

White Paper

Digging Deeper Into Deep Packet Inspection (DPI)

© 2007 Allot Communications Ltd. Allot Communications, NetEnforcer and the Allot logo are registered trademarks of Allot Communications. NetXplorer is trademark of Allot Communications. All other brand or product names are trademarks of their respective holders. All information in this document is subject to change without notice. Allot Communications Ltd., and/or its affiliates (collectively "Allot Communications") assume no responsibility for any errors that appear in this document.

Abstract

DPI is the foremost technology for identifying and authenticating protocols and applications (IP flows or sessions in general) conveyed by IP. This paper reviews DPI technology – what it does, how it works, methods of analysis it uses, and the opportunities it offers ISPs and carriers, such as new sources of revenue, solving bottleneck issues and guaranteeing Quality of Service (QoS).

"DPI helps operations improve the performance of interactive applications, preventing certain application traffic from unduly hogging resources and contributing to congestion. It also mitigates the effects of network attacks, which helps reduce capex and opex while increasing subscriber satisfaction, leading to lower churn."

James Crawshaw, Research Analyst, Light Reading Insider

Shallow vs. Deep Packet Inspection

The standard packet inspection process (a.k.a. shallow packet inspection) extracts basic protocol information such as IP addresses (source, destination) and other low-level connection states. This information typically resides in the packet header itself and consequently reveals the principal communication intent.




Figure 1: Shallow packet inspection – data from packet headers

The inspection level in the shallow inspection process is insufficient to reach any application-related deductions. For example, if a packet is the result of an application trying to set up additional connections for its core operation, an examination of the source or destination addresses as they appear within the packet header itself will not reveal any useful information regarding the connections to be used in the future, as requested by the application. Furthermore, as in this example, it is very common that the necessary information is spread over several packet transactions; and once again, examination of the header information alone overlooks the complete transaction perspective.

DPI, on the other hand, provides application awareness. This is achieved by analyzing the content in both the packet header and the payload over a series of packet transactions. Consequently, DPI provides the ability to analyze network usage and optimize network performance, thereby playing a crucial role in the equation between supply and demand faced by every network operator.

Application and Protocol Signatures

What is a Signature?

In order to deal with numerous Internet and network applications and protocols, a methodical and systematic identification process must be employed. Similar to a lawful operation in which fingerprints are used to identify the involvement of individuals in specific incidents, signatures are used to identify applications and protocols.

In their most broad sense, signatures are pattern recipes which are chosen for uniquely identifying an associated application (or protocol). When a new application or protocol is encountered, it is analyzed and an appropriate signature is developed and added to a database (typically referred to as a signature library).

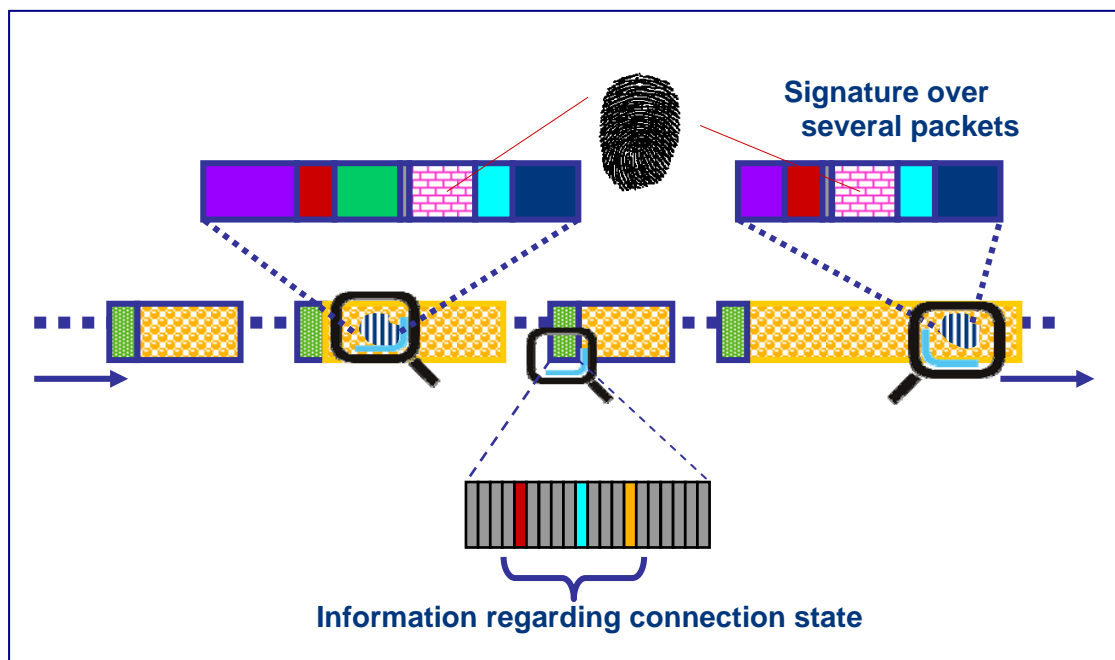


Figure 2: Deep Packet Inspection – analysis of encapsulated content over many packets

Application signatures must be checked on a regular basis as they tend to vary as new application updates or protocol revisions occur. For example, BitTorrent, eMule and Skype tend to upgrade their client software on a regular basis and encourage (and in some cases even force) users to move on to the new release. The use of these new releases with non-up-to-date signatures will dramatically affect classification performance.

False Positives and False Negatives

Although a signature is developed with the intention to uniquely and completely identify its related application or protocol, there are cases in which the signature is not robust (a.k.a. weak signature) and classification problems arise.

False positives is the basic terminology referring to *misclassification* – or in simple terms - the likelihood that an application will be identified as something it is not. If DPI is being used for guiding a subscriber management tool, this may lead to wrongful actions. A typical example of such a wrongful action could be the mistaken lowering of priorities to time-sensitive streaming traffic and resultant introduction of unwanted latency or even packet loss. Consequently, when developing signatures, every effort must be made to achieve zero percent of false positives. A common way to strengthen a weak signature is to use a combination of more than one pattern.

False negatives refers to those cases where it is not possible to consistently identify an application – sometimes the identification is classified, while other times it is missed by the classification tool. There are various reasons for this phenomenon, the most common of which is the fact that some applications can accomplish similar outcomes in several ways in different deployment scenarios. For example, some applications will behave differently if the client software operates through a proxy or firewall compared to the simpler case in which the client interacts with the web directly. Therefore, in these irregular cases, if the signature was developed under the assumption of direct communications, it is likely that the application will not be correctly classified in the case of a proxy or firewall.

"DPI technology allows service providers to improve performance of IP-based multimedia applications, and stops unwanted traffic from consuming network resources and causing congestion. DPI also mitigates network attacks to help reduce capex and opex, while enhancing subscriber satisfaction to reduce churn."
Mark Bieberich, VP Broadband,
Yankee Group

Methods of Signature Analysis

There are several possible methods of analysis used to identify and classify traffic. These range from analysis by port, by string match, by numerical properties, by behavior and heuristics.

Analysis by Port

Analysis by port is probably the easiest and most well known form of signature analysis. The reasoning is the simple fact that many applications use either default ports or some chosen ports in a specific manner. A good example is POP3 used for an email application. The incoming POP3 typically uses port 110, and if it is secure, it will use port 995. The outgoing SMTP is port 25.

However, since it is very easy to detect application activity by port, this is in fact a weakness, particularly because many current applications disguise themselves as other applications. The most notorious example is the Port 80 syndrome, where many applications camouflage as pure HTTP traffic.

As noted above, some applications select random ports instead of using fixed default ports. In this case, there is often some pattern involved in the port selection process - for example, the first port may be random, but the next will be the subsequent one, and so forth. However in some cases the port selection process may be completely random.

For all these reasons, it is often not feasible to use analysis by port as the only tool for identifying applications, but rather as a form of analysis to be used together with other tools.

Analysis by String Match

Analysis by string match involves the search for a sequence of textual characters or numeric values within the contents of the packet. Furthermore, string matches may consist of several strings distributed within a packet or several packets.

For example, many applications still declare their names within the protocol itself, as in Kazaa, where the string "Kazaa" can be found in the User-Agent field with a typical HTTP GET request. From this example, it is possible to understand the importance of DPI for correct classification. If analysis is performed by port analysis alone, then port 80 may indicate HTTP traffic and the GET request will further corroborate this observation. However, since the User-Agent field information is missing, this analysis will result in inaccurate classification i.e., HTTP and not Kazaa.

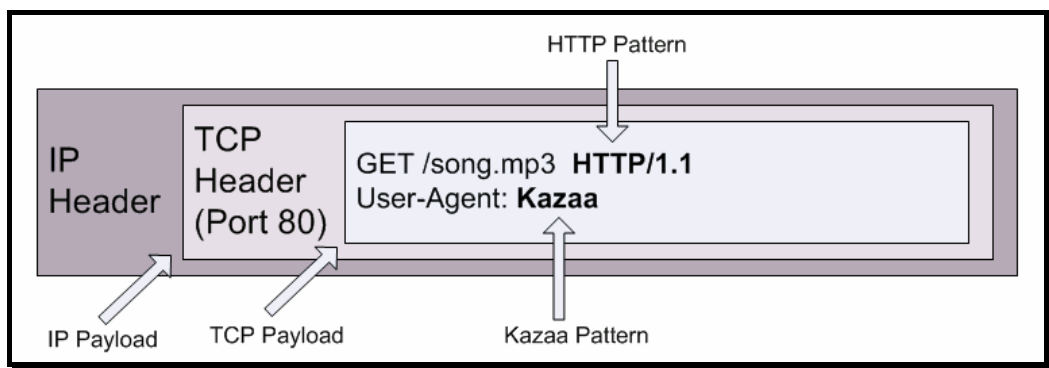


Figure 3: KazaA string match analysis

This example emphasizes once again that several analysis tools are required for assuring proper classification.

Analysis by Numerical Properties

Analysis by numerical properties involves the investigation of arithmetic and numerical characteristics within a packet, and of a packet or several packets. Some examples of properties analyzed include payload length, the number of packets sent in response to a specific transaction, and the numerical offset of some fixed string (or byte) value within a packet.

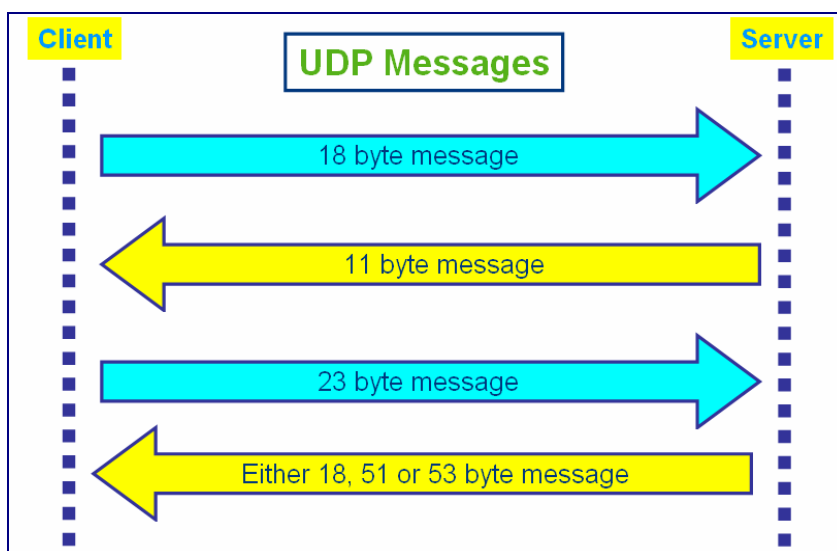


Figure 4: Skype (versions prior to 2.0) numerical properties analysis

For example, consider the process for establishing a TCP connection using some UDP transactions in Skype (versions prior to 2.0). The Client sends an 18 byte message, expecting in return an 11 byte response. This is followed by the sending of a 23 byte message, expecting a response which is 18, 51 or 53 bytes.

Similar to analysis by port and analysis by string match, analysis by numerical properties alone is insufficient, and can often lead to many false positives.

Analysis by Behavior and Heuristics

Behavioral analysis refers to the way a protocol acts and operates. Heuristic analysis typically boils down to the extraction of statistical parameters of examined packet transactions. Often, behavioral and heuristic analysis are combined to provide improved assessment capabilities.

For example, actions leading to other actions can clearly indicate a behavioral pattern which can be traced, as in the case where an active UDP connection eventually transforms into a TCP connection (using the same IP and port settings).

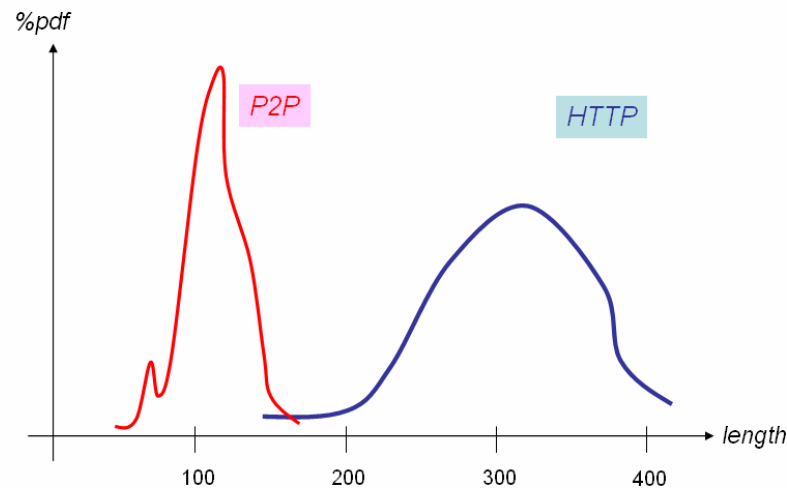


Figure 5: HTTP vs. P2P

Another example of behavior and heuristic analysis is shown in Figure 5, which compares HTTP and a typical P2P file sharing application. If the packet length histogram (PDF) alone is examined while ignoring the file download or upload transaction itself (which tends to use large packet lengths), it becomes apparent that while pure HTTP packets tend to concentrate around a few hundred bytes in length, P2P control layer information tends to use shorter packet lengths. In this way, by examining some short-term statistics, it is possible to conclude whether a port 80 connection carries pure HTTP traffic or other P2P-related traffic.

Use of Encryption and Obfuscation by Applications and Protocols

In the DPI world, life is becoming much more difficult with the use of encryption - the concealing of data to guarantee security and privacy. Encryption is a procedure which renders the contents of a message or a file unintelligible to anyone not authorized to read it. Using very complex mathematical algorithms, encryption makes it very difficult to trace back the information without a key.

Some historical examples of encryption include the Enigma machine used by Germany in WWII, and RC4, a more recent achievement from about 10-15 years ago, which is currently the most popular encryption algorithm used by most P2P protocols. Since the key lengths used are very large, it is almost impossible to reverse engineer and gain some information, like guessing a password.

Encryption should not be confused with obfuscation. Many application control and subscriber management devices utilize DPI with the intention to shape end-user traffic patterns (a.k.a., throttling) and in some cases may even inhibit the use of specific applications. As a result, the creators of several applications have chosen to conceal their operation by scrambling their related communications. The industry term used for this concealment method is *obfuscation* (concealing actions, by making things much more complex than necessary). Some examples of such applications are eMule (version 0.47c) and BitTorrent. In other cases, some over-the-top applications tend to scramble their communications due to the use of proprietary technology or for guaranteeing operation regardless of the actual network topology and end user environment (e.g., Skype).

The use of encryption for obfuscation purposes typically prevents the use of content-based inspection as it is, by definition, scrambled. On the other hand, other analysis methods are still valid and continue to detect such applications.

DPI for ISPs and Carriers

New Sources of Revenue

DPI is currently one of the hottest issues for ISPs and carriers. Constantly seeking new sources of revenue and to reduce churn, they are deploying or forming projects for the deployment of triple play and new services.

This is where DPI comes in:

- To analyze their current network situations and their readiness to receive rich, demanding, consuming and real time traffic.
- To analyze their subscribers' behavior, such as traffic patterns generated per hour/day/week and measure the over-the-top services being used by subscribers.
- To set up global application control policies - such as the total quantity of P2P or VoIP/Skype traffic - at the various peering points where they purchase bandwidth from upstream providers.
- To set up per subscriber SLAs/policies, in order to enforce smarter services, volume/duration-based billing, be more competitive, provide better QoE, and increase ARPU.

"Cable, DSL, satellite, fixed broadband wireless, WiFi, and other network providers are deploying DPI technology to support management of peer-to-peer (P2P) application traffic, protect against a broad range of network attacks, and more recently introduce quality of service (QoS) or prioritized services to generate new revenue streams."

Rona Shuchat, IDC

Solving Bottleneck Issues

Consider what is important to make traffic flow easily on a major city ring-road or reduce congestion in city centers. Enlarge the roads? Install traffic lights and road signals? Set up specific lanes for priority vehicles? Implement cameras, sensors and helicopters to centrally manage collected data and react quickly when an accident occurs? Establish a web site providing real time data on where to go and where not to go? Optimize car loads with a maximum number of travelers? Establish fees for people wanting to drive faster in a reserved lane?

The answer is a combination of all the above and more. It all focuses on bringing intelligence to the subject, in terms of collecting data to understand when a situation is becoming problematic or serious; in terms of creating self-regulations and priorities to control and optimize usage. This analogy is synonymous to IP requirements, which are lacking this essential intelligence. And it is the in-depth data made available by DPI which provides the intelligence to solve and control broadband bottleneck issues in real time.

Guaranteeing Quality of Service (QoS)

The ability to know what is running on the network enables ISPs and carriers to offer new, differentiated services to subscribers and guarantee their QoS. Many wonder whether subscribers are willing to spend more in monthly fees in exchange for this. This naturally depends on the subscriber's profile. Subscribers seeking priority for Skype or QoS for on-line gaming will be prepared to be pay for these services. Other subscribers would like to pay on demand (just like in the mobile phone subscription billing world according to special packages for SMS, frequently dialed numbers and weekend rates).

To gauge the market and offer the appropriate services, ISPs and carriers must have the ability to survey and constantly monitor their subscribers' habits. This ability can only be provided by DPI. Once this data is known, they can segment their subscribers into different types of user groups and put together customized service packages designed to meet the specific requirements of each group.

DPI Market Solutions

The Challenge to Implement DPI Technology

Designing a serious DPI device is a challenge, which is actually a list of multiple challenges:

- The capacity to build a comprehensive library of signatures and behavior.
- The capacity to quickly execute inspections on a nearly real time basis, to introduce the least possible latency, and to perform at multi-gigabit speed.
- The capacity of creating and making available all the necessary counters to feed a database with statistics and usage figures, thereby offering both high-level and detailed views of network activity.

DPI Meets the Market Need

The DPI market landscape is still evolving. This is being fired by a number of factors - primarily the burgeoning number of social applications such as IM, voice, video on demand and network gaming; the proliferation of more instant publications like blogs, wikis and "YouTube" type self-publishing zones; and the constantly growing dependence on the Internet in general, as seen in Google Earth, on-line banking and Web2.0.

These drivers are forcing ISPs and carriers to constantly evolve and compete to facilitate heavy downloads and uploads (video/pictures/backup/ltune/podcasting), real time traffic (voice - telephony), social traffic (IM and gaming), billing by volume and/or duration (pay on usage, pay per view), on-demand services via portals, security, availability and Quality of Experience.

To meet these needs, the industry will have to innovate, particularly in the areas of differentiated services, real time subscriber management and billing, and a lot of "on-demand" possibilities offering "click and start" capabilities. All this can be achieved using DPI.



© Allot Communications, 2007. All rights reserved. Allot Communications and the Allot logo are registered trademarks of Allot Communications. All other brand or product names are trademarks of their respective holders.
