White Paper

The Role of DPI in an SDN World

Prepared by

Graham Finnie
Chief Analyst, *Heavy Reading*
www.heavyreading.com

on behalf of

QOSMOS
Your Network is Information

www.qosmos.com

December 2012

Executive Summary

Deep packet inspection (DPI) has been widely deployed by all types of network operators over the past decade, and its use has accelerated over the past two years as new use cases emerge, centered in particular on policy management and analytics. According to *Heavy Reading*'s **Policy & DPI Tracker**, the policy-related DPI market grew just more than 20 percent in 2011 to reach $550 million, and is expected to continue to grow through the coming years.

The fundamental driver for deployment of DPI and associated techniques has been the need among operators to gain a better insight into IP traffic patterns and user behavior, and to act on that information to improve network performance, reduce the cost of bandwidth, control congestion and enhance subscriber quality of experience (QoE). At root, DPI helps operators regain control over a network that is now primarily carrying third-party applications and services, by accurately identifying those applications in real time.

Separately, a new development has rapidly emerged out of enterprise and academic environments, called software-defined networking (SDN). SDN aims first to separate out all network "control" functions from the simple data forwarding function in network switches and routers, and enable the network to be treated as a programmable resource.

Though the two developments are unrelated, they have similar objectives. Simply put, DPI seeks to make the network application-aware, while SDN seeks to make applications network-aware. On the face of it, at least at an intuitive level, the two are well-matched. However, much of the detail in SDN still remains to be resolved. In its campus/enterprise heartland, the Open Networking Foundation's (ONF) OpenFlow specification has emerged as the main mechanism for separating data forwarding from control functions, but it's not clear whether this will also be at the heart of carrier deployments of SDN; OpenFlow is only one way to do SDNs.

In its more radical variants, SDN brings revolutionary change to network architectures, raising questions about where capabilities such as policy, security and DPI will be located, and what they are for. In particular, will the current use cases for DPI still be relevant in a fully-developed SDN environment? And how can vendors and users of DPI prepare for SDN, given that the timing of SDN remains vague?

Meanwhile, a network operator initiative called **Network Functions Virtualization** has been launched with aims that are complementary to the ONF's. As the name implies, its aim is to encourage the virtualization of a wide range of network functions to reduce equipment and power costs and improve service velocity.

This white paper explores the issues raised by these initiatives and offers tentative conclusions about the likely role of DPI in future SDN-based and virtualized carrier networks. **Section 1** examines the key principles, drivers and likely trajectory for SDNs in existing networks, and looks at the symbiotic relationship between SDN and cloud services. It also provides a brief look at ongoing work by the IETF and ONF in this area, early product launches and vendor roadmaps, and draws interim conclusions about likely deployment scenarios for SDN. **Section 2** reviews the role of DPI in current networks and places it in the context of SDN, looking specifically at the current use cases and their applicability in an SDN environment, and at the question of where DPI is best located in an SDN network. **Section 3** presents a view from the industry, in which leading DPI vendor Qosmos sets out its views about of the role of DPI and "network intelligence" in the ongoing SDN revolution.

The SDN Revolution

Software-defined networking has two founding principles:

- To create a layered network architecture in which all control functions (in particular routing) are separated out from the simple network forwarding function
- To then make networking functions available as programmable resources, via a logically centralized "controller," connected via standardized interfaces (typically APIs) to "applications."

Note that centralization of the controller is a core principle, providing an end-to-end view of the network to users and applications. Note also that applications in SDN-speak include not just conventional business and consumer applications but also things such as optimization, policy, load balancing, and so on. These basic principles are set out visually in **Figure 1**.

Figure 1: Basic SDN Architecture

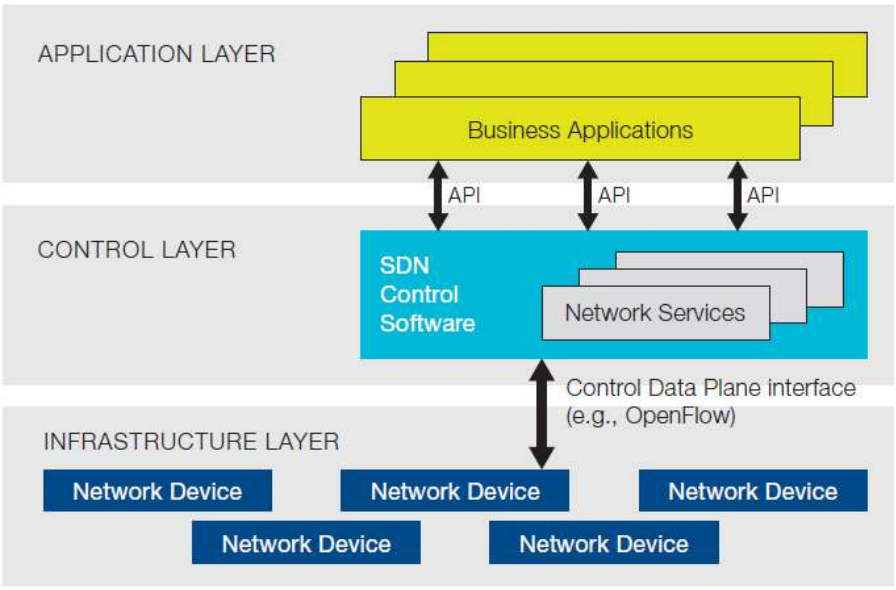

Source: ONF

The main benefits of SDN are:

- Applications and network services are no longer tied to particular network elements and physical infrastructure. This makes for a much more agile and flexible (or "elastic") network and services environment in which resources can be turned off and on, on demand, and applications can get what they need quickly and easily, making it much easier to launch new services into the network. For this reason, it is often associated with another key next-generation network objective, virtualization.
- It allows the network devices to be built using low-cost, high-performance commercial off-the-shelf hardware, and merchant silicon rather than cus-

tom ASICs, replacing complex routers in which data and control planes are today often integrated and proprietary.

- It allows resources such as routing, security, policy management, etc. to be made available as services via APIs. Initially these APIs will likely be vendor-specific, but they should eventually become standardized.
- It facilitates operator plans to offer a wide range of functions "as a service," including functions such as firewalls, load balancers, IMS features and policy management. This means that SDN is the ideal basis on which telcos can build cloud services.

ONF & Beyond

The SDN concept in its modern version was invented by the ONF, an organization whose board includes large telcos (Deutsche Telekom, NTT and Verizon), large Web players (Facebook, Google, Microsoft) and academics (Stanford, UC Berkeley). There are no vendors on the board, and the ONF's stated aim is "The transformation of networking through the development and standardization of a unique architecture called Software-Defined Networking (SDN), which brings direct software programmability to networks worldwide."

The ONF has developed the most important technical SDN standard to date, called OpenFlow, now in version 1.3.0. OpenFlow structures communications between an OpenFlow controller and, for example, an OpenFlow-equipped Ethernet switch (see **Figure 1**), which can be built on standard servers and merchant silicon.

It's worth noting that SDN, as a concept, is part of a broader drive toward the separation of control and data transport layers, which has been going on for some years. For example, the Internet Engineering Task Force (IETF) has had a longstanding interest in some of the key concepts now being explored in the ONF, and has defined several standards that predate SDNs and OpenFlow, but are highly relevant to SDN:

- RFC 4655 defines a Path Computation Element-Based architecture that moves only one of the functions of network elements into the cloud – path computation.
- RFC 3746, called ForCES (Forwarding and Control Plane Separation) sets out a new SDN-type architecture for network devices, and was published in 2004. ForCES does not specify that control must be handled in a centralized controller, and it envisages scenarios in which the two planes might be logically separate but reside in the same vendor equipment.

At the same time, some vendors, recognizing the potential importance (as both a threat and an opportunity) of SDN, are beginning to push their own variants on the SDN theme. For example Cisco, in its Open Network Environment proposal, speaks of "*multi-layer platforms*…" and of "*adapting and augmenting*..." what is already deployed, rather than simply replacing existing network gear – an approach that is less radical than that implied in OpenFlow. A key feature of Cisco's approach is the One Platform Kit (onePK), which includes control APIs covering policy, routing and data path, among other things. Brocade, an SDN pioneer, has already introduced hybrid gear that enables Layer 2/3 switches and OpenFlow to run in the same switches.

The ONF itself has established a Hybrid Working Group, headed by Cisco and tasked with establishing the requirements for a hybrid programmable forwarding plane that includes both orthodox switches and OpenFlow switches.

Meanwhile, a separate initiative emerged in October 2012, led by a group of 12 major network operators including AT&T, BT, Deutsche Telekom, Orange, Telefónica and Verizon, under the rubric of Network Functions Virtualization. In an introductory white paper, the group called for "international collaboration to accelerate development and deployment of [virtualized network functions] based on industry-standard servers." The white paper draws a clear distinction between network virtualization and SDN, but sees the two as "highly complementary."

The aims of this initiative are ambitious. Housed within ETSI (though not directly developing standards itself), it lists 11 equipment categories that could lend themselves to virtualization, including switching elements, network nodes (e.g., P-GW), home gateways, DPI, IMS, policy control, firewalls and optimization gear.

Embarking on SDN: Data Centers Are Center Stage

For the time being, however, attention is focused on OpenFlow as a standard that is already delivering tangible results. The many potential benefits of SDN have attracted strong interest from major network operators such as AT&T, BT, Deutsche Telekom, France Telecom, NTT and Verizon, as well as wholesale carriers such as Colt and Interoute. Some are already using OpenFlow in limited use cases, or trialing it. Meanwhile, vendors are lining up to share their visions or roadmaps.

Both groups are being chivvied along by the threat of competition. For operators, Google's decision to use OpenFlow to connect its huge data centers highlighted the threat posed to operator margins by new networks that are being operated at a fraction of the cost of conventional networks. Meanwhile, SDN startup vendors, most prominently Nicira (now part of VMWare) and Big Switch, are putting the heat on conventional vendors such as Cisco and Juniper, forcing them to respond, and there are many more startups now ramping up. Most major vendors have published roadmaps of some kind for SDN, and most will at a minimum support OpenFlow in existing equipment this year or next.

As the **sidebar** notes, a strong desire among network operators to create XaaS cloud-based services is helping stoke interest in SDN because of the strong synergies between the two.

Because these services reside at the edge, in the data center, this is the main focus for SDNs initially, since this is the area where the biggest gains could be realized. While compute functions have already been virtualized in many cases, network functions are tied to closed, proprietary networking gear, so

SIDEBAR: SDN and the Cloud

SDN and the cloud enjoy a near-symbiotic relationship, since the cloud depends on an environment in which functions that were formerly locked in network appliances and related equipment can be abstracted, virtualized and made available "as a service." This is exactly what an SDN, at least on paper, can provide. The end-game: All resources comprise one single logical data center in which everything, however geographically dispersed, is connected at LAN level. Moreover, virtualization means that applications are no longer tied to specific network servers and can be moved if necessary in response to demand.

For larger conventional network operators with dense national networks, this creates a unique opportunity: Their highly distributed networks give them many more options on where to locate resources, conferring a competitive advantage when delivering certain kinds of services.

this is the area where separation of functions could yield the biggest benefit.

However, this is by no means the only area where SDN is attracting interest. The relentless increase in bandwidth demand, driven by proliferating smartphones and OTT applications, is leading mobile operators to look for new, more cost-effective ways to both use available spectrum and bandwidth (for example, by applying virtualization and cloud techniques) and launch applications more quickly into the network to enable greater service differentiation and new sources of revenue.

So powerful is the SDN concept, and so compelling the need, that the core principles embodied in it now look certain to be deployed. However, there are still many areas of uncertainty, suggesting that SDN may prove to be a slow-burning revolution in conventional telco networks, rather than an overnight sensation. **Figure 2** provides a view of the possible timing of these changes.

Figure 2: SDN Timeline

Source: Heavy Reading

One big issue, as we noted in the last section, is whether the SDN architecture as set out by the ONF is appropriate for service providers, especially large network operators with millions of customers. Detractors argue that OpenFlow is too radical in its implications for existing network operators (though this doesn't preclude it being used by new operators), since it implies the complete replacement of their existing networks – a massive undertaking.

Heavy Reading believes that a hybrid approach is more likely to be deployed by major network operators, enabling them to make the transition to SDN in their own time, often maintaining established vendor relationships, and with a mixture of "open" and "closed" equipment. And ultimately, SDN will enable programmability at multiple, yet-to-be-defined layers of the network, rather than just between the control layer and infrastructure layer, as in OpenFlow.

Big early movers such as Verizon and Deutsche Telekom are already signaling a preference for hybrid SDN switches/routers that do not require wholesale replacement of the existing network infrastructure.

It's worth noting in this respect that in mobile networks based on 3G and 4G, some of the principles of SDN are already embodied (or at least implied) in the existing standards. The 3G model already separates out many functions and creates standard interfaces between them (as well as a standardized signaling model, Diameter) that should smooth the way to an SDN future.

Although the layering shown in **Figure 3** differs in some important respects from the ONF architecture (especially in how it defines "applications," "control" and "transport,") it does acknowledge the need to break up vertical technology stacks,

and this approach has allowed many specialized companies to flourish, using standard 3GPP interfaces to connect to network switches and routers.

Figure 3: 3GPP Layered Architecture

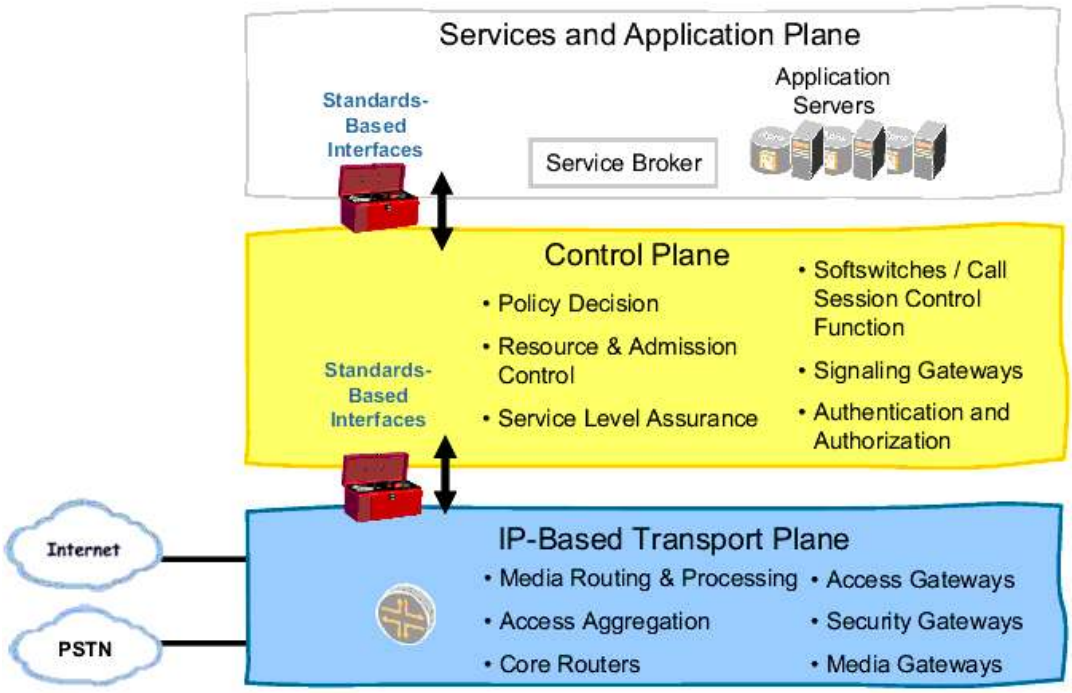

Source: Juniper Networks

However it is implemented, an SDN at root is a concept in which everything in networks becomes available as a programmable resource. This includes management functions, application servers, IMS services, session border controllers and radio access networks, among other things.

In this regard, it's important to note that the "applications" in the SDN applications layer will include not only traditional enterprise and consumer apps (including OTT apps), but also telco business and operational apps, such as billing and provisioning; core communication apps such as voice and IMS; and networking apps such as firewalls, load balancers and policy servers.In this respect, the separate Network Functions Virtualization initiative could prove important, since it explicitly aims to virtualize most of these network functions, deploying them as software on industry-standard (e.g., x86) hardware.

This SDN plan strongly suggests that one vital requirement will be the collection, analysis and presentation, as a usable resource, of detailed intelligence from the network – a function for which DPI is well suited, and already playing a vital role in the majority of networks today. Core current concepts such as load balancing, Layer 4-7 switches, policy management and application delivery controllers (ADCs) – which rely on a deep, real-time insight into higher layers that identifies applications and other metadata on traffic – are likely to play an even bigger role in an SDN network than they do today. As things stand, however, there is no clear guidance yet from the ONF on how this gap is to be bridged. The next two sections consider this issue and how it might be addressed.

DPI in an SDN World

As the previous section showed, SDN has potentially revolutionary implications for network infrastructure and for the vendors who supply it – but the precise nature and timing of the changes described remain quite obscure at this point.

How, then, should vendors using DPI, and network operators deploying it, plan to transition to SDN? Although much of the detail remains to be worked out, we can begin to see how this area might evolve. Before considering that, though, it is worth reviewing where things currently stand with DPI.

A "Must-Have" Technology

In previous survey work with both network operators and equipment vendors, we have shown that DPI is now a "must-have" technology (see **Figure 4**). Use of DPI by network operators has expanded rapidly in the past five years, and most operators have equipment in their network.

Figure 4: How Important Is DPI for You Today?

Source: Heavy Reading Survey of Network Equipment Vendors, 2011

What has driven DPI's ascent? Space precludes a full discussion of this question, but the key catalysts are:

- The need to identify and manage high-impact applications such as P2P and video streaming traffic in order to both reduce or ameliorate their effects, especially on network congestion, and improve the QoE of applications such as video streaming.
- The need to understand what customers are using and doing (using the monitoring and analytics capabilities that are now a part of most DPI software) in order to offer more appropriate and personalized service packages to them.
- Widespread deployment of the 3GPP Policy Control and Charging (PCC) standard, with DPI feeding back information to policy servers to help with decision-making; typically the same hardware or software also handles policy enforcement.

- The desire to create more sophisticated service tiers for mobile data customers – in which, for example, specific allowances are applied to some applications, or certain content or sites are zero-rated.
- The desire to monetize more of the value in IP services and applications, including the development of so-called two-sided business models that engage with third-party applications and content owners.
- The need to protect networks from spam, viruses, DDoS attacks and harmful or illegal content, using DPI to identify threats; relatedly, regulatory drivers such as the requirement for lawful intercept; and for parental or enterprise content control systems.

DPI is now a standard option in 3G gateway GPRS support nodes (GGSNs) and 4G LTE packet gateways (P-GWs), and has also helped build a flourishing specialist appliance industry and a supporting industry building DPI engines that can be deployed in a wide variety of equipment. According to *Heavy Reading*'s **Policy Control & DPI Tracker**, the network operator DPI market was worth approximately $550 million (excluding security-only use) in 2011, and grew more than 20 percent over the prior year.

With the imminent transition to 4G LTE, the requirement to develop, deploy and adapt fine-grained policies on a per-subscriber, per-service and per-application basis will become even greater

The fundamental underlying driver for DPI is the transition to flat all-IP networks and an open Internet applications environment, already ubiquitous in fixed networks but now spreading rapidly through mobile networks. With deployment of 4G LTE now underway, the requirement to develop, deploy and adapt fine-grained policies on a per-subscriber, per-service and per-application basis will become even greater, enhancing the role of DPI and related technologies further.

Although DPI was not specified by name in 3GPP standards, it was implicit, and in Release 11 of the standards (completed this year), a new entity called the Traffic Detection Function (TDF) was included for the first time, along with a standard interface, Sd, connecting it to the standard policy server, the PCRF.

One other important point: DPI vendors, and DPI itself, are continually evolving in response to industry requirements. For example, vendors now routinely track flows and packets to identify applications and trends. And some go far beyond simple identification of the application and can feed back information on device type, length of connection, frequency of connection and other types of "metadata."

The emergence of specialized DPI engines that can be incorporated into any type of networking equipment suggests that DPI will be even more widely distributed in the next year or two than it is today. It can already be found in GGSNs and P-GWs; probes and other similar equipment; policy enforcement appliances; video and Web optimization equipment; analytics software; Layer 4-7 switches; load balancers; ADCs; and specialized security equipment.

Fitting DPI Into SDN

Use of DPI is widespread and growing, and the range of ways in which it is used continues to evolve. However, the core capability remains what it has always been: to identify types of traffic (applications) running in the network, in real-time, and to associate that information at an increasingly granular level with other data such as subscriber, device, location, and so on.

The core question is, therefore, will operators of an SDN still need to be able to identify and analyze, in real-time, the applications and traffic running on their networks? As noted earlier, while the ONF work acknowledged the critical importance of the applications layer, it has not specified how this layer will connect with the other layers, or precisely how network and applications intelligence (as well as subscriber intelligence) will be collected, correlated and disseminated. Until now, the focus has been on Layers 2-3, not Layers 4-7.

Yet there is no doubt that Layers 4-7 will become a critical focus for the future. Indeed, if SDN is successful, the nuts and bolts of networking will fade into the background, and the needs of applications and subscribers will move into the foreground. On this reading, network service providers will gradually morph into applications service providers – with all that this implies.

DPI and related techniques will be at the heart of that transformation. It will create a virtuous circle or feedback loop in which a stream of real-time information on performance, application use trends, user behavior, congestion events, device trends and much else besides is fed back to the SDN controller and to the various network and consumer applications connected to it. Using policy and related tools (e.g., optimization software), this will allow for continual adjustment to circumstance, optimizing both the efficiency with which resources are consumed and the quality of the end-user experience – goals that match the ultimate aims of SDN, as well as closely related developments such as virtualization.

In an SDN, planners anticipate that the applications themselves (broadly defined – see definition earlier) would request certain capabilities in the network (e.g., QoS parameters) before launching, and this would shape the way that network resources are allocated (in principle, this could happen in real time). So policies for each app would be set at the time the app is developed in or migrated to a telco Platform as a Service (PaaS), and would cover e.g., performance SLAs, security/compliance requirements (e.g., geofencing), and so on. In the current network model, DPI is effectively a reactive technique: it identifies particular applications inline in real-time at certain key nodal points, such as a GGSN in a 3G mobile network, and applies policies (if necessary) that are downloaded from the policy server. Policies might include blocking, optimizing, prioritizing, and so on.

Despite this difference, most service providers expect to need increasingly sophisticated insight into the behavior of individual subscribers to make better policy decisions and build better service offers and features, including on-the-fly, real-time service offers. This is why analytics is currently the fastest-growing use case for DPI. In the more dynamic service environment implied by cloud/SDN, real-time analytics will become ever more important in helping operators distinguish themselves in ever more dynamic and fast-moving market environments.

Service providers expect to need increasingly sophisticated insight into the behavior of individual subscribers to make better policy decisions and build better services, including on-the-fly service offers. This is why analytics is currently the fastest-growing use case for DPI.

DPI-like capabilities will be used to feed a rich stream of information to the big data analytics packages that are more and more important to telcos as they seek to gain an understanding (increasingly automated) of

what end-users are doing, and shape service offerings accordingly. QoE has moved rapidly up the operator agenda in the past one to two years, but a QoE dashboard is only as good as the information fed into it.

Allied to that is the need to meet (and provide proof that they are meeting) QoE objectives – something that requires that detailed application performance data is collected on a per-subscriber-basis. This is a nascent need in many networks, but will become much more important as operators seek to provide cloud and other services to a range of industry verticals; the upsurge in M2M will also require the ability to offer and monitor SLA performance.

A related benefit here is that this could help operators build stronger relationships with third-party applications and content providers – an essential objective if they are to justify investment in cloud and SDN, not to speak of LTE and other infrastructure upgrades. DPI can provide a continuous stream of real-time information on performance, user behavior (e.g., hold times, frequency, etc.), transactions and other data that can help the third party tune its products to the market.

Figure 5: DPI Use Cases in an SDN

USE CASE	IN AN SDN
Detecting spam, malicious apps	DPI is likely to play a central role in policing applications and data entering the network; equally, it will be required by law enforcement and increasingly by parental control systems.
Analytics	Role of DPI would probably increase as operators focus more on applications, and seek to better tune the relationship between customers, applications and network.
Policy enforcement	Policy control would be even more important in an SDN, and DPI-like capabilities would likely be required for enforcement and as part of the information feedback loop.
Optimization	DPI will feed information to optimization applications to help decide how and when to optimize.
Traffic management	DPI will continue to help refine traffic management—for example, by offloading particular kinds of traffic from 3G/4G to Wi-Fi, or by optimizing video in congested cells.
Service differentiation	Will be closely aligned with the shift to use analytics, with the latter increasingly used to fine-tune services in real time; feeds from DPI data will be key.

Source: Heavy Reading

One critical aspect in all of this is that the applications and associated control elements need a **holistic** view of infrastructure conditions. This is a central goal of the ONF SDN plan, and something that DPI, in principle, can provide, by gathering information throughout the network and feeding it back to the control layer (controller) and to the applications so as to ensure that the right resources and capabilities are made available. This may turn out to be the killer app for DPI.

Locating DPI in an SDN Network

This all begs the question as to where DPI would be located in an SDN. Since the main focus of DPI deployment currently is in mobile networks, it is worth considering how DPI is deployed now, and how this might evolve in future. In a 3G network, DPI is most often to be found in the GGSN, or in a dedicated appliance or blade built by a specialist DPI vendor, often collocated with the GGSN. DPI can also be found in network probes, which are often are at the edge of the network. Other equipment that use DPI include load balancers and specialist security software.

In an environment in which (as discussed earlier) there is still no agreement about the types of hardware, software and services that would be deployed in an SDN, deciding how this might change in an SDN environment is inevitably somewhat speculative. But based on our description of areas where DPI would likely still be needed, we anticipate DPI-like capabilities would be less of a product function, and more of a distributed network capability or resource available on a controlled basis via APIs or other interfaces to any application. In an SDN, the GGSN in its current form would (presumably) be logically split apart, with much of the intelligence migrating to the controller or to the applications layer. Likewise DPI would no longer be associated with a dedicated appliance or blade.

In the operator paper on Network Function Virtualization, the authors specifically refer to DPI as a strong candidate for virtualization, and suggest that this should enable it to be distributed throughout the network:

> "A software-based DPI, providing advanced traffic analysis and multi-dimensional reporting, and showing the possibility of making off-the-shelf hardware work at actual line rates. **Software-based DPI can be pervasively deployed in the network, providing much better analysis capabilities, as well as simpler mechanisms for deployment, update, testing, and to scale it to changing workloads."** *[our emphasis]*

In an SDN, DPI could potentially be collocated with network devices (e.g., as software running in virtual switches), or it could be in the control layer (e.g., in the controller that mediates between applications and switches). In principle, collocating DPI capability in switches would be a more efficient way to run DPI, especially in view of the high CPU resource requirement. One way in which this information might traverse the network is as an extension to the flow table, as described in the next section. Application management systems could then extract relevant information to make real-time service decisions. In principle, this functionality could also be standardized, at least up to a point.

DPI is likely to be at least as important in an SDN as it is today, and making the right decisions about where and how to deploy it is among the more important choices the industry faces.

DPI might also be collocated with the controller, which could be directly fed the information to make key decisions or could share it via APIs with both end-user and network applications (e.g., firewall, video optimizer) that need it. However, much remains still to be decided. There are many interested parties in this debate and as things stand we do not yet have a clear view of where the applications/metadata detection will take place in an SDN. What we do know is that DPI is at least as important in an SDN as it is today, and making the right decisions about where and how to deploy it is among the more important choices the industry faces.

An Industry View From Qosmos

L4-7 Network Intelligence as a Key Enabler for SDN

At Qosmos, we believe that networks need L4-7 network intelligence to become fully service-aware. Our mission is to provide best-of-breed network intelligence technology to equipment vendors, ISVs and integrators that build these next-generation networks and solutions.

Defining L4-7 Network Intelligence

L4-7 network intelligence is created by techniques such as DPI, which analyzes traffic and provides intelligence in the form of application identification (App ID) and additional data about each traffic flow in the form of metadata attributes:

- **Examples of App ID:** SIP, SMTP, YouTube, Facebook, BitTorrent, Skype, etc.
- **Examples of extracted metadata:** URL, file name, browser type, cookies, DNS queries, video codec, IMSI, SIP caller/callee, user ID, login, etc.
- **Examples of computed metadata:** delay, jitter, application response time, etc.

One of the key challenges in SDN is the lack of application awareness in controllers or virtual switches (vSwitches), which prevents them from making smart decisions. L4-7 DPI and metadata provide this much-needed awareness.

Another challenge is that SDN is limited to L2-4 visibility, and cannot be used for efficient traffic steering needed to enable service insertion, since switches cannot differentiate traffic between various types of L7 applications. This means that each specialized system, for example video optimization, has to analyze the entire traffic in order to pick out the relevant flows and process them. With L4-7 intelligence, a switch can redirect each application flow to each specialized service processing. This not only enables efficient service insertion but also facilitates hybrid approaches by making physical and proprietary equipment (non-OpenFlow) SDN-compatible.

In SDN there could also be a tendency to duplicate DPI processing inside numerous applications, each consuming compute resources. But in an SDN environment, L4-7 DPI and protocol metadata can become a shared resource used by controller and applications to save on resources. **Figure 6** (next page) summarizes the key issues and the solutions provided by L4-7 network intelligence.

Implementing L4-7 Network Intelligence in an SDN Architecture

L4-7 software components can run inside [v]switches, where protocol information and metadata are fed northbound to the controller and to the applications. Alternatively, L4-7 inspection software can be embedded in the controller, feeding traffic information to applications through northbound API, or consumed directly by the controller.

In both cases, L4-7 DPI embedded at a few strategic locations in an SDN architecture creates common-format L4-7 traffic intelligence, which can be consumed by different SDN elements.In addition to [v]switches and controllers, some applica-

tions will also embed their own dedicated L4-7 analysis for specific service processing.

Figure 6: SDN Challenges & Solutions

SDN CHALLENGES	SOLUTION WITH L4-7 DPI & METADATA
No L4-7 application awareness in controller nor in vSwitches	L4-7 DPI and metadata engine provides controller and its applications with App IDs and metadata to make smarter decisions.
Difficult to insert services through integration of specialized physical L4-7 switches (FW, LB, video optimization, etc.)	L4-7 DPI and metadata engine analyses traffic and only relevant application flows are re-directed to specialized service processing. This enables hybrid approaches by making physical and proprietary equipment (non-OpenFlow) SDN-compatible.
Applying strict SDN architecture in the data plane may lead to latency and inefficient use of bandwidth	L4-7 DPI and metadata engine co-located in each switch avoids duplicated, remote DPI processing and inefficient use of network resources. L4-7 DPI and metadata engine analyses traffic and only relevant application flows are sent to appropriate specialized equipment / applications. This enables each, specialized equipment to focus its processing on relevant traffic only.
Difficult to configure switches due to inconsistent App and metadata IDs	In an SDN environment, L4-7 DPI and metadata can become a unified, common-format resource used by controller and applications. Each application uses only relevant metadata, structured in a consistent format to optimize cross-application interaction.
Duplicate DPI processing inside numerous applications consuming resources	In an SDN environment, L4-7 DPI and metadata can become a shared resource used by controller and applications to save on total CPU consumption.

Source: Qosmos

Standardization

A way to make L4-7 network intelligence a key enabler for future infrastructure is to create a standard by extending OpenFlow fields with additional information such as App IDs and corresponding metadata for each flow. This would create a common format that could be used by switches, controllers and applications. In addition, it would even work in proprietary, non-SDN environments: Virtual switches could use extended DPI fields even before the official extension to OpenFlow.

L4-7 Network Intelligence & Network Functions Virtualization

The objective of Network Functions Virtualization is to leverage standard IT virtualization technology to consolidate network equipment onto industry-standard servers, switches and storage. With Network Functions Virtualization, L4-7 DPI can migrate from being embedded in each network appliance to being a shared function residing in standard switches and servers.