observations ESTI (ref [1]) created an NfV working group to study and make recommendations on an end-to-end framework, including provisioning, testing, monitoring and scaling for NfV solutions.

Running Over a Hypervisor

There is a "virtualization tax" in terms of performance when running over a hypervisor, but there are clearly benefits as well. What are the benefits of running a network function on top of a hypervisor? Clearly there is the original benefit of virtualization: having a single binary software run on multiple types of physical hardware, but there is more. The hypervisor allows multiple virtual CPUs (vCPU) to run over a single physical CPU. The hypervisor can also move a function that was running on one server fairly seamlessly to another server. What these two capabilities provide generally fall into these two categories:

- Availability
- Efficient resource usage

In the following sections we will look in more details to what these capabilities mean in the NfV world.

NfV Benefits

When looking at NfV from a pure engineering point of view of power/performance/space it may not look that appealing. After all it's pretty intuitive that custom-built hardware would be more efficient then a general purpose solution and that hardware optimization would be more efficient then software optimization. However, when taking the operator point of view and especially from a CAPEX point of view a different picture emerges:

- Power efficiency: We need to consider the peak-to-average power. A big router sitting idle at 4:00AM draws more power then an NfV appliance sitting idle or turned off completely. Some of general purpose CPUs are so optimized to save power (mobile devices and such) that security experts claim that they can tell what a CPU does based on how much power it draws.
It's also worth noting that when a CPU company rates a CPU at a certain value it's a worst-case estimate with the floating point processor and other features running at once, which is typically not the case for NfV. On top of all that we have the ability to shrink and expand the number of NfV instances based on demand so that CPUs can be tuned off if not used. Because of the above the average consumption per-hour of an NfV appliance may end up being attractive.
- Granular scaling: similar to the point above but from a different angle: An NfV solution has very fine grained scaling, and so some power/space efficacy is derived from the fact that an operator can use exactly (!) the right amount of resources. Based on the monetization model this may mean that the operator pays exactly for the resources used and that translates to a cost benefit.
- Reuse of resources: A general-purpose server can perform other services at off-times when NfV functions are not needed. For example, it can do packet processing by day and payroll processing by night. The power/cost/space advantages come from the fact that the compute resources are optimally utilized.
- Availability: the ability to quickly move virtual machines from one server to the other is playing a key role in providing high availability. If a virtual machine fails it can be re-instantiated on a different server. Geographical redundancy is also possible - a server at a remote location can take over in case of a catastrophic even if the local data center is down.
- Feature velocity: network appliances are embedded systems that require embedded systems disciplines. Development on general purpose servers is easier because most software developer are familiar with that environment and the software development tools associated with it, therefor the time to develop and test software is shorter.
- Fast qualifications: because of the granularity of NfV services its possible to test new code drops for a very small population and slowly increase the deployment of the new code drop as it proves itself. This can reduce qualification times by being more aggressive in moving to a production environment because the risk of global failure is greatly reduced.

Scale out vs. Scale up

The discussion on scaling benefits with NfV warrants a separate discussion. The key message is the following:

NfV is not about performance. Its about scaling

What does the above statement mean?
The traditional approach in physical network appliances is to deal with increasing

bandwidth demands and feature requests by "scaling up", i.e. building a family of small/medium/large appliances with a larger number of linecards and faster network processors as demand grows. In contrast, the data center approach is to "scale out" which means to build a basic service instance that gets replicated (dynamically) with resource demands. So, for example if a virtual CCAP reaches its maximal packet forwarding capacity the operator – or more precisely an orchestration system - can instantiate a new virtual CCAP to handle the extra load.

The cloud world offers a feature called "cloud bursting"; when the local cloud runs our of resources the system can use another cloud, such as a public cloud offered by several companies, to take over. This is helpful if the peak conditions are relatively rare. For virtual CCAP there is another type of "cloud burst" that is possible; one where the average demand is handled by a physical device and only peak conditions by the cloud. This allows building a very cost optimized solution making the best of both physical and virtual versions of CCAP.

NfV Network Architecture

The basic physical building block of an NfV solution is a server. The server is typically a collection of multi-core CPUs, one or more network interfaces and storage. The multi-core CPUs can host multiple virtual machines, which in our case might be multiple instances of a virtual network function. But how do all these machines connect to the Ethernet and how can they exchange information between each other? The entity that helps with that is the vSwitch (Figure 1):

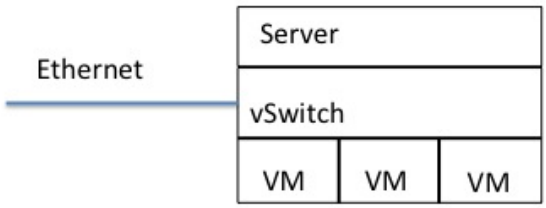

Figure 1 Server and vSwitch

With the magic of virtualization each one of the VMs "thinks" it owns the Ethernet to the outside world. The vSwitch, which typically uses one of the CPU cores, switches traffic between the VMs and traffic in/out of the physical Ethernet port connection.

The next level of how an NfV assisted cable network may be connected is depicted in Figure 2. Traffic from the Internet is directed at the front end of the data center that handles

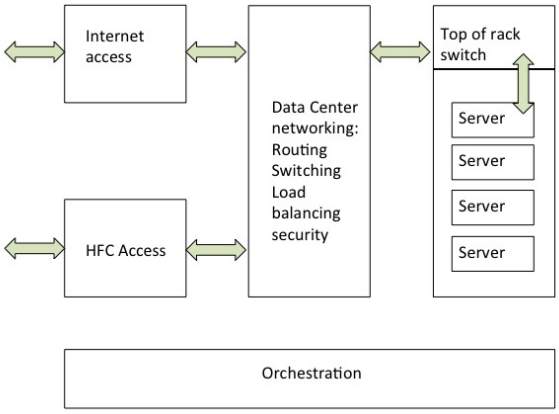

Figure 2 NfV packet processing

basic functions such as load balancing traffic between servers and first level sanitization of the traffic - for example denial of service protection. The next stop is typically the "top of rack switch" that sends the packet stream to the right server (in a data center the servers are stacked in a rack and connected to a top of rack switch). Once the packet stream is intercepted by a server it gets switched to the right VM via a vSwitch. The access to the HFC is essentially a mirror image of the process described above.
Considering how dynamic the data center is and how VMs and functions can move

around, the need for automation in this system is obvious. This automation is referred to as "orchestration". The following user story helps outline the role of orchestration in the workflow of creating a virtual CCAP.

A User Story: Virtual CCAP instance workflow

The following section will outline how a CCAP instance can be created:

1. A cable operator notices an increase in traffic demand and the need to add a virtual CCAP instance.
2. The operator can look in a "service catalogue" for a virtual CCAP function (as we cover in the next section other services may be firewall, deep packet inspection, parental controls etc.).
3. Assuming that an HFC access is already available then all the cable operator has to do is add the virtual CCAP function to the HFC segment that is under load.
4. The Orchestration system takes care of creating the right path between the Internet, HFC access, data center switching/routing, fiber interconnect and vSwitch.
5. Service up and running. No need to locate new linecards or CCAP chassis, no need for an install or hook up any equipment.

Virtual CCAP Q&A

The following sections are answers to frequently asked questions about virtual CCAP.

Is the virtual environment stable enough for packet processing?

The virtual environment has a reputation of being somewhat unpredictable, however, this reputation is not justified. Issue with variable latency and packet drops start showing up when we over subscribe (or close to oversubscribe) the server – not that different then issues with oversubscribing bandwidth on a physical appliance. So when provisioning NfV one should be careful to allocate enough resources and possibly set up VM scheduling for reliable operation of the NfV appliance as well as take into account other users of CPU in the sever, in particular the vSwitch and the hypervisor itself. If provisioned correctly the virtual environment can be as reliable as a "bare-metal" environment where the OS rides directly over the hardware.

How is a virtual CCAP related to SDN?

SDN, in a nutshell, is about separating the control/management plane from the packet forwarding in the data plane. The key to NfV is packet processing in a virtualized environment. NfV appliances can be controlled by an SDN control plane or can operate independently so that an NfV appliance may or may not be part of an SDN solution. In other words NfV and SDN are two independent technologies that are designed to solve different problems. It so happens that both can run in the cloud (aka "virtualize") but that's about all that's common between them.

Is the vCCAP a CCAP replacement?

One common question about the vCCAP is if it's going to replace the physical CCAP. The short answer is "no". The vCCAP should be viewed as another packaging option for a CCAP that is cost effective in certain cases, most likely the areas where a small scale CCAP fits in today. Another use case of the vCCAP is to handle peak usage, i.e. a physical CCAP can handle average traffic loads scenarios but the vCCAP can kick in with extra capacity is needed.

What if the data center is not in the normal data path?

Different operators may have different designs for their data center. The two base options are to (a) build a massively centralized data center a-la Google/amazon (b) build a distributed data center.
The first option could be a challenge for NfV. If the data center is built around video distribution and is physically remote from Internet peering points then packets would have to loop in the network in order to provide service through a virtual CCAP. Even in that environment a virtual CCAP can be used for certain applications (help with test/qualification, geographical redundancy etc.) but it would be harder to deploy at scale. A distributed data center with servers closer to the end users lands itself more easily to a scale deployment of virtual CCAP since the packet flows are similar to those with a physical CCAP network architecture.

How is CCAP Related to Service chaining?

NfV appliances handle dedicated functions. For example, one NfV appliance can be dedicated to deep packet inspection, another to a firewall and a third to a virtual CPE. Service chaining is the ability to set a pre-determined path that is per-application (or subscriber or service) along that path, e.g. one application can go through parental control and deep packet inspection and another application can be directed at only a parental control service. A CCAP appliance, either physical or virtual, is very likely to be both at the entry and the exit of a service chain. At the entry the service flow classification can assist with the selection of a service chain, and at the exit of the service chain the application of Quality of Service (QoS) is the last feature to be applied.

How are virtual CPE and virtual CCAP related?

The CPE can be swept along with CCAP in the virtualization wave. Once we have both CCAP and the CPE in the cloud they start appearing as elements in a service chain, where CCAP might take care of some aggregate policies and shared resources while the CPE is tailored around a specific subscriber needs. In other words, the traditional breakup of functions between a CPE and a CCAP device will be challenged once both move to the cloud.

Conclusion

The following table would help place a boundary around what virtual CCAP is and is not:

Virtual CCAP Is	Virtual CCAP Is not
Scaling optimized	Performance optimized
Connected to node	Integrated CCAP
Dependent on the data center	Standalone device
Dependent on orchestration	Self managed
Data plane processing in the cloud	Separation of only control plane

Network function virtualization is a viable implementation and deployment option in certain use cases and a virtual CCAP can fit well inside an NFV ecosystem. The operational benefits that made virtualization a success in storage and compute apply to networking in general and more specifically to the virtual CCAP where it can be the right tool to address certain operational challenges in a cost effective and scalable way.

References

1. http://www.etsi.org/technologies-clusters/technologies/nfv