CHAPTER

HOW SDN WORKS

4

In previous chapters we have seen why SDN is necessary and what preceded the actual advent of SDN in the research and industrial communities. In this chapter we provide an overview of how SDN actually works, including the basic components of a software defined networking system, their roles, and how they interact with one another. In the first part of this chapter we focus on the methods used by Open SDN. We also examine how some *alternate* forms of SDN work. As SDN has gained momentum, some networking vendors have responded with alternate definitions of SDN, which better align with their own product offerings. Some of these methods of implementing SDN-like solutions are new (some are not), and are innovative in their approach. We group the most important of these alternate SDN implementations in two categories, *SDN via APIs* and *SDN via Hypervisor-Based Overlay Networks*, and discuss them separately in the latter half of this chapter.

4.1 FUNDAMENTAL CHARACTERISTICS OF SDN

As introduced in Chapter 3, software defined networking, as it evolved from prior proposals, standards, and implementations such as ForCES, 4D, and Ethane, is characterized by five fundamental traits: *plane separation*, *a simplified device*, *centralized control*, *network automation and virtualization*, and *openness*.

4.1.1 PLANE SEPARATION

The first fundamental characteristic of SDN is the separation of the forwarding and control planes. Forwarding functionality, including the logic and tables for choosing how to deal with incoming packets, based on characteristics such as MAC address, IP address, and VLAN ID, reside in the forwarding plane. The fundamental actions performed by the forwarding plane can be described by how it dispenses with arriving packets. It may *forward*, *drop*, *consume*, or *replicate* an incoming packet. It may also transform the packet in some manner before taking further action. For basic forwarding, the device determines the correct output port by performing a lookup in the address table in the hardware ASIC. A packet may be dropped due to buffer overflow conditions or due to specific *filtering* resulting from a QoS rate-limiting function, for example. Special-case packets that require processing by the control or management planes are consumed and passed to the appropriate plane. Finally, a special case of forwarding pertains to multicast, where the incoming packet must be replicated before forwarding the various copies out different output ports.

Software Defined Networks. http://dx.doi.org/10.1016/B978-0-12-804555-8.00004-1
Copyright © 2017 Elsevier Inc. All rights reserved.

The logic and algorithms that are used to program the forwarding plane reside in the control plane. Many of these protocols and algorithms require global knowledge of the network. The control plane determines how the forwarding tables and logic in the data plane should be programmed or configured. Since in a traditional network each device has its own control plane, the primary task of that control plane is to run routing or switching protocols so that all the distributed forwarding tables on the devices throughout the network stay synchronized. The most basic outcome of this synchronization is the prevention of loops.

While these planes have traditionally been considered logically separate, they co-reside in legacy Internet switches. In SDN, the control plane is moved off of the switching device and onto a centralized controller. This is the inspiration behind Fig. 1.6 in Chapter 1.

4.1.2 SIMPLE DEVICE AND CENTRALIZED CONTROL

Building on the idea of separation of forwarding and control planes, the next characteristic is the simplification of devices, which are then controlled by a centralized system running management and control software. Instead of hundreds of thousands of lines of complicated control plane software running on the device and allowing the device to behave autonomously, that software is removed from the device and placed in a centralized controller. This software-based controller may then manage the network based on higher-level policies. The controller provides primitive instructions to the simplified devices when appropriate in order to allow them to make fast decisions about how to deal with incoming packets.

4.1.3 NETWORK AUTOMATION AND VIRTUALIZATION

Three basic abstractions forming the basis for SDN are defined in [1]. This asserts that SDN can be derived precisely from the abstractions of *distributed state*, *forwarding*, and *configuration*. They are derived from decomposing the actual complex problem of network control faced by networks today into simplifying abstractions. For a historical analogy, note that today's high-level programming languages represent an evolution from their machine language roots through the intermediate stage of languages like C, where today's languages allow great productivity gains by allowing the programmer to simply specify complex actions through programming abstractions. In a similar manner, Shenker [1] purports that SDN is a similar natural evolution for the problem of network control. The distributed state abstraction provides the network programmer with a global network view that shields the programmer from the realities of a network that is actually comprised of many machines, each with its own state, collaborating to solve network-wide problems. The forwarding abstraction allows the programmer to specify the necessary forwarding behaviors without any knowledge of vendor-specific hardware. This implies that whatever language(s) emerges from the abstraction needs to represent a sort of lowest common denominator of forwarding capabilities of network hardware. Finally, the configuration abstraction, which is sometimes called the *specification* abstraction, must be able to express the desired goals of the overall network without getting lost in the details of how the physical network will implement those goals. To return to the programming analogy, consider how unproductive software developers would be if they always needed to be aware of what is actually involved in writing a block of data to a hard disk when they are instead happily productive with the abstraction of file

input and output.[1] Working with the network through this configuration abstraction is really network virtualization at its most basic level. This kind of virtualization lies at the heart of how we will define Open SDN in this work.

The centralized software-based controller in SDN provides an open interface on the controller to allow for automated control of the network. In the context of Open SDN, the terms *northbound* and *southbound* are often used to distinguish whether the interface is to the applications or to the devices. These terms derive from the fact that in most diagrams the applications are depicted above (i.e., to the north) the controller while devices are depicted below (i.e., to the south) the controller. An example of a southbound API is the OpenFlow interface that the controller uses to program the network devices. The controller offers a northbound API, allowing software applications to be plugged into the controller, and thereby allowing that software to provide the algorithms and protocols that can run the network efficiently. These applications can quickly and dynamically make network changes as the need arises. The northbound API of the controller is intended to provide an abstraction of the network devices and topology. There are three key benefits that the application developer should derive from the northbound API: (1) it converts to a syntax that is more familiar to developers (e.g., REST or JSON are more convenient syntaxes than are TLVs); (2) it provides abstraction of the network topology and network layer allowing the application programmer to deal with the network as a whole rather than individual nodes; and (3) it provides abstraction of the network protocols themselves, hiding the application developer from the details of OpenFlow or BGP. In this way, applications can be developed that work over a wide array of manufacturers' equipment that may differ substantially in their implementation details.

One of the results of this level of abstraction is that it provides the ability to virtualize the network, decoupling the network service from the underlying physical network. Those services are still presented to host devices in such a way that those hosts are unaware that the network resources they are using are virtual and not the physical ones for which they were originally designed.

4.1.4 OPENNESS

A characteristic of Open SDN is that its interfaces should remain standard, well documented, and not proprietary. The APIs that are defined should give software sufficient control to experiment with and control various control plane options. The premise is that by keeping open both the northbound and southbound interfaces to the SDN controller, this will allow for research into new and innovative methods of network operation. Research institutions as well as entrepreneurs can take advantage of this capability in order to easily experiment with and test new ideas. Hence, the speed at which network technology is developed and deployed is greatly increased, as a much larger pool of individuals and organizations are able to apply themselves to today's network problems, resulting in better and faster technological advancement in the structure and functioning of networks. The presence of these open interfaces also encourages SDN-related open source projects. As discussed in Section 1.7 and Section 3.6, harnessing the power of the open source development community should greatly accelerate

[1] Granted, programmers who *are* aware of technical details of device driver interactions with hardware usually make better programmers and are better able to optimize performance. While the abstractions discussed here will increase the efficiency of future network programmers, we are not suggesting that deep technical knowledge of how packets move through networking switches and communication links will cease to be needed by elite network programmers.

innovation in SDN [2]. In 2014, Lloyd Carney, CEO of Brocade made the following statement on openness [3]: "*Vertical architectures are giving way to open solutions that allow customers to select best-of-breed technologies that enable new services and networks optimized for the customer's own specific requirements. I can say without hesitation that openness is an integral part of nearly every discussion I have with customers, and I firmly believe it is creating meaningful business opportunities for Brocade.*"

In addition to facilitating research and experimentation, open interfaces permit equipment from different vendors to interoperate. This normally produces a competitive environment which lowers costs to consumers of network equipment. This reduction in network equipment costs has been part of the SDN agenda since its inception.

DISCUSSION QUESTION:

Of the five fundamental traits of SDN we define here, not all are viewed as essential to SDN within the networking industry. Which do you think are more controversial for certain industry players and why?

4.2 SDN OPERATION

At a conceptual level, the behavior and operation of an SDN is straightforward. In Fig. 4.1 we provide a graphical depiction of the operation of the basic components of SDN: the SDN devices, the controller and the applications. The easiest way to understand the operation is to look at it from the bottom up, starting with the SDN device. As can be seen in Fig. 4.1, the SDN devices contain forwarding

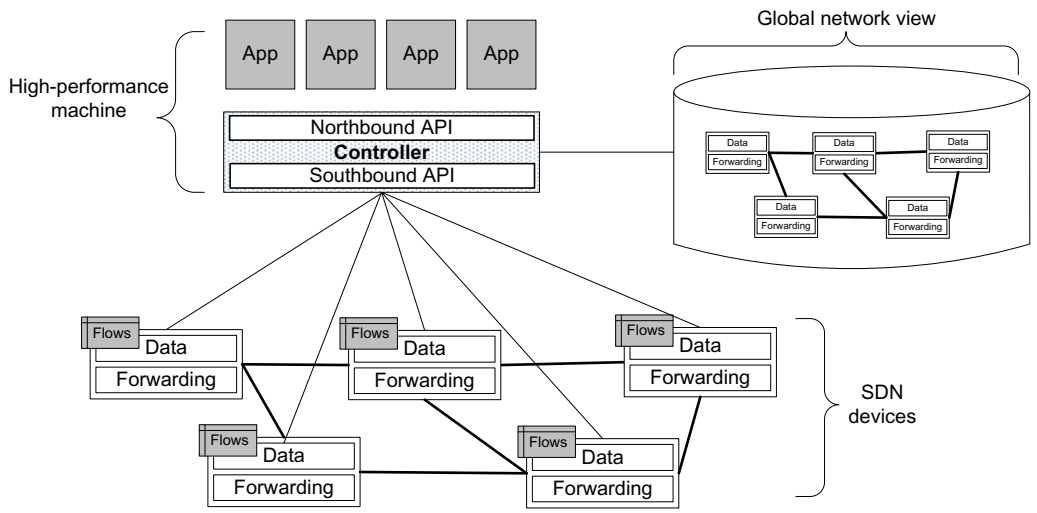


FIG. 4.1

SDN operation overview.

functionality for deciding what to do with each incoming packet. The devices also contain the data which drives those forwarding decisions. The data itself is actually represented by the flows defined by the controller, as depicted in the upper left portion of each device.

A flow describes a set of packets transferred from one network endpoint (or set of endpoints) to another endpoint (or set of endpoints). The endpoints may be defined as IP address-TCP/UDP port pairs, VLAN endpoints, layer three tunnel endpoints, and input port, among other things. One set of rules describes the actions that the device should take for all packets belonging to that flow. A flow is unidirectional in that packets flowing between the same two endpoints in the opposite direction could each constitute a separate flow. Flows are represented on a device as a flow entry.

A flow table resides on the network device and consists of a series of flow entries and the actions to perform when a packet matching that flow arrives at the device. When the SDN device receives a packet it consults its flow tables in search of a match. These flow tables had been constructed previously when the controller downloaded appropriate flow rules to the device. If the SDN device finds a match, it takes the appropriate configured action, which usually entails forwarding the packet. If it does not find a match, the switch can either drop the packet or pass it to the controller, depending on the version of OpenFlow and the configuration of the switch. We will describe flow tables and this packet matching process in greater detail in Section 4.3 and Section 5.3.

The definition of a flow is a relatively simple programming expression of what may be a very complex control plane calculation previously performed by the controller. For the reader less familiar with traditional switching hardware architecture, it is important to understand that this complexity is such that it simply cannot be performed at line rates and instead must be digested by the control plane and reduced to simple rules that can be processed at that speed. In Open SDN, this digested form is the flow entry.

The SDN controller is responsible for abstracting the network of SDN devices it controls and presenting an abstraction of these network resources to the SDN applications running above. The controller allows the SDN application to define flows on devices and to help the application to respond to packets which are forwarded to the controller by the SDN devices. In Fig. 4.1 we see on the right-hand side of the controller that it maintains a view of the entire network that it controls. This permits it to calculate optimal forwarding solutions for the network in a deterministic, predictable manner. Since one controller can control a large number of network devices, these calculations are normally performed on a high performance machine, with an order-of-magnitude performance advantage over the CPU and memory capacity than is typically afforded to the network devices themselves. For example, a controller might be implemented on an eight-core two-GHz CPU versus the single-core one-GHz CPU more typical on a switch.

SDN applications are built on top of the controller. These applications should not be confused with the application layer defined in the seven-layer OSI model of computer networking. Since SDN applications are really part of network layers two and three, this concept is orthogonal to that of applications in the tight hierarchy of OSI protocol layers. The SDN application interfaces with the controller, using it to set *proactive* flows on the devices and to receive packets which have been forwarded to the controller. Proactive flows are established by the application; typically the application will set these flows when the application starts up, and the flows will persist until some configuration change is made. This kind of proactive flow is known as a *static* flow. Another kind of proactive flow is where the controller decides to modify a flow based on the traffic load currently being driven through a network device.

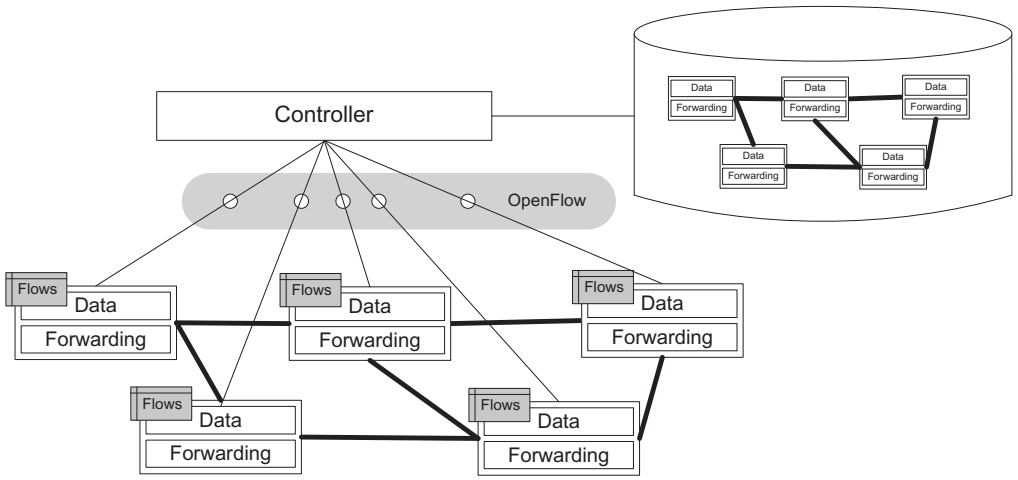


FIG. 4.2

Controller to device communication.

In addition to flows defined proactively by the application, some flows are defined in response to a packet forwarded to the controller. Upon receipt of incoming packets that have been forwarded to the controller, the SDN application will instruct the controller as to how to respond to the packet, and, if appropriate, will establish new flows on the device in order to allow that device to respond locally the next time it sees a packet belonging to that flow. Such flows are called *reactive* flows. In this way, it is now possible to write software applications which implement forwarding, routing, overlay, multipath, and access control functions, among others.

There are also reactive flows that are defined or modified as a result of stimuli from sources others than packets from the controller. For example, the controller can insert flows reactively in response to other data sources such as *Intrusion Detection Systems* (IDS) or the NetFlow traffic analyzer [4].

Fig. 4.2 depicts the OpenFlow protocol as the means of communication between the controller and the device. While OpenFlow is the defined standard for such communication in Open SDN, there are alternative SDN solutions discussed later in this chapter which may use vendor-specific proprietary protocols. The next sections discuss in greater detail SDN devices, controllers and applications.

4.3 SDN DEVICES

An SDN device is composed of an API for communication with the controller, an abstraction layer, and a packet processing function. In the case of a virtual switch, this packet processing function is packet processing software as shown in Fig. 4.3. In the case of a physical switch, the packet processing function is embodied in the hardware for packet processing logic, as shown in Fig. 4.4.

The abstraction layer embodies one or more flow tables, which we discuss in Section 4.3.1. The packet processing logic consists of the mechanisms to take actions based on the results of evaluating

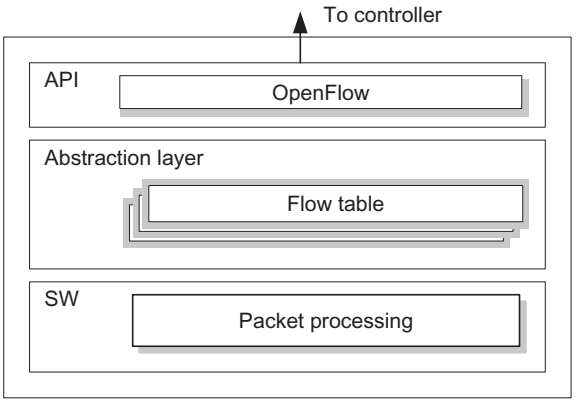


FIG. 4.3

SDN software switch anatomy.

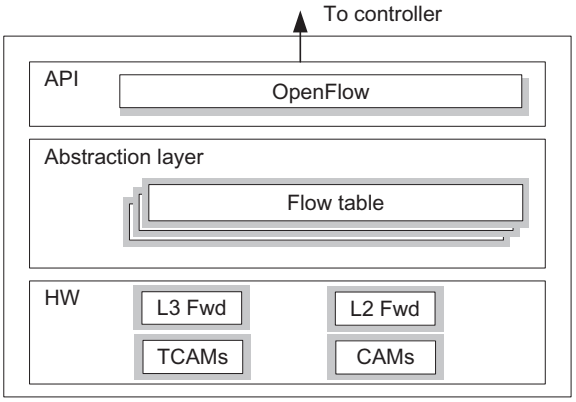


FIG. 4.4

SDN hardware switch anatomy.

incoming packets and finding the highest priority match. When a match is found, the incoming packet is processed locally, unless it is explicitly forwarded to the controller. When no match is found, the packet may be copied to the controller for further processing. This process is also referred to as the controller *consuming* the packet. In the case of a hardware switch, these mechanisms are implemented by the specialized hardware we discuss in Section 4.3.3. In the case of a software switch, these same functions are mirrored by software. Since the case of the software switch is somewhat simpler than the hardware switch, we will present that first in Section 4.3.2. Some readers may be confused by the distinction between a hardware switch and a software switch. The earliest routers that we described in Chapter 1 were indeed just software switches. Later, as we discussed in Section 2.1, we explained

that over time the actual packet forwarding logic migrated into hardware for switches that needed to process packets arriving at ever-increasing line rates. More recently, a role has reemerged in the data center for the pure software switch. Such a switch is implemented as a software application usually running in conjunction with a hypervisor in a data center rack. Like a VM, it may be instantiated or moved under software control. It normally serves as a virtual switch and works collectively with a set of other such virtual switches to constitute a virtual network. We will discuss this concept in greater depth in Chapter 8.

4.3.1 FLOW TABLES

Flow tables are the fundamental data structures in an SDN device. These flow tables allow the device to evaluate incoming packets and take the appropriate action based on the contents of the packet that has just been received. Packets have traditionally been received by networking devices and evaluated based on certain fields. Depending on that evaluation, actions are taken. These actions may include forwarding the packet to a specific port, dropping the packet, and flooding the packet on all ports, among others. An SDN device is not fundamentally different, except that this basic operation has been rendered more generic and more programmable via the flow tables and their associated logic.

Flow tables consist of a number of prioritized flow entries, each of which typically consists of two components, *match fields* and *actions*. Match fields are used to compare against incoming packets. An incoming packet is compared against the match fields in priority order, and the first complete match is selected. Actions are the instructions that the network device should perform if an incoming packet matches the match fields specified for this flow entry.

Match fields can have wildcards for fields that are not relevant to that particular match. For example, when matching packets based just on IP address or subnet, all other fields would be wildcarded. Similarly, if matching on only MAC address or UDP/TCP port, the other fields are irrelevant, and, consequently, those fields are wildcarded. Depending on the application needs, all fields may be important, in which case there would be no wildcards. The flow table and flow entry constructs allow the SDN application developer to have a wide range of possibilities for matching packets and taking appropriate actions.

Given this general description of an SDN device, we now look at two embodiments of an SDN device: first, the more simple software SDN device, and then a hardware SDN device.

4.3.2 SDN SOFTWARE SWITCHES

In Fig. 4.3 we provide a graphical depiction of a purely software-based SDN device. Implementation of SDN devices in software is the simplest means of creating an SDN device because the flow tables, flow entries, and match fields involved are easily mapped to general software data structures, such as sorted arrays and hash tables. Consequently, it is more probable that two software SDN devices produced by different development teams will behave consistently than is the case for two different hardware implementations. Conversely, implementations in software are likely to be slower and less efficient than those implemented in hardware, as they do not benefit from any hardware acceleration. Consequently, for network devices that must run at high speeds, such as 100Gbs, only hardware implementations are feasible with current technology.

Due to the use of wildcards in matching, which poses a problem for typical hash tables, the packet processing function depicted in Fig. 4.3 uses sophisticated software logic in order to implement efficient match field lookups. Hence, in the early days of SDN, there was a wide variance in the performance of different software implementations, based on the efficiency with which these lookups are accomplished. Fortunately, software SDN device implementations have matured. The fact that there are two widely recognized software reference implementations (see Section 4.3.4), both of which use sophisticated and efficient methods of performing these lookups, has resulted in greater uniformity in software SDN device performance.

Software device implementations also suffer less from resource constraints, since considerations such as processing power and memory size are not an issue in typical implementations. Thus, whereas a hardware SDN device implementation will support only a comparatively limited number of flow entries, the ceiling on the number of flow entries on a software device may be orders of magnitude larger. As software device implementations have more flexibility to implement more complex actions, we expect to see a richer set of actions available on software SDN device implementations than on the hardware SDN devices that we examine in the next section.

Software SDN device implementations are most often found in software-based network devices, such as the hypervisors of a virtualization system. These hypervisors often incorporate a software switch implementation which connects the various virtual machines to the virtual network. The virtual switch working with a hypervisor is a natural fit for SDN. In fact, the whole virtualization system is often controlled by a centralized management system, which also meshes well with the centralized controller aspect of the SDN paradigm.

4.3.3 HARDWARE SDN DEVICES

Hardware implementations of SDN devices hold the promise of operating much faster than their software counterparts and, thus, are more applicable to performance-sensitive environments, such as in data centers and network cores. In order to understand how SDN objects such as flow tables and flow entries can be translated into hardware, we will briefly review some of the hardware components of today's networking devices.

Currently, network devices utilize specialized hardware designed to facilitate the inspection of incoming packets and the subsequent decisions that follow, based on the packet matching operation. We see in Fig. 4.4 that the packet processing logic shown in Fig. 4.3 has been replaced by this specialized hardware. This hardware includes the layer two and layer three forwarding tables, often implemented using *Content-Addressable Memories* (CAMs) and *Ternary Content-Addressable Memories* (TCAMs). The price-performance gains enjoyed by RAM technology makes its use more popular for layer two and layer three Forwarding Tables in contemporary designs. The layer three Forwarding Table is used for making IP-level routing decisions. This is the fundamental operation of a router. It matches the destination IP address against entries in the table, and, based on the match, takes the appropriate routing action (e.g., forward the packet out interface B3). The layer two Forwarding Table is used for making MAC-level forwarding decisions. This is the fundamental operation of a switch. It matches the destination MAC address against entries in the table, and, based on the match, takes the appropriate forwarding action (e.g., forward out interface 15).

The layer two forwarding table is typically implemented using regular CAMs or hardware-based hashing. These kinds of associative memories are used when there are precise indices, such as a

forty-eight bit MAC address. TCAMs, however, are associated with more complex matching functions. TCAMs are used in hardware to check not only for an exact match, but also for a third state, which uses a mask to treat certain parts of the match field as wild cards. A straightforward example of this is matching an IP destination address against networks where a *longest prefix match* is performed. Depending on subnet masks, multiple table entries may match the search key and the goal is to determine the closest match. A more important and innovative use of TCAMs is for potentially matching some but not all header fields of an incoming packet. These TCAMs are thus essential for functions such as *policy-based routing* (PBR).

This hardware functionality allows the device to both match packets and then take actions at a very high rate. However, it also presents a series of challenges to the SDN device developer. Specifically:

- How best to translate from flow entries to hardware entries; for example, where and when should the hardware developer utilize CAMs, TCAMs or more straightforward RAM-based hash tables?
- Which of the flow entries to handle in hardware, versus how many to fall back to using software. Most implementations are able to use hardware to handle some of the lookups, but others are handed off to software to be handled there. Obviously, hardware will handle the flow lookups much faster than software, but hardware tables have limitations on the number of flow entries they can hold at any time, and software tables could be used to handle the overflow.
- How to deal with hardware action limitations which may impact whether to implement the flow in hardware versus software; for example, certain actions such as packet modification may be limited or not available if handled in hardware.
- How to track statistics on individual flows. When using devices such as TCAMs which may match multiple flows, it is not possible to use that device to count individual flows separately. Also, gathering statistics across the various tables can be problematic because the tables may count something twice or not at all.

These and other factors will impact the quality, functionality, and efficiency of the SDN device being developed, and must be considered during the design process. For example, hardware table sizes may limit the number of flows, and hardware table capabilities may limit the breadth and depth of special features supported. The limitations presented by a hardware SDN device might require adaptations to SDN applications in order to interoperate with multiple heterogeneous types of SDN devices.

While the challenges to the SDN device designer are no doubt formidable, the range of variables confronting the SDN application developer is vast. The first generation SDN device developer is corralled into basically retrofitting existing hardware to SDN, and thus does not have many choices and indeed may be unable to implement all specified features. The SDN application developer, on the other hand, must deal with inconsistencies across vendor implementations, with scaling performance on a network-wide basis, and a host of other more nebulous issues. We discuss some of these application-specific issues in Section 4.5.

This section provided an overview of the composition of SDN devices and what considerations must be taken into account during their development and their use as part of an SDN application. We provide further specifics on flow tables, flow entries and actions in Chapter 5.

DISCUSSION QUESTION:

Why is the use of devices like TCAMs more challenging in an SDN environment than in a traditional switch?

4.3.4 EXISTING SDN DEVICE IMPLEMENTATIONS

There are a number of SDN device implementations available today, both commercial and open source. Software SDN devices are predominantly open source. Currently, two main alternatives are available, *Open vSwitch* (OVS) [5] from Nicira and *Indigo* [6] from Big Switch. Incumbent *Network Equipment Manufacturers* (NEMs), such as Cisco, HP, NEC, IBM, Juniper, and Extreme, have added OpenFlow support to some of their legacy switches. Generally, these switches may operate in both legacy mode as well as OpenFlow mode. There is also a new class of device called *white box switches*, which are minimalist in that they are built primarily from merchant silicon switching chips and a commodity CPU and memory by a low-cost *Original Device Manufacturer* (ODM) lacking a well-known brand name. One of the premises of SDN is that the physical switching infrastructure may be built from OpenFlow-enabled white box switches at far less direct cost than switches from established NEMs. Most legacy control plane software is absent from these devices, since this is largely expected to be provided by a centralized controller. Such white box devices often use the open source OVS or Indigo switch code for the OpenFlow logic, and then map the packet processing part of those switch implementations to their particular hardware. Chuck Robbins, Cisco CEO, recognizes the impact of SDN when he states "*Everything we build will be programmable*" [7].

4.3.5 SCALING THE NUMBER OF FLOWS

The granularity of flow definitions will generally be more fine as the device holding them approaches the edge of the network and more general the more the device approaches the core. At the edge, flows will permit different policies to be applied to individual users and even different traffic types of the same user. This will imply in some cases multiple flow entries for a single user. This level of flow granularity would simply not scale if it were applied closer to the network core, where large switches deal with the traffic for tens of thousands of users simultaneously. In those core devices, the flow definitions will be generally more coarse, with a singe aggregated flow entry matching the traffic from a large number of users whose traffic is aggregated in some way such as a tunnel, a VLAN or a MPLS LSP. Policies applied deep into the network will likely not be user-centric policies but policies that apply to these aggregated flows. One positive result of this is that there will not be an explosion of the number of flow entries in the core switches due to handling the traffic emanating from thousands of flows in edge switches.

4.4 SDN CONTROLLER

We have noted that the controller maintains a view of the entire network, implements policy decisions, controls all the SDN devices that comprise the network infrastructure, and provides a northbound API for applications. When we have said that the controller implements policy decisions regarding routing, forwarding, redirecting, load balancing, and the like, these statements referred both to the controller and to the applications that make use of that controller. Controllers often come with their own set of common application modules, such as a learning switch, a router, a basic firewall, and a simple load balancer. These are really SDN applications, but they are often bundled with the controller. Here we focus strictly on the controller.

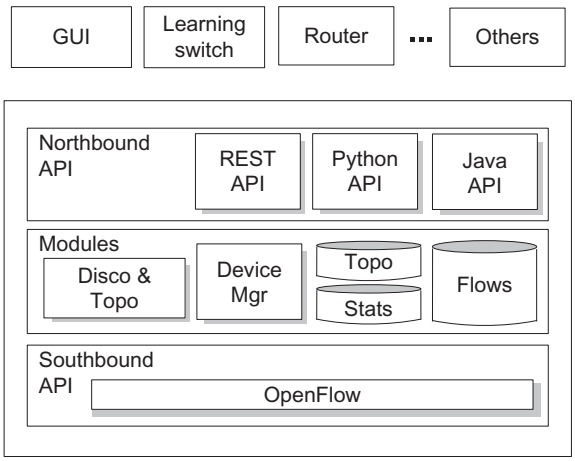


FIG. 4.5

SDN controller anatomy.

Fig. 4.5 exposes the anatomy of an SDN controller. The figure depicts the modules that provide the core functionality of the controller, both a northbound and southbound API, and a few sample applications that might use the controller. As we described earlier, the southbound API is used to interface with the SDN devices. This API is OpenFlow in the case of Open SDN or some alternative such as BGP in other SDN solutions. It is worth noting that in some product offerings both OpenFlow and alternatives coexist on the same controller. Early work on the southbound API has resulted in more maturity of that interface with respect to its definition and standardization. OpenFlow itself is the best example of this, but de facto standards like the Cisco CLI and SNMP also represent standardization in the southbound-facing interface. OpenFlow's companion protocol OF-Config [8] and Nicira's *Open vSwitch Database Management Protocol* (OVSDB) [9] are both open protocols for the southbound interface, though these are limited to configuration roles.

Currently, the Open Source SDN (http://www.opensourcesdn.org) group is proposing a northbound counterpart to the southbound OpenFlow standard. While the lack of a standard for the controller-to-application interface is considered a current deficiency in SDN, organizations like the OpenSource SDN group are developing proposals to standardize this. The absence of a standard notwithstanding, northbound interfaces have been implemented in a number of disparate forms. For example, the Floodlight controller [10] includes a Java API, and a *Representational State Transfer* (RESTful) [11] API. The OpenDaylight controller [12] provides a RESTful API [13] for applications running on separate machines. The northbound API represents an outstanding opportunity for innovation and collaboration amongst vendors and the open source community.

4.4.1 SDN CONTROLLER CORE MODULES

The controller abstracts the details of the SDN controller-to-device protocol so that the applications above are able to communicate with those SDN devices without knowing the nuances of those devices.

Fig. 4.5 shows the API below the controller, which is OpenFlow in Open SDN, and the interface provided for applications. Every controller provides core functionality between these raw interfaces. Core features in the controller will include:

- **End-user Device Discovery**: Discovery of end-user devices, such as laptops, desktops, printers, mobile devices, etc.
- **Network Device Discovery**: Discovery of network devices which comprise the infrastructure of the network, such as switches, routers, and wireless access points.
- **Network Device Topology Management**: Maintain information about the interconnection details of the network devices to each other, and to the end-user devices to which they are directly attached.
- **Flow Management**: Maintain a database of the flows being managed by the controller and perform all necessary coordination with the devices to ensure synchronization of the device flow entries with that database.

The core functions of the controller are device and topology discovery and tracking, flow management, device management and statistics tracking. These are all implemented by a set of modules internal to the controller. As shown in Fig. 4.5, these modules need to maintain local databases containing the current topology and statistics. The controller tracks the topology by learning of the existence of switches (SDN devices) and end-user devices and tracking the connectivity between them. It maintains a *flow cache* which mirrors the flow tables on the different switches it controls. The controller locally maintains per-flow statistics that it has gathered from its switches. The controller may be designed such that functions are implemented via pluggable modules such that the feature set of the controller may be tailored to an individual network's requirements.

Many companies implementing SDN look to these core modules to help them model their network and create abstraction layers. Level 3's CTO Jack Waters described his company's efforts with SDN: "*The combined company (Level 3 and Time Warner Telecom) has done a number of network trials. With respect to SDN, we're in the development phase around a few NFV betas so it's early but we have use cases up and running already. The combined company is laser focused on how to leverage the technology really to help build a network abstraction layer for things like provisioning, configuration management, and provisioning automation is something we think will be great for us and our customers.*" [14]

4.4.2 SDN CONTROLLER INTERFACES

For SDN applications, a key function provided by the SDN controller is the API for accessing the network. In some cases, this northbound API is a low level interface, providing access to the network devices in a common and consistent manner. In this case, that application is aware of individual devices, but is shielded from their differences. In other instances the controller may provide high level APIs that provide an abstraction of the network itself, so that the application developer need not be concerned with individual devices, but with the network as a whole.

Fig. 4.6 takes a closer look at how the controller interfaces with applications. The controller informs the application of *events* that occur in the network. Events are communicated from the controller to the application. Events may pertain to an individual packet that has been received by the controller or some state change in the network topology such as a link going down. Applications use different *methods* to

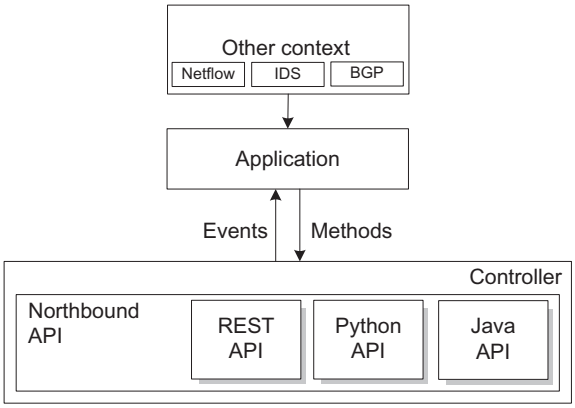


FIG. 4.6

SDN controller northbound API.

affect the operation of the network. Such methods may be invoked in response to a received event and may result in a received packet being dropped, modified and/or forwarded, or the addition, deletion or modification of a flow. The applications may also invoke methods independently without the stimulus of an event from the controller, as we explain in Section 4.5.1. Such inputs are represented by the *other context* box in Fig. 4.6.

As of this writing, a search of the web reveals more than twenty SDN controllers available on the market. Interestingly, there are no standards for the northbound API, also called the *northbound interface* (NBI). Why is the standardization of the NBI important? From a customer perspective, a standard NBI would allow for the SDN application market to materialize. SDN application developers no longer would have to modify their application to multiple proprietary interfaces across SDN controllers. If one acquires an SDN application, the ability to plug-n-play different controllers into the environment without big changes to code would become a reality. In addition, investments made by the enterprise or carrier in creating custom SDN applications to differentiate their own end services would be more flexible because of the ability to plug-n-play SDN controllers.

4.4.3 EXISTING SDN CONTROLLER IMPLEMENTATIONS

There are a number of implementations of SDN controllers available on the market today. They include both open source SDN controllers and commercial SDN controllers. Open source SDN controllers come in many forms, from basic C-language controllers, such as NOX [15], to Java-based versions, such as Beacon [16] and Floodlight [10]. There is even a Ruby-based [17] controller called Trema [18]. Interfaces to these controllers may be offered in the language in which the controller is written or other alternatives, such as REST or Python. The open source controller OpenDaylight (ODL) [12] has been built by a consortium of vendors. An important open source controller, *Open Network Operating System* (ONOS) [19] is rapidly gaining traction in the carrier market. Other vendors offer their own commercial version of an SDN controller. Vendors such as NEC, IBM, and HP offer controllers that are primarily OpenFlow implementations. Most other NEMs offer vendor-specific and proprietary SDN

controllers that include some level of OpenFlow support, though these controllers vary considerably in terms of which version of the OpenFlow specification they support. In general, these commercial controllers do not support the latest version of the specification, which evolves more quickly than the NEMs' ability and motivation to keep up. As of this writing, ODL and ONOS have assumed dominant roles for commercial applications. For this reason, we provide deeper treatment of these two controllers in Section 7.3.

There are pros and cons to the proprietary alternative controllers. Although proprietary controllers are more closed than the nominally open systems, they do offer some of the automation and programmability advantages of SDN while at the same time providing a *buck stops here* level of support for the network equipment. They permit SDN-like operation of legacy switches, obviating the need to replace older switching equipment in order to begin the migration to SDN. They do constitute closed systems, however, which ostensibly violates one of the original tenets of SDN. They also may do little to offload control functionality from devices, resulting in the continued high cost of network devices. These proprietary alternative controllers are generally a component of the alternative SDN methodologies we introduce in Section 4.6.

4.4.4 POTENTIAL ISSUES WITH THE SDN CONTROLLER

In general the Open SDN controller suffers from the birthing pains common to any new technology. While many important problems are addressed by the concept and architecture of the controller, there have been comparatively few large-scale commercial deployments thus far. As more commercial deployments scale, more real-life experience in large, demanding networks will be needed. In particular, experience with a wider array of applications with a more heterogeneous mix of equipment types is needed before widespread confidence in this architecture is established. Achieving success in these varied deployments will require that a number of potential controller issues be adequately addressed. In some cases, these solutions will come in multiple forms from different vendors. In other cases, a standards body such as the ONF will have to mandate a standard. In Section 3.2.6 we stated that a centralized control architecture is needed to grapple with the issues of latency, scale, high availability, and security. In addition to these more general issues, the centralized SDN controller will need to confront the challenges of *coordination between applications*, *the lack of a standard northbound API*, and *flow prioritization*.

Scale will be a very important aspect of SDN efforts at large carriers like AT&T and Verizon. John Donovan, Chief Infrastructure Executive at AT&T recently said [20]: "*No part of our network will be unaffected. We've catalogued the hundreds of network functions that we manage, and decided which will be relevant in the future and which are becoming obsolete. Of those 200 critical future functions, we plan to move 75% to this new software defined architecture in the next five years. The 5% of network functions that we plan to virtualize in 2015 includes elements of our mobile packet core, the heart of our mobile network.*" AT&T is not alone in its challenge to scale. In November of 2014, Vodafone Australia's CTO indicated that the company's goal is to virtualize most of its core network within five years [21].

DISCUSSION QUESTION:

What might be some solutions for dealing with issues of scale in an Open SDN environment?

There may be more than one SDN application which is running on a single controller. When this is the case, issues related to application prioritization and flow handling become important. Which application should receive an event first? Should the application be required to pass along this event to the next application in line, or can it deem the processing complete, in which case no other applications get a chance to examine and act on the received event?

An emerging standard for the northbound API holds promise for efforts to develop applications that will be reusable across a wide range of controllers. Early standardization efforts for OpenFlow generally assumed such a northbound counterpart would emerge, and much of the efficiency gain assumed to come from a migration to SDN will be lost without it. Late in 2013 the ONF formed a workgroup that focuses on the standardization of the northbound API. Today, in collaboration with OpenSource SDN, that effort has produced an *intents-based* NBI, called *Boulder*. Boulder has created the semantics and information models that support applications in communicating the network changes while abstracting away the underlying implementation [22]. The hope is that significant momentum will be gained in SDN via the adoption of this unifying NBI by multiple SDN controllers.

The reader should note that there is not universal agreement that a single, standard NBI is always needed across all the mainstream controllers. An alternative view reveals an ecosystem growing around each controller. This is reminiscent of how Microsoft Windows and Apple OSX each have their own ecosystems with some applications remaining unique to each platform while other applications are ported across multiple platforms. Moreover, an application written inside a particular controller will likely retain a performance advantage over applications written northbound of that controller.

Flows in an SDN device are processed in priority order. The first flow which matches the incoming packet is acted upon. Within a single SDN application, it is critical for the flows on the SDN device to be prioritized correctly. If they are not, the resulting behavior will be incorrect. For example, the designer of an application will put more specific flows at a higher priority (e.g., match all packets from IP address 10.10.10.2 and TCP port 80), and the most general flows at the lowest priority (e.g., match everything else). This is relatively easy to do for a single application. However, when there are multiple SDN applications, flow entry prioritization becomes more difficult to manage. How does the controller appropriately interleave the flows from all applications? This is a challenge and requires special coordination between the applications.

4.5 SDN APPLICATIONS

SDN applications run above the SDN controller, interfacing to the network via the controller's northbound API. SDN applications are ultimately responsible for managing the flow entries that are programmed on the network devices, using the controller's API to manage flows. Through this API the applications are able to: (1) configure the flows to route packets through the best path between two endpoints; (2) balance traffic loads across multiple paths or destined to a set of endpoints; (3) react to changes in the network topology such as link failures and the addition of new devices and paths, and (4) redirect traffic for purposes of inspection, authentication, segregation, and similar security-related tasks.

Fig. 4.5 includes some standard applications, such as a GUI for managing the controller, a learning switch and a routing application. The reader should note that even the basic functionality of a simple layer two learning switch is not obtained by simply pairing an SDN device with an SDN controller. Additional logic is necessary to react to the newly seen MAC address and update the forwarding tables in the SDN devices being controlled in such a way as to provide connectivity to that new MAC

address throughout the network while avoiding switching loops. This additional logic is embodied in the learning switch application in Fig. 4.5. One of the perceived strengths of the SDN architecture is the fact that switching decisions can be controlled by an ever richer family of applications that control the controller. In this way, the power of the SDN architecture is highly expandable. Other applications that are well-suited to this architecture are load balancers and firewalls, among many others.

These examples represent some typical SDN applications which have been developed by researchers and vendors today. Applications such as these demonstrate the promise of SDN—being able to take complex functionality that formerly resided in each individual network device or appliance, and allowing it to operate in an Open SDN environment.

4.5.1 SDN APPLICATION RESPONSIBILITIES

The general responsibility of an SDN application is to perform whatever function for which it was designed, be it load-balancing, firewalling, or some other operation. Once the controller has finished initializing devices and has reported the network topology to the application, the application spends most of its processing time responding to events. While the core functionality of the application will vary from one application to another, application behavior is driven by events coming from the controller as well as external inputs. External inputs could include network monitoring systems, such as Netflow, IDS, or BGP peers. The application affects the network by responding to the events as modeled in Fig. 4.6. The SDN application registers as a *listener* for certain events, and the controller will invoke the application's callback *method* whenever such an event occurs. This invocation will be accompanied by the appropriate details related to the event. Some examples of events handled by an SDN application are: *End-user Device Discovery*, *Network Device Discovery*, and *Incoming Packet*. In the first two cases, events are sent to the SDN application upon the discovery of a new end-user device (i.e., MAC address) or a new network device (e.g., switch, router, or wireless access point), respectively. Incoming Packet events are sent to the SDN application when a packet is received from an SDN device, due to either a flow entry instructing the SDN device to forward the packet to the controller, or because there is no matching flow entry at the SDN device. When there is no matching flow entry the default action is usually to forward the packet to the controller though it could be to drop the packet, depending on the nature of the applications.

There are many ways in which an SDN application can respond to events that have been received from the SDN controller. There are simple responses, such as downloading a set of default flow entries to a newly discovered device. These default or *static* flows typically are the same for every class of discovered network device, and, hence, little processing is required by the application. There are also more complex responses which may require state information gathered from some other source apart from the controller.

This can result in variable responses depending on that state information. For example, based on a user's state, an SDN application may decide to process the current packet in a certain manner, or else it may take some other action, such as downloading a set of user-specific flows.

4.6 ALTERNATE SDN METHODS

Thus far in this chapter we have been examining what we consider the original definition of SDN, which we distinguish from other alternatives by the term *Open* SDN. Open SDN most certainly has the

broadest support in the research community and amongst the operators of large data centers such as Google, Yahoo! and their ilk. Nevertheless, there are other proposed methods of accomplishing at least some of the goals of SDN. These methods are generally associated with a single networking vendor or a consortium of vendors. We define here two alternate categories of SDN implementations, *SDN via APIs* and *SDN via Hypervisor-Based Overlay Networks*. The first of these consists of employing functions that exist on networking devices that can be invoked remotely, typically via traditional methods, such as SNMP, CLI, or NETCONF. Alternatively, newer mechanisms such as RESTful APIs may be used. In SDN via Hypervisor-Based Overlay Networks the details of the underlying network infrastructure are not relevant. Virtualized *overlay* networks are instantiated across the top of the physical network. We make the distinction that the overlays be hypervisor-based since network overlays exist in other, nonhypervisor-based forms as well. One early example of this sort of network virtualization is VLAN technology. Another type of overlay network that is not related to our use of the term are *P2P/Overlay* networks such as Napster and BitTorrent. We will further clarify the distinction between SDN and P2P/Overlay networks later in Section 9.7.

In the following sections we provide an introduction of how these two alternate SDN methods work, and defer a more detailed treatment of these as well as the introduction of some others to Chapter 6.

4.6.1 SDN VIA APIs

If a basic concept of SDN is to move control functionality from devices to a centralized controller, then this can be accomplished in other ways than the OpenFlow-centric approach coupled to Open SDN. Looking at the value provided by SDN as described in earlier chapters, a couple of key functional improvements are *programmability* and *centralized control*. Programmability provides the opportunity for applications to be written in order to automate processes and to provide a level of dynamic control which goes beyond what is available today with static, CLI-based networking updates. Centralized control provides the ability for applications running on a controller or independently to have views of the entire network, and hence increases the ability to make optimal policy-based forwarding and routing decisions.

Fig. 4.7 depicts a number of levels at which APIs can be provided for SDN applications. These levels are *Device*, *Controller*, and *Policy*, for which we provide detail in the following sections.

SDN via Device APIs

In *SDN via Device APIs*, the strategy is to provide a richer set of control points on the devices, so that centrally located software can manipulate those devices and provide the intelligent and predictable behavior that is expected in an SDN-controlled network. Consequently, many vendors offer SDN solutions by improving the means of affecting configuration changes on their network devices.

We depict *SDN via Device APIs* graphically in Fig. 4.8. The diagram shows a set of centralized applications communicating with devices via device-level APIs. Often with SDN via Device APIs solutions, vendors provide an enhanced level of APIs on their devices rather than just the traditional CLI and SNMP. The architecture shown in the figure is reminiscent of the earlier diagrams in this chapter. The most obvious difference with the earlier diagrams is that we here depict the controller as optional. This is because even when the controller is present in this paradigm, it adds little or no value to the communication between the application and device. That is, it serves as a pass-through

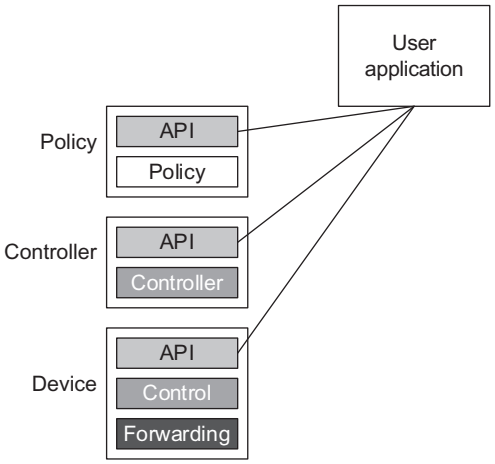


FIG. 4.7

SDN via different levels of APIs.

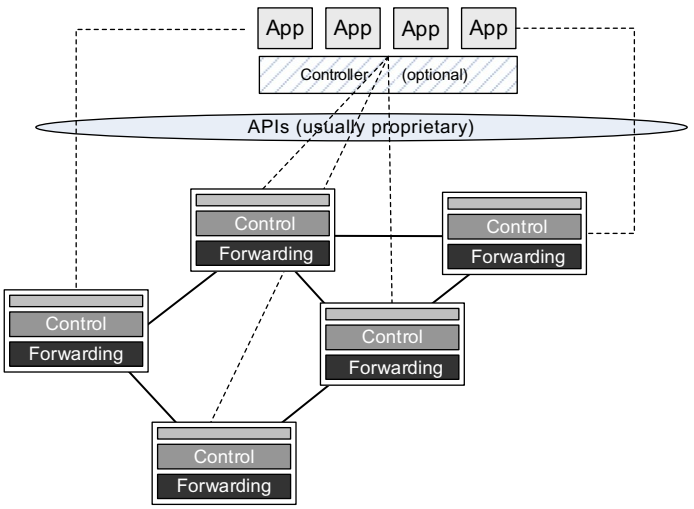


FIG. 4.8

SDN via Device APIs.

only. Indeed, the applications may sometimes communicate with the devices directly without using a controller at all.

Since the early days of the commercial Internet, it has been possible to set configuration parameters on devices using methods such as CLI and SNMP. While these mechanisms have long been available, they can be cumbersome and difficult to maintain. Furthermore, they are geared toward relatively rare

static management tasks, not the dynamic, frequent, and automated tasks required in environments such as today's data centers. Newer methods of providing the means to make remote configuration changes have been developed in the last few years.

One such approach used by some vendors is the use of RESTful APIs. REST has become the dominant method of making API calls across networks for computational tasks. REST uses *HyperText Transfer Protocol* (HTTP), the protocol commonly used to pass web traffic. RESTful APIs are simple and extensible and have the advantage of using a standard TCP port and, thus, require no special firewall configuration to permit the API calls to pass through firewalls. We provide a more detailed discussion of RESTful APIs in Section 6.2.4.

Another related API is the IETF's *Interface to the Routing System* (I2RS) [23] I2RS provides an interface between routing protocols and the *Routing Information Base* (RIB). The four key drivers inspiring this new protocol are:

- The need for an interface that is programmatic, asynchronous, and offers fast, interactive access for atomic operations.
- To provide access to structured routing information and state that is frequently not directly configurable or modeled in existing implementations or configuration protocols.
- To provide the ability for network management and other applications to subscribe to structured, filterable event notifications from the routing system.
- To facilitate extensibility and provide standard data-models to be used by network applications.

While I2RS seems to be very synergistic with SDN, it remains to be seen whether this new architecture will gain traction in commercial use.

Note in the figure that communication to devices using these APIs (rather than OpenFlow) can be done either directly, or through a controller. In either case, the solution is using improved APIs which provide the ability to control the forwarding plane through mechanisms made available on the device by the manufacturer.

SDN via Controller APIs

In order to promote programmability from a centralized location, SDN via Controller-level APIs can provide a platform from which to build SDN applications. Solutions using this type of SDN utilize the set of APIs provided by the controller. These APIs are open and available for application developers. Since there is no standard NBI, applications written for one controller's set of APIs may not run on a different controller.

Fig. 4.9 depicts the SDN via Controller APIs model. The reader should note that this model does not appear to contrast fundamentally from those we presented earlier in Section 4.4 in our discussion about Open SDN. The fundamental distinction comes in the nature of the southbound APIs the controller uses. In the case of Open SDN, the southbound protocol is OpenFlow. In the case of SDN via Controller APIs, the southbound protocol consists of one or more legacy protocols that are being re-purposed to provide SDN-like functionality via the centralized controller.

One goal of implementing APIs at the controller level is to provide a level of abstraction between the devices and the application, abstracting details of each specific device and each protocol, and making it possible for the application to interact with the controller at a higher level. One example of this is the case of gathering topology information about devices in the network and presenting

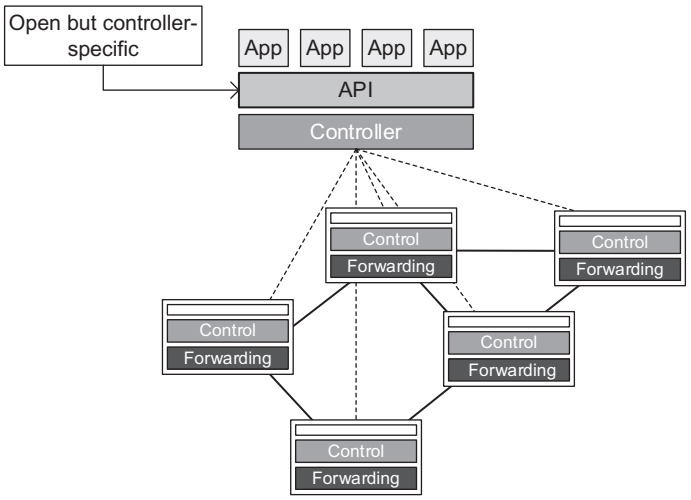


FIG. 4.9

SDN via Controller APIs.

it via the controller API. In this case, the application benefits from the work done by the controller and is able to read information about nodes and links from the controller without having to perform the actual device and host discovery and topology. This is precisely how existing routing protocols such as *Border Gateway Protocol* (BGP), *Border Gateway Protocol Link State* (BGP-LS), and *Path Computation Element Protocol* (PCE-P) are being used today. Implementations of SDN using these device protocols make use of technology existing in devices in order to provide the ability to create SDN applications which control routes and paths throughout the network. The use of these protocols with SDN will be described in detail in Section 6.2.3.

In other cases, the controller performs only light processing of the messages between the application and device. An example of this is found in one of the typical uses of NETCONF as an SDN API. NETCONF is one of the device management and configuration protocols most favored by vendors and customers today. It is used especially in higher-level devices such as core routers. These devices expose their configurable information via data models defined using the YANG data definition language. (NETCONF and YANG will be discussed in greater detail in Section 6.2.2.) In our context of SDN via Controller APIs, the application typically makes RESTCONF requests to the controller. RESTCONF is a protocol wherein NETCONF requests are bundled in easier-to-use REST messages. The controller in turn translates the RESTCONF request into the NETCONF protocol, with the contents of the request remaining fundamentally the same.

One SDN controller that fits squarely in the *SDN via Controller APIs* category is ODL. A prime feature of ODL is its ability to support multiple southbound protocols. Included in the long list of supported protocols are OpenFlow (giving support for Open SDN), NETCONF (supporting SDN via APIs), and BGP-LS/PCE-P (supporting SDN via APIs). Fig. 4.10 shows ODL support for multiple southbound protocol plugins. ODL evangelists argue that northbound APIs can be developed that would allow applications to manage hybrid sets of devices, some managed on the southbound side

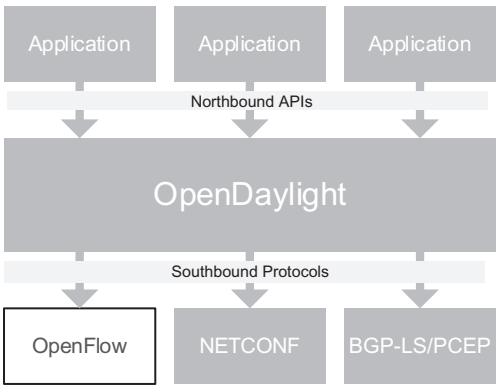


FIG. 4.10

OpenDaylight controller APIs.

via OpenFlow and some managed by legacy APIs. While this would not afford the fine-grained control over all flows afforded by a pure OpenFlow solution, it offers some intriguing possibilities for intelligent control of hybrid networks. Fig. 4.10 shows that the distinction between an Open SDN controller as described earlier in this chapter and SDN via Controller APIs can become blurred.

SDN via Controller level APIs remains a popular choice for SDN application developers, since it provides SDN functionality in the form of programmability and centralized control, as well as supporting legacy devices.

SDN via Policy-level APIs

Another approach is to provide a level of APIs which reside at a layer above the controller level. These APIs are created at a level of abstraction such that they address *policies*, rather than merely individual device or network capabilities. These APIs come in different flavors and may be directed at different target domains, but they all attempt to address network configuration from a *declarative* perspective, rather than an *imperative* one.

These concepts are defined as follows:

- *Imperative:* Imperative systems and APIs require the user to input exactly *how* to do a particular task.
- *Declarative:* Declarative systems and APIs request the user to input exactly *what* is to be accomplished. It is the responsibility of the system to determine *how* to do it.

Fig. 4.11 shows SDN via Policy APIs; notice the policy layer running on top of the controller, providing functionality and a level of abstraction which shields the SDN application from the nuances and details of interacting with a specific device.

These policy-based API solutions are gaining traction not only in the SDN via APIs category, but also in the Open SDN category as well. An example of this type of API would be the NBI concept of *intents*, which raises the API level of abstraction such that requests are truly declarative. The notion of

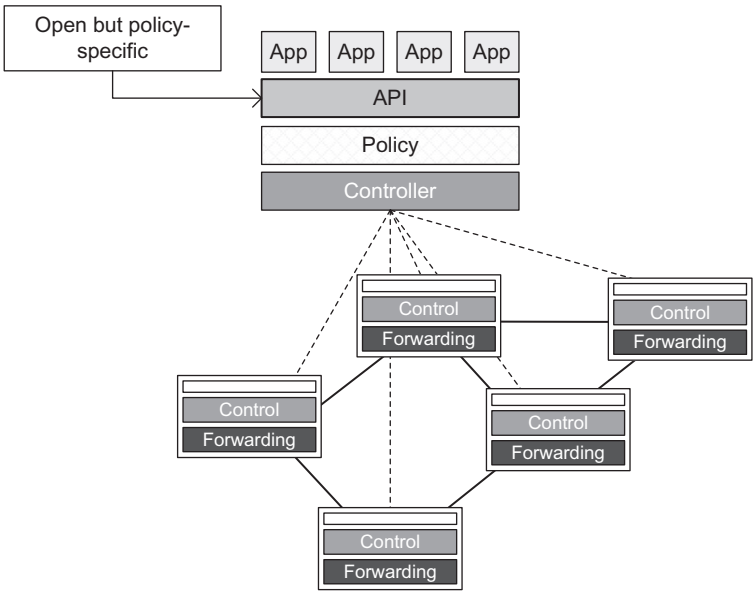


FIG. 4.11

SDN via Policy APIs.

controlling network behavior via policy is an active area of research in SDN. We cover relevant policy languages in Section 7.6.

4.6.2 BENEFITS AND LIMITATIONS OF SDN VIA APIS

There are a number of benefits of SDN via APIs. One distinct advantage of this approach is that, because it uses legacy management interfaces, it therefore works with legacy switches. Thus, this solution does not require upgrading to OpenFlow-enabled switches. Another benefit of this approach is that it allows for some improvement in agility and automation. These APIs also make it easier to write software such as orchestration tools which can respond quickly and automatically to changes in the network (e.g., the movement of a virtual machine in a data center). A third advantage is that it allows for some amount of centralized control of the devices in the network. Therefore, it is possible to build an SDN solution using the provided APIs on the distributed network devices. Finally, there is potential for increased openness in the SDN via APIs approach. Although the individual interfaces may be proprietary to individual vendors, when they are exposed to the applications, they are made open for exploitation by applications. The degree of openness will vary from one NEM to another.

Of course, the API-based SDN methods have their limitations. First, in some cases *there is no controller at all*. The network programmer needs to interact directly with each switch. Second, even when there is a controller, it may not provide an abstract, *network-wide* view to the programmer. This is certainly true when dealing with NETCONF-based APIs on the controller. In these situations, the programmer needs to think in terms of individual switches. (Note that it is possible to develop a layer above NETCONF that abstracts away device details, so it would not be fair to conclude that

the use of NETCONF or other similar APIs precludes an abstraction layer above.) Third, since there is still a control plane operating on each switch, the controller and, more importantly, the programmer developing applications on top of that controller must synchronize with what the distributed control plane is doing. Another drawback is that the solution is often proprietary. When these APIs are nonstandard (as opposed to a protocol such as OpenFlow), SDN-like software applications using this type of API-based approach will only work with devices from that specific vendor, or a small group of compatible vendors. This limitation is sometimes circumvented by extending this approach to provide support for multiple vendors' APIs. This masks the differences between the device APIs to the application developer, who will see a single northbound API despite the incompatible device interfaces on the southbound side. Obviously, this homogeneity on the northbound interface is achieved by increased complexity within the controller.

DISCUSSION QUESTION:

What do you see as the major advantages of the SDN via APIs solution? Do you see these as advantages going forward into the future of networking, or are they short-term benefits?

4.6.3 SDN VIA HYPERVISOR-BASED OVERLAY NETWORKS

Another more innovative alternate SDN method is what we refer to as *Hypervisor-Based Overlay* networks. Under this concept the current physical network is left as it is, with networking devices and their configurations remaining unchanged. Above that network, however, *hypervisor-based virtualized* networks are erected. The systems at the edges of the network interact with these virtual networks, which obscure the details of the physical network from the devices that connect to the overlays.

We depict such an arrangement in Fig. 4.12, where we see the virtualized networks *overlaying* the physical network infrastructure. The SDN applications making use of these overlay network resources

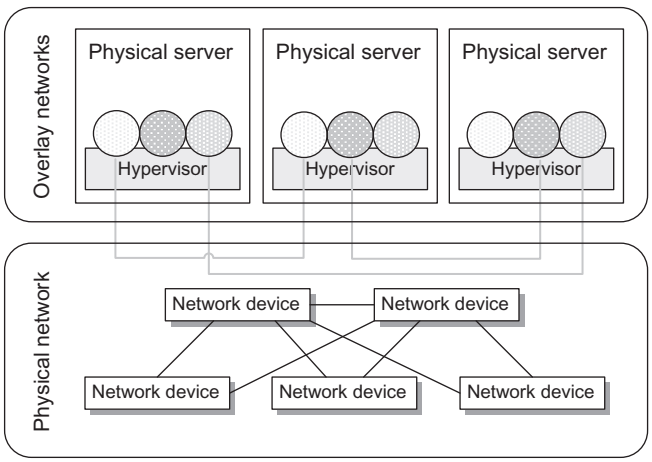


FIG. 4.12

Virtualized networks.

are given access to virtualized networks and ports, which are abstract in nature and do not necessarily relate directly to their physical counterparts below.

As shown in Fig. 4.12, conceptually the virtual network traffic runs *above* the physical network infrastructure. The hypervisors inject traffic into the virtual network and receive traffic from it. The traffic of the virtual networks is passed through those physical devices, but the endpoints are unaware of the details of the physical topology, how routing occurs or other basic network functions. Since these virtual networks exist above the physical infrastructure, they can be controlled entirely by the devices at the very edge of the network. In data centers, these would typically be the hypervisors of the VMs that are running on each server.

The mechanism that makes this possible is tunneling, which uses encapsulation. When a packet enters the edge of the virtual network at the source, the networking device (usually the hypervisor) will take the packet in its entirety and encapsulate it within another frame. This is shown in Fig. 4.13. Note that the edge of the virtual network is called a *Tunnel Endpoint* or *Virtual Tunnel Endpoint* (VTEP).

The hypervisor then takes this encapsulated packet and, based on information programmed by the controller, sends it to the destination's VTEP. This VTEP decapsulates the packet and forwards it to the destination host. As the encapsulated packet is sent across the physical infrastructure, it is being sent from the source's VTEP to the destination's VTEP. Consequently, the IP addresses are those of the source and destination VTEP. Normally, in network virtualization, the VTEPs are associated with hypervisors.

This tunneling mechanism is referred to as *MAC-in-IP* tunneling, because the entire frame, from MAC address inwards, is encapsulated within this unicast IP frame, as shown in Fig. 4.13. Different vendors have established their own methods for MAC-in-IP tunneling. Specifically, Cisco offers VXLAN [24], Microsoft uses NVGRE [25], and Nicira's is called STT [26].

This approach mandates that a centralized controller be in charge of making sure there is always a mapping from the actual destination host to the destination hypervisor which serves that host.

Fig. 4.14 shows the roles of these VTEPs as they serve the source and destination host devices. The virtual network capability is typically added to a hypervisor by extending it with a virtual switch. We introduced the notion of virtual switch in Section 4.3.2, and it is well suited to the overlay network concept. The virtual network has a virtual topology, consisting of the virtual switches interconnected by virtual point-to-point links. The virtual switches are depicted as the VTEPs in Fig. 4.14 and the

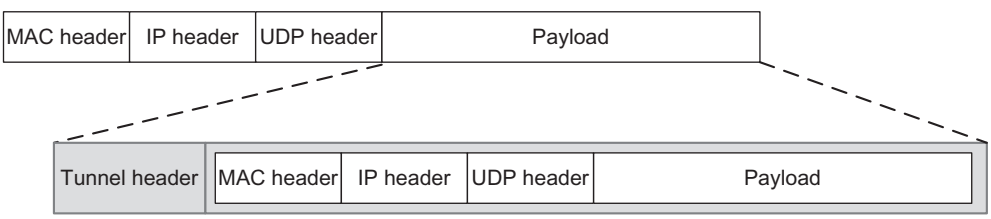


FIG. 4.13

Encapsulated frames.

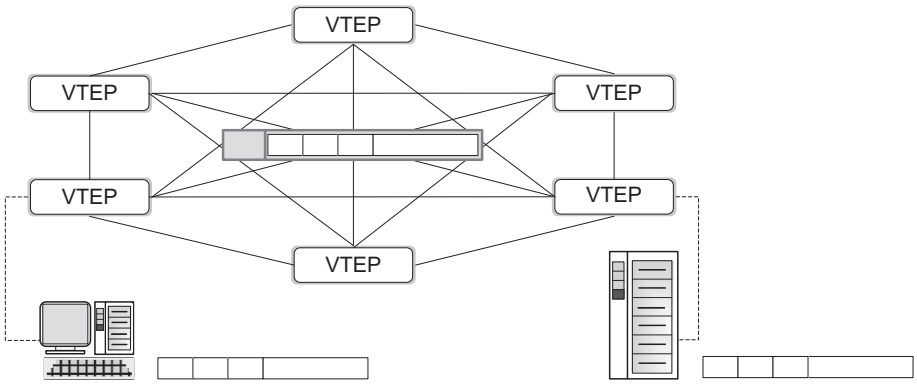


FIG. 4.14

Virtual tunnel endpoints.

virtual links are the tunnels interconnecting them. All the traffic on each virtual network is encapsulated as shown in Fig. 4.13 and sent VTEP-to-VTEP. The reader should note that the tunnels depicted in Fig. 4.14 are the same as the links interconnecting hypervisors in Fig. 4.12. As Fig. 4.12 indicates, there can be multiple overlay networks existing independently and simultaneously over the same physical network.

In summary, SDN via Hypervisor-Based Overlay Networks is well-suited to environments such as data centers already running compute and storage virtualization software for their servers. It does address a number of the needs of an SDN solution. First, it addresses MAC address explosion in data centers and cloud environments, because all those host MAC addresses are hidden within the encapsulated frame. Second, it addresses VLAN limitations, because all traffic is tunneled and VLANs are not required for supporting the isolation of multiple tenants. Third, it addresses agility and automation needs, because it is implemented in software and these virtual networks can be constructed and taken down in a fraction of the time that would be required to change the physical network infrastructure.

Nevertheless, these overlay networks do not solve all the problems that can be addressed by an Open SDN solution. In particular it does not address existing issues within the physical infrastructure, which still requires manual configuration and maintenance. Moreover, it fails to address traffic prioritization and efficiency in the physical infrastructure, so confronting STP's blocked links and QoS settings continue to challenge the network engineer. In the introduction to this section we said that with overlays the current physical network is left as it is, with networking devices and their configurations remaining unchanged. Under this assumption, hypervisor-based overlays do not directly address the desire to open up network devices for innovation and simplification, since those physical network devices have not changed at all. Indeed, while one of the benefits of overlays is the ability to build the virtual network on top on an existing physical one, the overlay strategy can also be used over a *new* physical network that incorporates, for instance, more simple switches. Hence, the overlay strategy does not itself preclude device simplification.

DISCUSSION QUESTION:

What do you find most compelling about the SDN via Overlays solution? Why do you think this solution is popular in certain domains?

4.7 CONCLUSION

This chapter has described the basic functionality related to the manner in which an SDN solution actually works. It is important to realize that there is no fundamental incompatibility between the hypervisor-based overlay network approach to SDN and Open SDN. In fact, some implementations use OpenFlow to create and utilize the tunnels required in this kind of network virtualization. It is not unreasonable to think of these overlay networks as stepping stones toward a more complete SDN solution which includes SDN and OpenFlow for addressing both the virtual as well as the physical needs of the network. In Chapter 6 we will delve more deeply into the SDN alternatives introduced in this chapter as well as other alternatives not yet discussed. First, though, in the next chapter we provide a detailed overview of the OpenFlow specification.

DISCUSSION QUESTION:

Of the three different types of SDN we've introduced—Open SDN, SDN via APIs, and SDN via Overlays—which do you think will be most popular in the near future, and why? Or longer term?

REFERENCES

[1] Shenker S. The future of networking, and the past of protocols. Open Networking Summit, October 2011, Stanford University.

[2] Industry Leaders Collaborate on OpenDaylight Project, Donate Key Technologies to Accelerate Software-Defined Networking. OpenDaylight, April 2013. Retrieved from: http://www.opendaylight.org/announcements/2013/04/industry-leaders-collaborate-opendaylight-project-donate-key-technologies.

[3] Kerner S. Brocade Doubles Down on Open SDN. Enterprise Networking Planet, May 23, 2014. Retrieved from: http://www.enterprisenetworkingplanet.com/netsysm/brocade-doubles-down-on-open-sdn.html.

[4] NetFlow traffic analyzer. Solarwinds. Retrieved from: http://www.solarwinds.com/netflow-traffic-analyzer.aspx.

[5] Production quality, multilayer open virtual switch. Open vSwitch. Retrieved from: http://openvswitch.org.

[6] Open thin switching, open for business. Big Switch Networks, June 27, 2013. Retrieved from: http://www.bigswitch.com/topics/introduction-of-indigo-virtual-switch-and-switch-light-beta.

[7] Kerner S. New Cisco CEO's Message? Faster, Faster, Faster. Enterprise Networking Planet, October 7, 2015. Retrieved from: http://www.enterprisenetworkingplanet.com/netsysm/new-cisco-ceos-message-faster-faster-faster.html.

[8] OpenFlow Management and Configuration Protocol (OF-Config 1.1.1). Open Networking Foundation, March 23, 2013. Retrieved from: https://www.opennetworking.org/sdn-resources/onf-specifications.

[9] Pfaff B, Davie B. The Open vSwitch Database Management Protocol. Internet Draft, Internet Engineering Task Force, October 2013.

[10] Wang K. Floodlight documentation. Project Floodlight, December 2013. Retrieved from: http://docs.projectfloodlight.org/display/floodlightcontroller.
[11] Learn REST: A RESTful Tutorial. Retrieved from: http://www.restapitutorial.com.
[12] Lawson S. Network heavy hitters to pool SDN efforts in OpenDaylight project. Network World, April 8, 2013. Retrieved from: http://www.networkworld.com/news/2013/040813-network-heavy-hitters-to-pool-268479.html.
[13] Open daylight technical overview. Retrieved from: http://www.opendaylight.org/project/technical-overview.
[14] Level 3's Waters on integrating tw telecom, SDN and deepening the metro network footprint. FierceTelecom, November 12, 2014. Retrieved from: http://www.fiercetelecom.com/special-reports/level-3s-waters-integrating-tw-telecom-sdn-and-deepening-metro-network-foot.
[15] Nox. Retrieved from: http://www.noxrepo.org.
[16] Erikson D. OpenFlow @ Stanford, February 2013. Beacon. Retrieved from: https://openflow.stanford.edu/display/Beacon/Home.
[17] Ruby: a programmer's best friend. Retrieved from: http://www.ruby-lang.org.
[18] Trema: full-stack OpenFlow framework in Ruby and C. Retrieved from: http://trema.github.io/trema.
[19] ONOS project partners with Linux Foundation. Retrieved from: http://onosproject.org/.
[20] Goldstein P. AT&T's Donovan: Traffic from 'millions' of wireless customers now running on SDN architecture. FierceWireless, October 8, 2015. Retrieved from: http://www.fiercewireless.com/story/atts-donovan-traffic-millions-wireless-customers-now-running-sdn-architectu/2015-10-08.
[21] Baldwin H. Are carriers really leading the charge to SDN? Infoworld, November 19, 2014. Retrieved from: http://www.infoworld.com/article/2850034/are-carriers-really-leading-the-charge-to-sdn.html.
[22] Project BOULDER: Intent Northbound Interface (NBI). Retrieved from: http://opensourcesdn.org/ossdn-projects/.
[23] Atlas A, Halpern J, et al. An architecture for the interface to the routing system. Internet Draft, Internet Engineering Task Force, February 20, 2016.
[24] Mahalingam M, Dutt D, Duda K, Agarwal P, Kreeger L, Sridhar T, et al. VXLAN: a framework for overlaying virtualized layer 2 networks over layer 3 networks. Internet Draft, Internet Engineering Task Force, August 26, 2011.
[25] Sridharan M, et al. NVGRE: network virtualization using generic routing encapsulation. Internet Draft, Internet Engineering Task Force, September 2011.
[26] Davie B, Gross J. STT: a stateless transport tunneling protocol for network virtualization (STT). Internet Draft, Internet Engineering Task Force, March 2012.