SDN NBI: Crossing the Control-Data Plane Divide?

Lori MacVittie, 2013-24-10

#SDN #NBI #API At what latitude does one start to violate the core control-data plane separation principle of SDN with NBI-based "applications"?

A significant source of chatter around SDN began with a focus on the northbound interface (apparently now dubbed NBI, FWIW) at Layer 123 SDN & OpenWorld Congress 2013. A quick overview of the latest focus in SDN can be read in this SDN Central article, "ONF Will Tackle SDN's Northbound Interface".

Of particular interest was the diagram shared that presents the depth and breadth of possible northbound interfaces ranging from network topology, virtual network management (OpenStack, Pflow VTN, etc..) to application-specific interfaces. That "LB" (load balancing for the uninitiated) is depicted in the same abstraction layer as a firewall and only a step above switching and routing tells me it's a focus on L4 (TCP) load balancing and not more advanced, application-driven load balancing but that's a nit I'll pick on later.

The key takeaway from the chart is that the level of abstraction as you approach application-specificity decreases (of course, necessarily) and the functions capable of being served at that layer are much narrower.

Here's what bugs me, though, in general about this NBI love fest that's about to begin. The closer you get to applications, i.e. every step that takes you above layer 4 (TCP), the closer you get to turning what's supposed to be a separated control path into part of the data path.

The more an SDN application interacts with the traffic comprising a specific application flow the more often it must be involved. At some point, you cross the line between control and data path and make the controller by virtue of its hosting the SDN application, part of the data path.

Consider the simplest case - HTTP. If you want to route HTTP traffic by type (images here, static content there) or if you're a service provider wanting to leverage header enrichment the controller must become an active participant in the data path. Remember, a single TCP **flow** - which is the level SDN proposes to manage (because of limitations on switching models, but that's a post for another day) - can be used to exchange multiple HTTP **messages**. Each of those messages might be a request for a different type of content - image, static HTML, audio, video, etc... If you want to steer (route) each request to a different service (or bypass services that aren't applicable) then you have to examine each request.

In an SDN architecture, that means an application plugged-in to the controller must become part of the data path in order to perform its function as part of the control plane.

That's no different in terms of the performance and scalability (and reliability) implications than sticking an application directly in the data path today. Only the application is in the network data path, which is supposed to be really fast and unimpeded by single points of failure.

That's part of the benefit of the SDN architecture. The separation of control and data planes are supposed to provide a measure of reliability in and of itself. If that "line" is crossed have we not violated one of the basic tenets of SDN architecture?

> The control plane functions include the system configuration, management, and exchange of routing table information. These are performed relatively infrequently. The route controller exchanges the topology information with other routers and constructs a routing table based on a routing protocol, for example, RIP (Routing Information Protocol), OSPF (Open Shortest Path Forwarding), or BGP (Border Gateway Protocol). It can also create a forwarding table for the forwarding engine. Since the control functions are not performed on each arriving individual packet, they do not have a strict speed constraint and are implemented in software in general. [emphasis mine]
>
> — High Performance Switches and Routers (2011) Chao & Liu

If applications at higher latitudes, i.e. tending toward application-specificity, execute control plane functions relatively frequently, they essential become part of the data path.

That's not to say more passive, application-specific applications can't or won't be implemented successfully via the NBI. In fact, they likely will and will provide significant value. But once we cross into active participation in the flow and try to apply application-specific policies and logic to those flows, we may be crossing a line that wasn't designed to be crossed.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com