IHS TECHNOLOGY WHITE PAPER

The Need for DPI in Cybersecurity Solutions

Jeff Wilson

21 March 2017

ihs.com

Contents

Exhibits

IHS™ TECHNOLOGY

COPYRIGHT NOTICE AND DISCLAIMER © 2016 IHS. Reprinted with permission from IHS.
Content reproduced or redistributed with IHS permission must display IHS legal notices and attributions of authorship. The information contained herein is from sources considered reliable, but its accuracy and completeness are not warranted, nor are the opinions and analyses that are based upon it, and to the extent permitted by law, IHS shall not be liable for any errors or omissions or any loss, damage, or expense incurred by reliance on information or any statement contained herein. In particular, please note that no representation or warranty is given as to the achievement or reasonableness of, and no reliance should be placed on, any projections, forecasts, estimates, or assumptions, and, due to various risks and uncertainties, actual events and results may differ materially from forecasts and statements of belief noted herein. This report is not to be construed as legal or financial advice, and use of or reliance on any information in this publication is entirely at client's own risk. IHS and the IHS logo are trademarks of IHS.

DPI and the Application Era

Security solution manufacturers are tasked with the difficult job of helping your customers better protect themselves while they also try to simplify and consolidate their security infrastructure. It can be difficult for everyone in the cybersecurity technology conversation to get on the same page, but there's now one thing everyone can agree on. In today's connected world, there is a single language that all users and devices speak: ***the language of the application***.

In terms of transmitting data across public and private networks, applications are critical because they generate content—the payload portion of any packet. When network control and security devices need to make intelligent decisions about moving packets around in a secure manner, there is a significant amount that can be done by inspecting headers, understanding the source and destination, and maybe even flagging types of applications—but in the end security will always also require knowledge of data and content (the payload).

If you really want to know if there's an exploit buried inside a Microsoft Word document, you have to inspect the content. If you really want to know if regulation-mandated personally identifiable information (PII) is leaving your network in an unsanctioned e-mail, you have to read the e-mail. If you really want to know if you users are using unsanctioned mobile applications to upload confidential work documents to their personal Dropbox accounts, you have to identify Dropbox traffic and then inspect the payload.

To do all of this, you need one key technology: deep packet inspection (DPI). DPI technology and solutions have been around for years, both as standalone platforms and technology embedded in a wide range of security, traffic management, and control devices. Very simply, DPI is defined as a form of computer network packet filtering that examines the data (payload) portion of a packet as it passes an inspection point, searching for traffic instructions, security violations, or any other defined criteria.

© 2017 IHS

In a typical DPI integration, network traffic is captured using, for example, Intel® DPDK libraries for fast packet processing. The flows are then ingested by the DPI engine, which inspects the header and payload using a combination of techniques including behavioral recognition, regexp, statistical analysis, flow correlation, and DNS matching. Next, the DPI engine classifies traffic flows up to the application layer (7) and extracts additional information in the form of metadata and content as required. This information is passed on to reporting and policy engines, which send flow processing instructions to the device using the engine. That device can now make flow processing (and security enforcement) decisions based on the DPI engine's analysis. The flow processing could be done by a firewall, a router, or any other device in which the DPI engine is integrated. The goal of this effort is to turn traffic chaos into useful information, providing context for security decisions such as: Which applications are generating the traffic? Are files embedded? Is protected content being moved around the network in a way that violates policy?

Exhibit 1 DPI Engine Integration Model

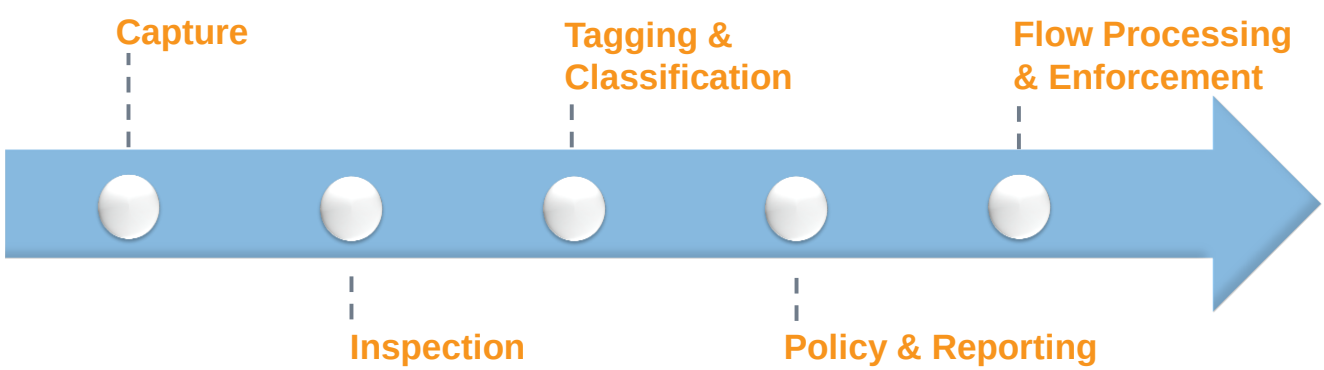

Source: IHS

© 2017 IHS

And while you would assume that DPI is a technology that every security product and platform currently in use supports, it isn't. A significant portion of the firewalls that are in use today are stateful-inspection firewalls that do not have DPI capability or are weakened by very limited DPI technology (for example, open source DPI). Many messaging and web security platforms don't yet support DPI. Many security solutions in virtualized environments don't support DPI. DPI is more broadly available in next-generation firewalls (NGFW) and threat intelligence platforms, but most enterprises and service providers don't have these advanced products—and it's often highly impractical to re-route all traffic that needs deep inspection back to an NGFW.

DPI engines are difficult to build and maintain, and when implemented improperly can be incredibly taxing on system performance. As a result, many core security platform and product vendors have not invested in building their own DPI engines. There are third-party engines available for integration, and the remainder of this paper will look at the compelling case for integrating one of these engines into your existing security product/platform. A commercial DPI engine can be a replacement or a complement to internally developed DPI technology.

What Drives Buyers to Demand DPI?

The decision to invest the time and resources required to integrate DPI into existing security solutions is not one to make lightly, and it's important to step back and look at the big forces causing your customers to make spending decisions. The drivers discussed apply to DPI, but they aren't unique to it. They are the forces that move the broader cybersecurity technology market, and any discussion about security technology can be framed using these drivers. DPI does, however, have a significant role to play in each driver.

Exhibit 2 Security Trends Orbiting DPI

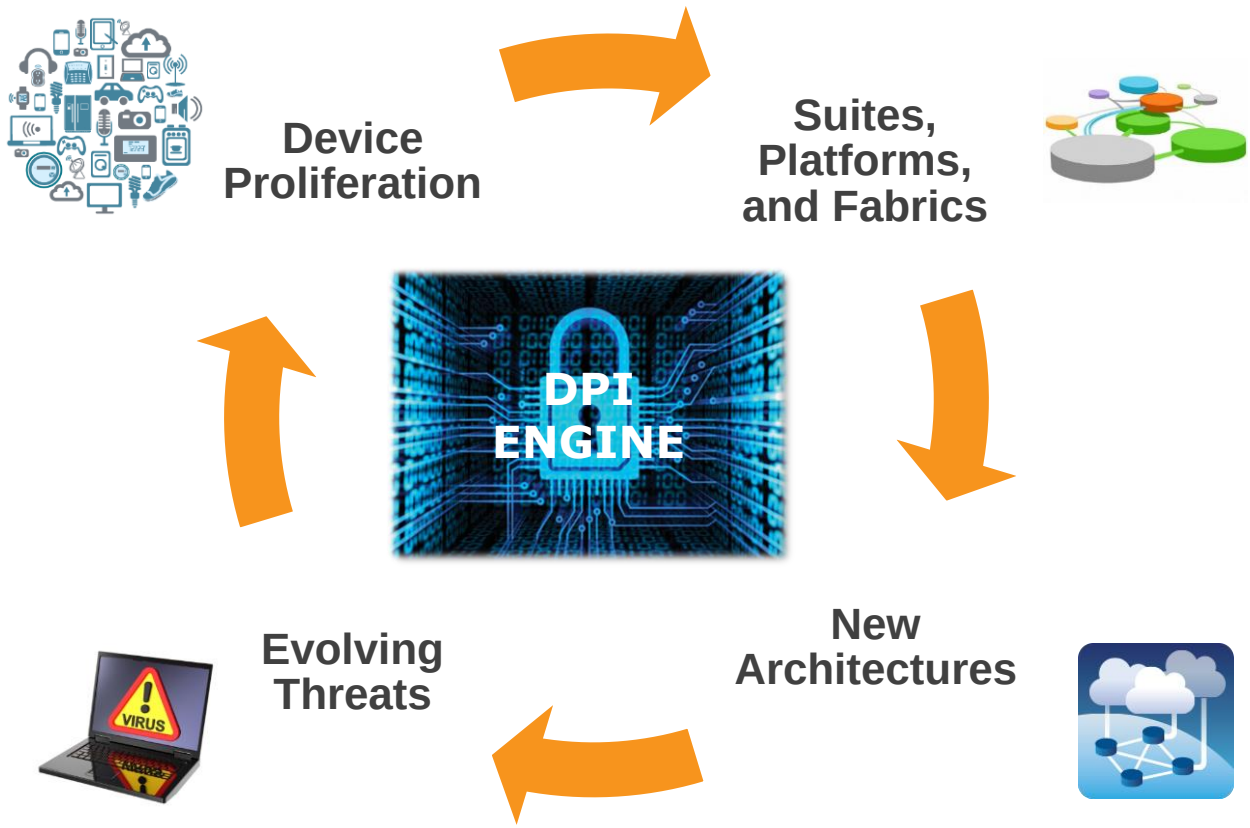

Source: IHS

© 2017 IHS

Device proliferation: Smartphones, tablets, machine-to-machine (M2M), the Internet of Things (IoT), and the Internet of Everything (IoE) are driving fundamental changes in how, where, and why security technology is deployed. There are security solutions for every part of the device chain: hardware, software, and network. More devices mean more traffic on the network and a greater need for traffic visibility. The "app" mindset that drives the need for DPI exists in large part because of mobile devices and, moving forward, the broad range of devices that will be the IoE. These IoE devices won't be able to install security client software and at best may support device authentication and encryption—but they likely won't be able to perform any kind of deep security inspection at the device level. This means that security solutions in the network, using DPI, will be required for deep understanding.

Suites, platforms, and fabrics: Defense tools and strategies have been built over time, layering new technology on top of old as new threats emerge, leaving most companies with a complicated infrastructure with many holes. Everyone from the smallest business to the largest carrier is trying to collapse defense layers and simplify security architecture and protection to improve effectiveness. At the same time, they want to have the best security available in the platforms that they do finally settle on, which means there will be a requirement to embed DPI inside security solutions.

New architectures: Security technology isn't evolving in a bubble; rather, it's tied to network architectures—and the emergence of virtualization, software-defined networking (SDN), network functions virtualization (NFV), and cloud services are driving significant changes in IT infrastructure and network architectures. These changes have a major impact on how security technology is consumed. That said, granular visibility and control of traffic inside virtualized environments becomes both more important and more difficult. Functional DPI in virtualized and cloud environments will be a key part of ensuring security parity in the old world and the new.

Evolving threats: Security technology innovation is driven by changes in the threat landscape, which is ever-changing. Security technology solutions have to be engineered to defeat human ingenuity and not physics or math, and as a result there is a cyclical pattern of threat protection technology development. New threats emerge, new technologies are built to combat those threats, those technologies are absorbed into larger platforms, and the process repeats infinitely. DPI is a building-block technology in the fight against evolving threats because it enables processes like heuristics and behavioral analysis.

Historically, DPI was something that large service providers cared about for things like traffic management and quality of service. And while those uses of DPI (and standalone DPI solutions) are still compelling, both enterprises and service providers have a significant need for DPI to be integrated into a range of security and network solutions.

Exhibit 3 DPI Applications

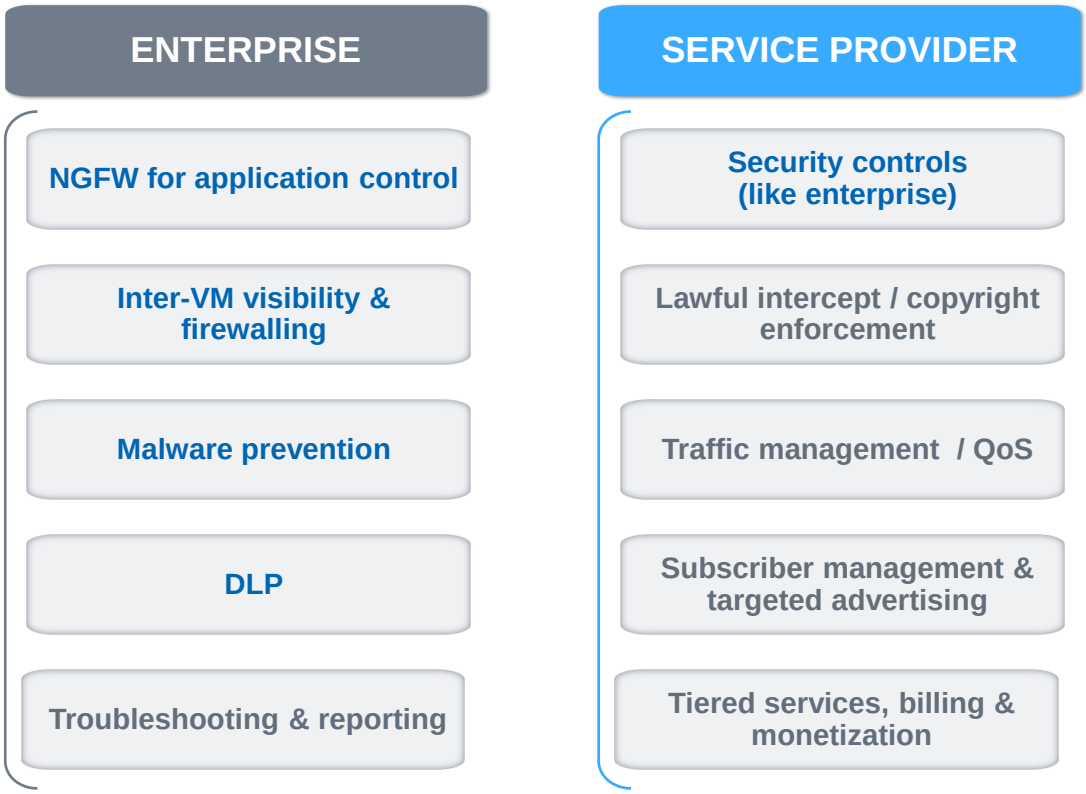

Source: IHS

© 2017 IHS

The applications may be different, but the core functions that enable them will be the same. Ultimately, this points to a large target audience of customers looking to simplify their networks and support deeper inspection across multiple platforms and uses. Next, we'll look at two sample implementations.

Sample Application: NGFW-Physical and Virtual

The difference between a stateful-inspection firewall and a next-generation firewall is, in a nutshell, DPI. When properly integrated into a firewall platform, DPI allows that firewall to classify and control applications running across the network and apply security policies, all at the speed of a traditional stateful-inspection firewall.

Exhibit 4 Firewall + DPI Engine = NGFW

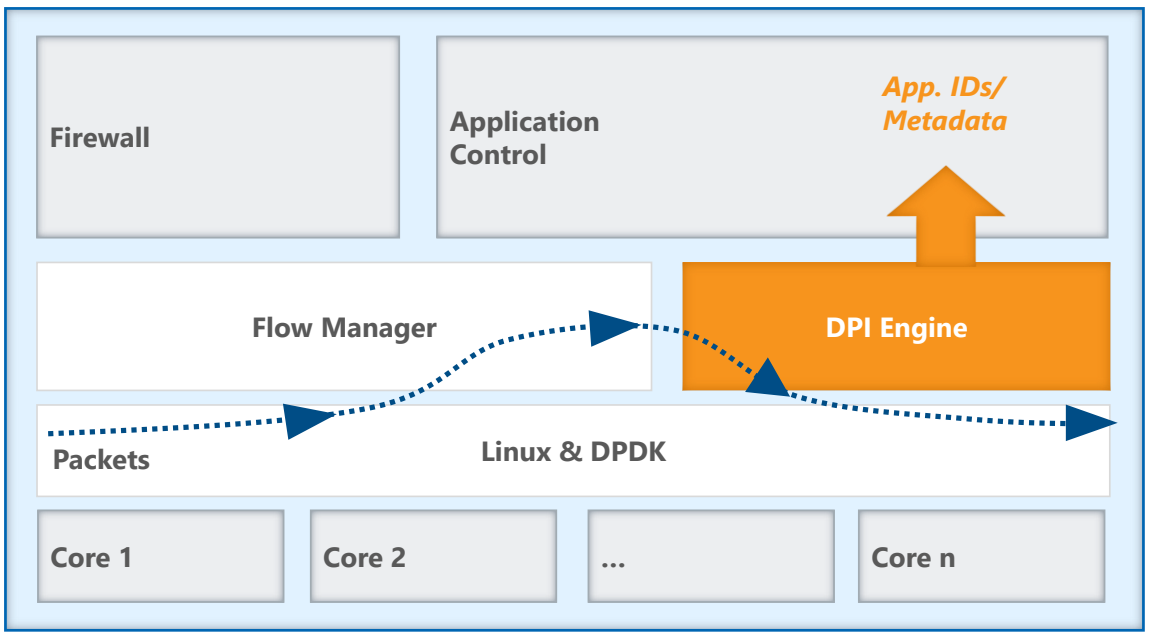

Source: IHS

© 2017 IHS

For this all to work properly, the DPI engine has to be able to identify and classify applications, as well as create metadata based on payload content and statistical traffic analysis. The metadata is key because it's the tool that allows the firewall to go from making simple decisions based on gross application ID (this is Facebook) to making more granular decisions based on the content of the Facebook traffic (this is a download within Facebook, which is not allowed). The type and variety of metadata that the DPI engine can extract, as well as the frequency at which metadata and app signatures are updated, is key to understanding the quality of the DPI engine. A DPI engine is able to extract a significant amount of valuable metadata, especially in controlled environments such as enterprise networks where it can have access to decrypted traffic based on systems such as secure sockets layer (SSL) interception proxies.

As the world moves rapidly toward private and public cloud infrastructure and more enterprise and service provider traffic moves through virtual machines, full layer-7 visibility within the virtualized environment will become increasingly important. Not only will this visibility allow you to see traffic up to layer 7, but ultimately you'll be able to build solutions that can implement firewalling rules using layer-7 data (application information and metadata based on payload inspection). Therefore, the benefits discussed above of combining traditional stateful-inspection firewalling with DPI can be realized in both traditional and virtualized environments.

An added benefit of implementing a solution like this as part of your virtualized security offering is that you don't really have to build anything new. A hallmark of SDN and data centers is that many are built to leverage open-source, or simply "open," tools that are programmable. If you have your own policy tools, or are leveraging things like Open Daylight, you can assemble a solution like this with the addition of a DPI-based layer-7 traffic classification tool using those open interfaces. In the end, this means that you can very quickly develop a working layer-7 visibility and control product for virtualized environments.

Sample Application: Extracting File Content for DLP and Malware Protection

A wide range of security platforms and applications beyond the firewall can also make use of a DPI engine's unique ability to access and stream raw packet content (attached files, e-mail message header/body) and metadata (URLs, IP addresses, sender/receiver info)—and this information can be used for a wide range of security decisions. For example, if a data loss prevention (DLP) system is looking at outbound e-mail attempting to identify content that shouldn't leave the network, the DPI engine can extract file information to check for restricted content such as social security numbers, credit card numbers and, other defined PII, and then the DLP system policy can decide what to do with those outbound e-mails.

Exhibit 5 Extracting File Content

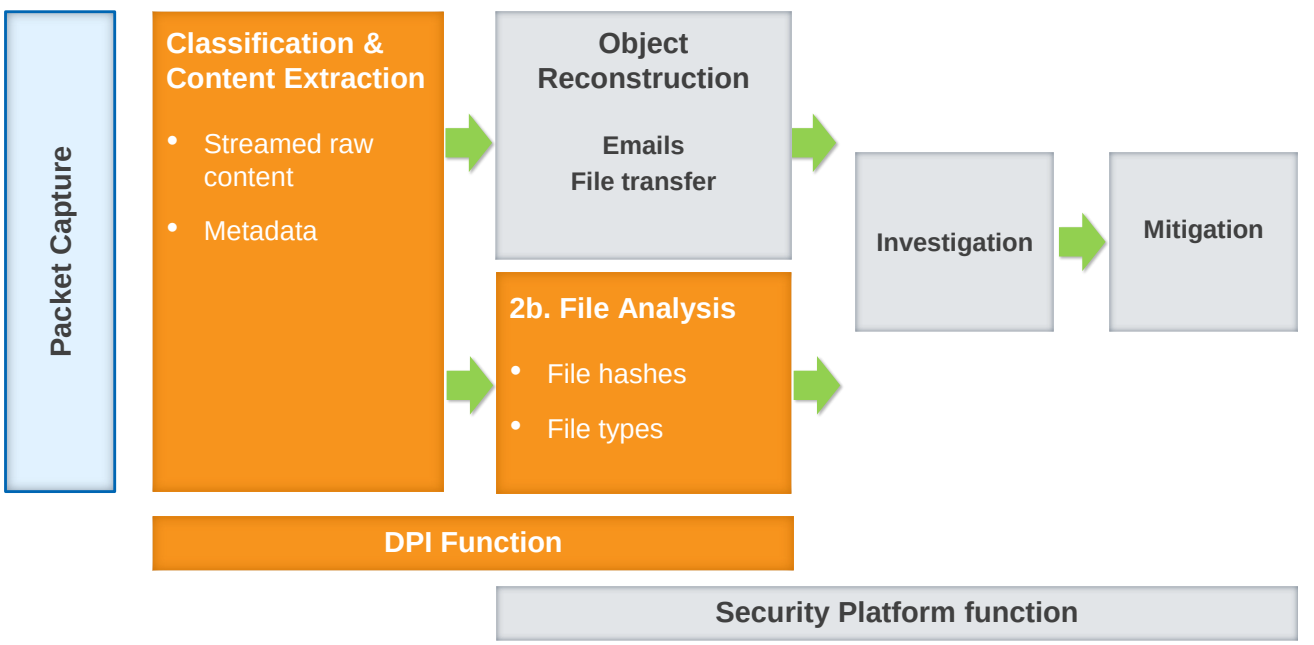

Source: IHS

© 2017 IHS

Any security system that wants to perform malware prevention can also use deep information about files attached to e-mails, sent in file transfers, or moved into on-premises or cloud storage. Identification of file types, file integrity checks, and streaming of raw file content and metadata can be handled by an integrated DPI engine, which then hands off the information to the security solution it has embedded for reconstruction. After reconstruction, the security system can use whatever investigation tools it has—virus scanning, sandboxing, pattern/signature matching—to enable alerts, policy enforcement and, ultimately, mitigation.

Final Thoughts

We will never live in a world with fewer devices and fewer threats than today. Scale and complexity will only increase, and the need for a wide variety of security tools to integrate DPI will only grow—especially as your customers look to simultaneously resist threats and simplify their security infrastructure. If you don't have the resources to build and maintain your own DPI engine and deal with all continuous protocol signature updates, the time is now to find one that you can start to integrate across your solution portfolio.

Commissioned by Qosmos to educate cybersecurity technology companies about the advantages of integrating third-party DPI solutions into their existing security platforms, this paper was written autonomously by analyst Jeff Wilson based on IHS independent cybersecurity technology research.

Contact

Jeff Wilson
Senior Research Director and Advisor,
Cybersecurity Technology
+1 408.583.3337
Jeff.Wilson@ihs.com

IHS Customer Care:

Americas: +1 800 IHS CARE (+1 800 447 2273); CustomerCare@ihs.com
Europe, Middle East, and Africa: +44 (0) 1344 328 300; Customer.Support@ihs.com
Asia and the Pacific Rim: +604 291 3600; SupportAPAC@ihs.com