

Modern Day Challenges for Lawful Intercept

By: Karl Wale, Director of Product Line Management

Overview

Law Enforcement Agencies must take into account the full spectrum of communication “applications” including voice, VoIP, E-mail, and the ever-increasing number of social networking, peer-to-peer (P2P), and even gaming applications that can be used to transmit information between parties. Detecting, analyzing, and intercepting content from these different media presents a very large challenge—but DPI is up to the task.

CONTENTS

Three-Stage Approach *pg. 2*

Tackling the Challenge *pg. 3*

In addition to the plethora of new communication mechanisms, the widespread availability and adoption of mobile broadband makes the task of tracing lawful intercepts back to individual subscribers even more complex. This problem will also get more difficult when one considers that mobile networks are looking to offload traffic from their core network as close to the edge as possible (e.g., at a Femto Gateway or at a Serving Gateway in LTE); see Figure 1.

The final dimension to the problem is the sheer amount of data that needs to be monitored in order to find the traffic of interest applicable to the specific target. This is particularly true for intercepts located at peering points and high capacity aggregation points in IP networks.

We should also distinguish between different applications of LI even though they share a great deal in terms of the challenges and, to a certain extent, the technology. Broadly speaking, we can consider LI as applied to criminal investigations and then LI as applied to national security. The requirements and specifications will differ, and some may require more advanced techniques such as “keyword detection.” In fact, with regard to national security, the technology is very flexible and has also been adapted to support cyber security requirements that protect many types of critical networks from attack, such as essential computerized monitoring and management services for electrical power, water supply, etc.

DPI systems use highly specialized algorithms to collect and analyze packets associated with a specific session or flow. Modern LI systems must scale from a few Gbps to many 10’s or even 100’s of Gbps depending on their location in the network. Given that this throughput is beyond the capacity of any single processor or stand-alone device, we have seen a strong move in this space to large bladed architectures such as AdvancedTCA. Not only can ATCA provide the scalability and performance required, it is also best suited to deployment in carrier-grade networks due to its cooling and mechanical specifications. In addition, the wide range of interfaces encountered in carrier networks—including Ethernet and optical links such as OC-12, OC-48, and OC-192—make ATCA an ideal platform choice for LI.

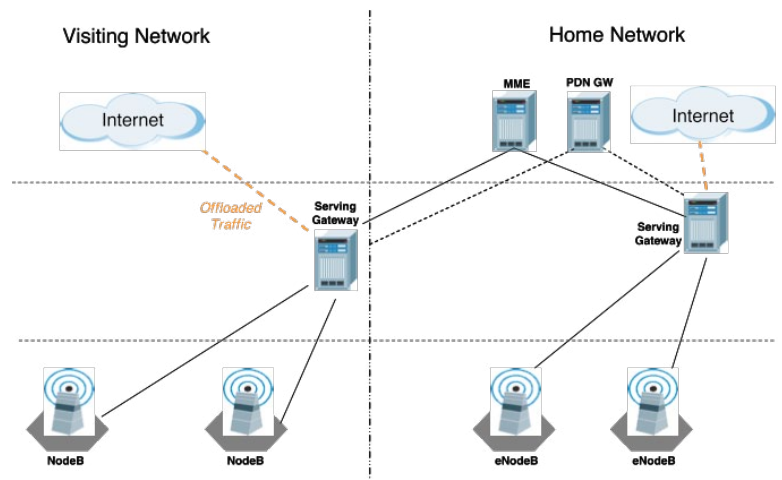


Figure 1. DPI-Enabled Internet Offload

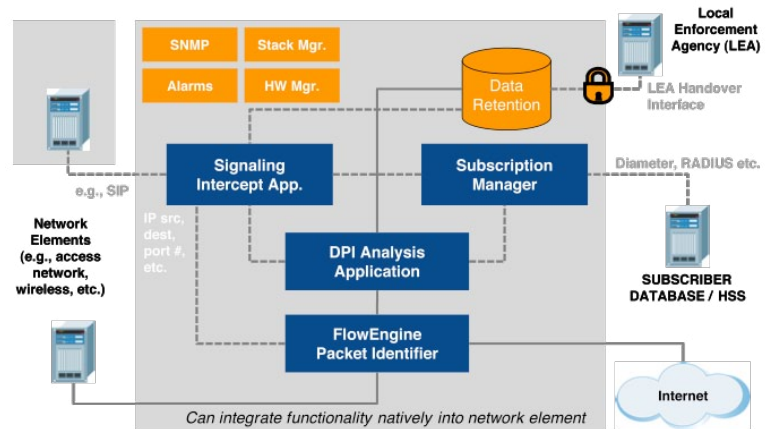


Figure 2. LI Application Overview (Stand-Alone or Integrated)

Three-Stage Approach

DPI-based solutions typically have three distinct processing, or filtering, stages on the traffic—and at each stage the traffic volume is reduced. The first stage is usually a front end load balancer which sees all traffic coming off the network. The load balancer may involve a fairly simple decision process to either forward traffic for further analysis (or not). The key for this phase is that it must be performed at line rate and it is not permissible to miss any packets—hence the desire to keep this function as simple as possible due to the volume of traffic involved. The decision could be based on a specific IP address, a number

of pre-defined IP addresses, or application type. An ATCA switch such as Radisys' ATCA-FM40 can perform this function at 80Gbps, and multiple switches can be combined to scale to several hundreds of Gbps.

The load balancer typically uses both source and destination IP addresses to collect sequences of packets in both directions, and ensures that all packets from a particular user go to the same application processor or DPI engine. In some circumstances one may want this front end decision to take into account more information about the session or content. In this case a packet processing blade can be used in conjunction with the switch at the front end. Such a configuration can consider more complex parameters and even dig into the packet to find specific attributes associated with the applications or session. The output from this stage is spread, or load balanced, across a number of packet processor blades which perform the in-depth traffic analysis.

The second stage is the DPI analysis engine which looks at each of the packets to perform a stateful analysis of the communication setup, handshaking, and how the traffic is passed. By comparing these measured parameters to known patterns and using other analysis techniques it is possible to determine the actual application being used such as a P2P session (e.g., BitTorrent), a VoIP session (e.g., Skype), or even a specific gaming application (e.g., Halo).

Once the application is identified it is possible to extract content or metadata from the session, which can include caller, person called, duration of call, content, and even the equipment used and the contact numbers. In a mobile network this can be used to identify the specific handset, IMEI number, and phone number. As a result, by using DPI it is possible to identify and track communications no matter what media or application is used—be it fixed or mobile.

The third stage of processing is typically more specialized and beyond the scope of this article. Often the final stage is highly application-specific and performed on Intel x86 or other general purpose processors. These CPU blades may also provide

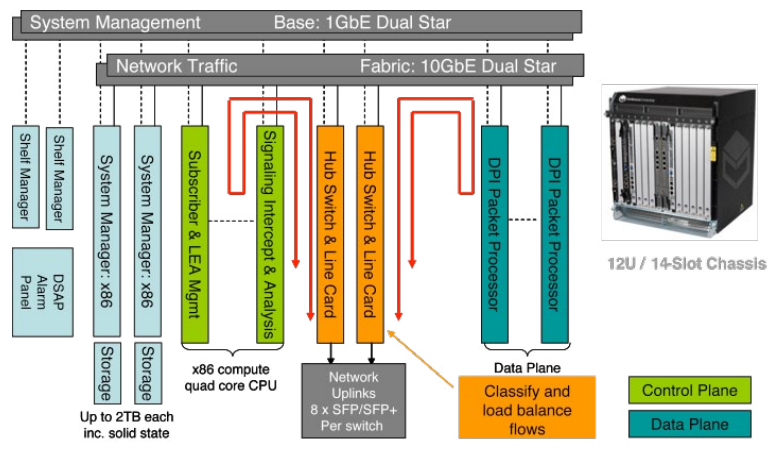


Figure 3. DPI LI System Configuration

other functions in the system such as bulk storage of intercepted traffic as well as provisioning of the industry-defined interface for the transfer of intercepted information to the Local Enforcement Agency (LEA).

Tackling the Challenge

One of the biggest challenges for DPI vendors, particularly for LI, is volume. With traffic shaping, there is only a handful of popular applications being used on the network to generate the majority of traffic, and these are the critical ones to control. With LI the problem is different because one needs to cover many, many more applications, which can result in a signature database covering literally thousands of applications. What's more, this range of applications is constantly evolving, and thus keeping up-to-date is critical. Anti-virus software is a good analogy, where one is constantly running to stay ahead of the threats.

To support this requirement we are seeing an ecosystem develop that includes vendors providing and maintaining libraries of application signatures, often pre-integrated with industry-standard hardware platforms and optimized for specific processors. We are also seeing a move away from specialized processors and hardware-based solutions (including ASICs and FPGAs) to multi-core MIPS-based solutions that are easily programmable for adding new applications and signatures on a regular basis.

Modern ATCA platforms use packet processors such as the RMI XLR732 or the Cavium OCTEON to perform such processing. It is also possible to use specialized offload engines to extract very high performance Regular Expression (RegEx) analysis, which is commonly used for keyword detection. These offload engines are extremely powerful and allow the traffic to be inspected for keywords and their context, which helps to reduce the frequency of “false positives.” RegEx can be performed in specialized processors but can also be achieved using TCAMs hosted on ATCA packet processor blades such as Radisys’ ATCA-PP50. Companies such as NetLogic (which acquired RMI) and Cavium (Nitrox) provide these types of capabilities.

In summary, the role of Lawful Intercept continues to become more complex as technology advances. The speeds involved and range of applications means that DPI is the right technology to help address these new requirements and capabilities. Furthermore, popular concerns about “Big Brother watching”—the aspects of DPI that cause the most controversy when applied to traffic shaping and targeted advertising—are in fact the key attributes that make DPI very attractive to law enforcement and national security agencies.

The Radisys logo consists of the word "radisys" in a lowercase, sans-serif font, with a registered trademark symbol (®) to the right. The text is white and is set against a dark red rectangular background.

Corporate Headquarters

5435 NE Dawson Creek Drive
Hillsboro, OR 97124 USA
503-615-1100 | Fax 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com | info@radisys.com

