

Does OpenFlow Make Sense in the Enterprise Networks?

With Arista, Avaya, Brocade, Cisco, Extreme Networks and HP
The Data Center LAN Evolution Series

Transcript and Digest, October 2011

• Introduction



Dr. Jim Metzler, Moderator, Co-Founder Webtorials Analyst Division

There is considerable discussion in the trade press about the value of OpenFlow. There has been a lot of interest recently in OpenFlow - a communications protocol that enables the separation of the control of packets from the forwarding of packets. By separation is meant that the forwarding of the packets occurs on an OpenFlow switch and the control of those packets occurs on a separate controller.

The OpenFlow specification itself is being developed by the Open Networking Foundation (ONF).

One of the things that is interesting about the ONF is that its founding and board member are Deutsche Telekom, Verizon, facebook, Google, Yahoo and Microsoft. At first it may seem strange that companies such as Google, facebook and Yahoo are so involved with the development of new communications protocols. However, given that separating the control and the forwarding of packets onto separate devices is somewhat of a radical idea, one could argue that the initial advocates would have to be non-traditional players.

The definitive paper on OpenFlow is entitled "[OpenFlow: Enabling Innovation in Campus Networks](#)." The paper was written in 2008 by researchers at some of the US's most prestigious universities; i.e., Stanford, Berkeley, Princeton and MIT.

The first sentence of that paper states "This whitepaper proposes OpenFlow: a way for researchers to run experimental protocols in the networks they use every day." That sentence sets up the theme for this month's discussion: ***Does OpenFlow actually enable the innovation and cost savings that the articles in the press have been talking about or is OpenFlow just a science experiment by some really bright people?***



- Q&A



Dr. Jim Metzler, Moderator, Ashton, Metzler & Associates

What potential benefits does OpenFlow offer for the typical enterprise network?



While OpenFlow is an interesting concept, it's not the first attempt at trying to obtain greater flow visibility or to enhance Layer 3 capabilities across a networking infrastructure; most industry veterans will remember the attempts the market had with ATM LAN Emulation, and also PBB-TE. The reality is that many of these concepts, although promising in a small, controlled, and experimental environment will present significant challenges when scaled-out for the real-world. For Enterprise businesses, or even more so Service Providers, attempting to deploy such a paradigm across a large and diverse infrastructure would require a demonstrable cost/benefit upside; for many, especially those burnt with previous incarnations, this is a challenge they will probably not be willing to again take on.

At its essence, OpenFlow articulates a separation of the Control and Data planes, and indeed can be a beneficial model in certain scenarios; typically with the promise of enhanced performance and optimized scalability. However is this indeed the right model to best optimize performance and provide the agility required for the dynamic Data Center and true virtualization..? Will a fully static approach deliver enough flexibility for tomorrow's Cloud solutions..? The general concept is something that Avaya is intimately familiar with, having has implemented this very model with its newly-released 802.11n WLAN architecture; we call this capability "Split-Plane". But going back to the broader question of OpenFlow, perhaps "is routing really broken..?" might be a more pertinent question. Do we gain something that we genuinely need..? If we to apply the all-important "so what" test, it's quite probably that we'd struggle to find a clear benefit proposition for this technology, when taking all of the additional complexity under due consideration. Something to watch maybe, but this is probably not the most significant problem that we need to be solving..?



The typical enterprise network is becoming complex with the proliferation of virtual machines, mobile devices, and network-attached devices such as surveillance cameras, etc. Like virtual machines can be deployed on servers, virtual (or logical) networks can be supported on top of the physical enterprise network allowing management of virtual networks (e.g., attachment of servers or virtual machines to a virtual network, etc.) independently of the management of the physical network. Network virtualization using OpenFlow can simplify the operation of such networks by creating virtual networking layers to manage authentication, security, and mobility separately from the physical layer. OpenFlow also addresses the challenges being encountered by Service Providers by enabling Hyper-scale data center solutions, network virtualization solutions, and flow management for the WAN.



As a long-established leader and visionary in the field of networking, Cisco sees software-defined networks playing a key role in the ongoing evolution of networking. SDN offer a way for customers to take easier to take advantage of

the sophisticated features of their infrastructure and to bring applications and infrastructure closer together. Cisco is supporting the Open Networking Foundation (ONF) and OpenFlow as a way to advance and standardize this technology. We expect the efforts of ONF to advance OpenFlow to the point that it is suitable for production environments. As part of this effort, Cisco is in active development of OpenFlow support in its Nexus portfolio. For more info: [OpenFlow: "Pulling networking into the application stack"](#)




Today, networks are more-or-less deployed and managed physically, using device-level management tools and traditional technologies like VLANs. This approach has resulted in networks that are static and don't respond well to changes. New connectivity and innovations that requires different policies and configurations takes a long time since those networks are too inflexible and can't be adapted fast enough. Because OpenFlow-enabled solution allows users to manage the network more proactively and in a more centralized way the network can be more dynamic and responsive to business needs and less costly to administer. OpenFlow allows administrators to programmatically control the traffic flow with centralized controllers to dynamically provision and orchestrate the behavior of the network.



Dr. Jim Metzler, Moderator, Ashton, Metzler & Associates

I was talking with Kyle Forster of Big Switch and he raised the issue of needing very fast switch control planes to run OpenFlow with high control plane performance. If OpenFlow becomes successful, what kind of switch architecture is going to be right for OpenFlow?

 A high performance, Layer 2/3 non-blocking data center Ethernet switch is required to support flattened network designs, be it for cloud, storage, HPC and similar objectives. OpenFlow technology can stress the control plane of a switch, depending on the way in which OpenFlow is deployed. For example, if flow set-up is reactive (based on traffic flow) instead of proactive (based on routes or switched paths). In this respect, OpenFlow is not unique in that there are other protocols which can stress the control plane in bursts (such as OSPF recalculation after a topology change in a large area). The most important characteristic needed to support OpenFlow on the network switch is the capability to quickly add, remove or re-order flow entries in the hardware. Overall, OpenFlow does not require any changes to the switch hardware architecture.



OpenFlow needs a switch architecture optimized for classification of traffic and flexibility for actions taken on the traffic. The HP Networking ASIC architecture makes use of flexible TCAMs to give broad classification capabilities and uses dynamic programmability to give flexibility in actions taken on the traffic.



Steven Taylor, Webtorials

OK... I'll jump in here on behalf of our Webtorials community.

- 1) Can you please define TCAM and what makes it special in general?
- 2) Can you please be a bit more specific about how the use of TCAM is specifically well-suited for an OpenFlow implementation? A specific example or two would be really useful.

Thanks!



The short answer is that it's probably too early to say what is genuinely possible, what might be truly optimal, and might ultimately be viable in the commercial context. There is an open question with OpenFlow, where the network takes on the roll of being the control plane and what implications this has for stability and reliability; much will depend upon how the standard evolves. More broadly, as an industry we need to be supportive of attempts at creating standardized interfaces that simplify and/or improve the implementation and operation of evolving applications and data infrastructures; our open Fabric-based infrastructure for Data Center and Campus being a good proof-point of what is possible when interoperability is a key tenet. While OpenFlow has yet to reach that point where deployment into today's real-world environments could be envisioned, we do see the concept of Software Defined Networking (SDN) as a logical development for the industry.

Virtualization of the network is finally enabling a paradigm shift away from traditional IT business models, but tighter integration between the network and the applications being transported by it is the crucial next step to removing organizational silos and providing a truly consumerized user experience in the business environment. As with any major shift, we expect to see technology alternatives appearing sooner than the general business population is ready to accept the cultural shift that accompanies this evolution. With that in mind, we are closely following this technology trend, while also exploring avenues for tighter integration of our own applications with the network; Avaya, being a multifaceted communications company, obviously has special interest in optimizing the end-to-end experience.



Joanie Wexler, Analyst/Editor, Webtorials

For those of us who are just learning about OpenFlow, can someone explain what the problem is that OpenFlow purports to solve? Routers and switches have supported separate control, data forwarding and management planes for years. Why is it necessary to separate these functions into different pieces of equipment? In fact, that seems to buck the trend of collapsing multiple virtual machines into single servers (or virtual routing/forwarding, VRF, instances into single routers). Thanks for any light you can shed! –Joanie



Dr. Jim Metzler, Moderator, Ashton, Metzler & Associates

Joanie, let's look at your question in two parts. The first part is does OpenFlow fit in with some of the megatrends in the industry? The answer is definitely yes. You mentioned server virtualization. What typically happens is that IT organizations take servers out of branch offices and put them into centralized data centers and then they virtualize them. IT organizations are motivated to do this to both save money, have better security and gain more control – in this case over their company's data resources. I will suggest in a moment that those factors are also driving OpenFlow. Another burgeoning trend in the industry is desktop virtualization. With desktop virtualization, the applications are centralized and the user's device could be a fully functional PC or a really dumb device. Again, a very similar situation with OpenFlow.

The paper on OpenFlow that is referenced at the top of this discussion (OpenFlow: Enabling Innovation in Campus Networks) makes the comment that "Virtualized programmable networks could lower the barrier to entry for new ideas, increasing the rate of innovation in the network infrastructure." The paper gives a number of possible scenarios. As an example of the breadth of the paper, one is mobile wireless VOIP clients.

So, one way to look at OpenFlow is that it is like a hypervisor on one of the virtualized servers you mentioned and provides a new level of control and security. As you know, hypervisors provide a lot of really powerful management functionality and have APIs to hundreds, if not thousands of other companies (e.g., the massive size of VMWorld?) so that they can add value on top of OpenFlow. That is one potential advantage for OpenFlow – open up networks to a wide range of innovation and integration.

Another potential advantage for OpenFlow goes back to my reference to virtual desktops. When you virtualize your desktop and keep applications in a central site, you might keep a fully functional desktop or you might just use a really dumb device. It is possible to envision a world in which OpenFlow is widely deployed and switches and routers perform pretty much the way they do now. It is also possible to envision a world in which OpenFlow is widely deployed and switches and routers are just dumb, low cost forwarding engines.



Dr. Jim Metzler, Moderator, Ashton, Metzler & Associates

I received the following comment from Dick Willson of Allied Telesis.

The definitive paper on OpenFlow "OpenFlow: Enabling Innovation in Campus Networks" was written in 2008 as a research paper. OpenFlow was developed so that experimental network protocols could be deployed at reasonable scale to run in the networks that was used every day. As a result of this research the OpenFlow protocol is offered to the industry as a technology that could virtualize the network infrastructure, in the same way that the "hypervisor" technology virtualized the server.

Thought Leadership

Does OpenFlow Make Sense in the Enterprise Networks?
October 2011



Today's routers and switches are bloated with proprietary "features" that vendors persuade their customers to use so that customers then become constrained into a single vendor end-to-end solution. This was done many years ago when the IT industry operated mainframes – proprietary hardware and proprietary operating system. The IT industry has moved on, the network industry should do the same.

OpenFlow only provides the interfaces that enable the tools to be developed that would virtualize the physical network infrastructure it does not provide any "applications".

It will still take a few years of pilot implementation and deployment, just like it took a few years before the hypervisor was accepted as legitimate deployable technology for servers.




Steven Taylor, Webtorials

The original paper(s) on OpenFlow involved its use in a "Campus" environment. While that environment would not necessarily be an academic environment per se, it does seem to imply a physically limited area.

It's no secret today that the data center today may involve multiple physical locations and a combination of premises-based equipment and cloud-based services.

Is OpenFlow better, worse, or the same as proprietary solutions for providing excellent data center operations in this virtualized environment?


 OpenFlow is still in its early stages yet it holds a lot of promise in terms of addressing complexity in many of today's data center networks. The challenges are particularly severe in large scale-out data center networks, such as cloud data centers, where issues such as multi-tenancy, virtual machine and service provisioning, security and traffic isolation are all very real challenges.

Being able to take these challenges, centralize the intelligence and manage the problem space from a single window pane via OpenFlow is particularly attractive. Interestingly, large service providers in the past have had similar network challenges. From complex forwarding and routing, to multi-tenancy, to service provisioning, these problems are not new. Carriers and services providers have traditionally addressed them through provisioning systems, but many of which were proprietary. Today one can think of OpenFlow as doing the same task in large data centers, but with promise of addressing the problems in a standardized manner with broad industry participation.




Dr. Jim Metzler, Moderator, Ashton, Metzler & Associates

In theory if OpenFlow were to be successful, switches and routers could become just relatively dumb forwarding engines and all of the requisite intelligence would reside in a controller. ***What is wrong with that view, or put another way, what intelligence is best left in the switches and routers and why?***

 A more practical approach would be to see what is the problem that is best solved using a centralized controller approach such as using OpenFlow, and then build OpenFlow solutions to address those (i.e. it is best to start with the problem). In today's networks this is even more relevant, where every network is different, with different problem spaces, and as such, the OpenFlow-based solutions for those will also be different.

This was very apparent at the recent Open Networking Summit hosted at Stanford University where different vendors, including Extreme Networks, demonstrated a diverse set of OpenFlow-based network solutions tackling a broad range of challenges based on the problem spaces each have envisioned. Therefore, a blanket approach of moving all intelligence out of the switches and routers may not necessarily be the most pragmatic one. If we once again to take a page out of the service provider's playbook, even though they have built complex provisioning systems which centralized a lot of the intelligence, the switches themselves by no means were dumb forwarding engines. That balance struck between centralized intelligence and distributed processing will also be struck in the data center and with OpenFlow. Some examples of problems which lend themselves to a centralized intelligence controller model include: overlay network provisioning, virtual machine mobility management, and comprehensive resource scheduling.

 While there is a community that believes OpenFlow will lead to the instant commoditization of hardware, we believe they will be sorely disappointed. OpenFlow allows you to do some very cool things in the control plane, but at some point, something still needs to handle the duties of the data plane. Its also important to note that OpenFlow is extensible, so there will continue to feature differentiation in the actual forwarding hardware. Both of those facts add up to there being continued differentiation in data plane hardware. By the way, that is not just our perspective; Martin Casado recently expressed some similar sentiments on his **recent blog**: (Martin's research helped lay the foundation for OpenFlow).

Finally, there is the unpleasant reality that data centers are heterogeneous environments, especially in the enterprise, and switching hardware must have the flexibility and extensibility to handle whatever technologies, old and new, get thrown at it.

From a Cisco perspective, one of the primary benefits of OpenFlow is breaking down the wall between applications and their underlying infrastructure and providing a programmatic interface to data center infrastructure. OpenFlow is put one aspect of the broader concept of software defined networks (SDN) which we believe will help reshape the evolution of networking be providing the programmability/extensibility and closer coupling of application and infrastructure.
