"Moore's law" for SDN and Cloud?

By Open Networking Summit - March 11, 2016 - 12:31am

Blog Post by Dr. Marcelo Spohn, Mukesh Kumar and Raghu Gangi, ADARA

"Moore's law" essentially stated that the number of components per integrated circuit would double every two years. The increased demand for higher processor speeds and processing power was a powerful and positive feedback which drove a significant increase in the supply of technological capability. This revolution advanced the industry in ways we still benefit from today. A similar revolutionary advance has been developed and is being demonstrated by ADARA, CALIENT, and ON.Labs with support from China Unicom, Centeris, and C Spire in ONS Booth #17, as well as the SDN Solutions Showcase (S3) and the March 17th ADARA CALIENT On.Lab Speaking Session.

"Moore's law" gave us the powerful general purpose processors (CPUs) and later the Systems on a Chip (Merchant Silicon SOCs) that we involved in SDN and NFV have come to rely on. There are lessons to be learned as we contemplate our own revolution in networking software design. What can those of us involved in SDN, NFV, and Cloud learn from our hardware colleagues? As it turns out, we can learn quite a lot in fact.

SDN and NFV are facing a situation that has parallels to the revolution in integrated circuits'. SDN and NFV have enabled Cloud Computing and Networking to become more reality than a concept. Cloud Computing and Networking has already evolved from an architecture and service which made the logical grouping of resources its primary basis for operations. Increasingly, Cloud Computing and Networking now must incorporate a real-time component; this need holds true for SDN and NFV, upon which cloud services are built. Performance-based networking and computing is now required to keep pace with industry. The trend in SDN, NFV, and in cloud services is toward computing and networking that is increasingly High Transaction, Low Latency, High Bandwidth, High Volume, and Highly Administratively Unpredictable. This trend toward performance as the basis of decision-making in SDN, NFV, and Cloud is a self -fulfilling prophecy; just as it was in the case of "Moore's law".

At its heart, SDN is the replacement of vendor implemented logic with remotely generated logic to facilitate programming virtual and physical network infrastructure. NFV is the replacement of purpose-built hardware with virtual appliances and functions. Both developments inherently increase operational flexibility and administrative convenience. We can acknowledge this and then ask ourselves the obvious questions (which never seem to be asked interestingly enough). How are SDN and NFV inherently better performing? How are these replacements, for programming logic or for purpose-built systems, any better performing when they rely upon the same set of legacy protocols and many of the same technologies? How is a virtual version of an appliance inherently better performing than its hardware ancestor? How is administratively partitioning (i.e. virtualization) of a finite spectrum of physical connectivity a true solution to a geometric spike in demand for real, physical bandwidth?

"Moore's law" accurately predicted the demand for performance increases would consequently drive the need for the supply of increased performance. When we apply this to software networking developments, specifically SDN and NFV in an analogous way, we easily recognize that new technological approaches are not only required, they are long overdue.

The original Internet (ARPANET) used a distributed routing scheme based upon measurements of queuing delay, i.e. performance, at each and every node. These performance values were propagated to all nodes and the lowest latency paths were calculated. The original scheme suffered oscillations under load; today we can still see such oscillations (e.g. route flapping). Thus demonstrating the difficulty of (a) utilizing performance as the basis of a networking computations and (b) achieving performance as the result of computation in a large scale networked topology. As networking developed from its' initial implementation, and as the industry recognized the difficulty in mastering the use of metrics which directly and reliably correlated to performance, performance-based networking was replaced. Substituted in its place was a networking scheme based upon metrics which did not correlate to performance. Since they did not correlate to performance, why were they adopted? Because they were easy to implement; they were largely static, and they offered stability under load. It may surprise many to learn that these schemes are our standards, still in use today. These standards use distance with some administrative weightings, (BGP/ iBGP; Router Hops / eBGP network hops), or interface speed (OSPF); statistics (heuristic's) or hashing; which are complete performance agnostic and yet is the preferred implementation of virtual networking (e.g.VxLAN) overlays; we even see physical constructs (SRLG / Shared Risk Link Groups, contiguous spans of fiber optic cables connecting two nodes) used; again these values are not rooted in performance, and they are not deterministic for performance. Here is further proof: according to the default OSPF metric cost value calculation, the default OSPF cost for a Fast Ethernet interface (100 Mbps) and Gigabit Ethernet interface (1 Gbps) are the same despite an order of magnitude difference. Clearly existing standards networking; hop counts, interface speeds, hash values, fiber strands is very loosely correlated to performance, if at all. A quick survey of every current vendor finds they all use BGP or OSPF or Hashing. As an industry, we have come to rely upon static and completely performance agnostic mechanisms'. Regarding performance based networking, the industry has stagnated. The development of externalized path computation applications does not solve the issue; you can implement a computation engine; but when the input to the computation is static non-performance metrics, the output will be just as performance agnostic. Externalizing the computation does not solve the problem by itself. In short any good developer knows bad input equals bad output. Here is definitive illustrative proof; every major provider, even the ones that implement flat, 1 and 2 Tier Clos (Fat Tree) pod scheme networks suffer significant congestion, especially from the elephant flows (Container and VM Migrations that are the hallmark of Cloud Computing and NFV). The approach has been to statically and administratively overprovision virtual network segments'; "Moore's law" hangs in the air stating analogously; how long will that delay the inevitable?

As the authors of OpenFlow correctly stated; networking had "ossified". OpenFlow was a pivotal event; it created a unified interface; initially to the packet layers (OpenFlow 1.0), and later to the optical layers; (OpenFlow1.5). We now have unified interfaces' for Multi-Layer SDN, the integration of the Virtual (VxLAN) and Physical (IP/MPLS) Packet Layers and the Virtual (DWDM) and Physical (OCS) Optical Layers as well as the Application and Session (TCP) Layers. Now the issues are coordination and continuous optimization of the Packet and Optical Layers. Each Layer represents differing communication mediums, with differing parameters and differing methods for computing costs; here again the SDN /NFV analog lesson of "Moore's law" confronts us. All Layers, with corresponding Control Planes, need to be federated; and the Real-Time Performance of each and every layer needs to be known and normalized to provide a computed response to Service requests; this is what is required for Multi-Layer Packet Optical Integration, or SD-PON (Packet Optical Networking). Standards protocols can and should be used for signaling, but advanced technology had to be developed for SDN Multi-Layer Packet Optical Integrated Networking. An additional benefit of this Multi-Layer performance-based knowledge is that the entire range of imaginable services, from SD-WAN to Cloud Data Center Interconnections, vProvider Edge, Intelligent Workload Placement, and Universal Access, as well as other High-Value Services are now Network Primitives; we get them for free as a fallout from this implementation, just as vendor choice and custom programmability is a fallout from Open Standards Interfaces (OpenFlow).

programmability is a fallout from Open Standards Interfaces (OpenFlow)

The approach is not theoretical; it has been developed. At ONS 2016, SDN Multi-Vendor, Multi-Domain, Multi-Layer Performance Based SD-WAN Packet Optical Integration is being demonstrated in Booth#17, in the SDN Solutions Showcase (S3), and it is also being presented at a Speaking Session. ON.Labs have led the way for this industry leap forward. The development of the ONOS Controller is a key advance that has made such a technological development possible. The open nature of ON.Labs Project has enabled a ready linkage to OpenStack for Performance-based Choreography and Orchestration via collaborators contributing externalized PCE's with Events Based Optical Intents. The integration of Cloud Data Center Computing and Networking to Service Provider WAN/Transport is immediate, reliable and logical. Because services can be impeded at any point in the End-to-End infrastructure, every layer and every point must be managed in a coordinated manner; now that is possible.

The demonstration highlights the Open Community Components; in a Multi-Vendor, Multi-Domain implementation that is an industry first.

- The architecture applies to any Cloud Data Center, Service Provider or Enterprise
- It can be implemented transparently, incrementally, and immediately to Brown Fields, or Green Fields
- The solution mitigates/eliminates deployment risk
- Manages normal operations, faults/restorations,
- Results in higher utilization rates, faster response times, higher availability
- Lower CAPEX/OPEX
- Zero Touch discovery
- Configuration, Provisioning, programming, and continuous optimization
- Manage the virtualization of resources (multi-tenant and micro-segmentations) as well as the creation and use of physical resources (real additional physical bandwidth) in a real-time performance-based manner that is fully coordinated between each and every layer, and each and every hop in the End-to-End infrastructure.

"Moore's law" is once again insightful, predictive, and that insight usefully applied. Hopefully, Gordon Moore would be pleased.

 ONS 2016's blog