WIRE-SPEED CAPTURES WITH PORTABLE DEVICES

by Francisco J. Hens, Vicente J. Bergas

Improvements of storage technology in terms of capacity / speed and continuous optimization of Field-programmable Gate Array (FPGA) integrated circuits are bringing a totally new wave of possibilities in data capture and processing applications. Expensive and complex appliances based on Hard Disk Drive (HDD) arrays are not required any more.

What you will learn:

- Latest advances in storage devices applied to high speed traffic capture.
- How SSD storage compares with traditional storage.
- How to use hardware accelerated packet filtering to capture network traffic.
- Applications of hardware time-stamping to time critical data analysis.
- Types of traffic replay and their applications.

What you should know:

- Basic TCP/IP architecture
- Some basic knowledge about Ethernet
- Fundamental switching and routing

FPGAs are perfectly suited for wire-speed processing of fast data sources and small form factor *Solid State Drives* (SSD) supply excellent performance, large storage capacity and they are perfectly adapted to operation in portable equipments. One of the reasons of migration to SSD is that current cost of storage solutions based on flash memory allows to supersede past objections about price of these devices.

Portable devices based on FPGA and SSD designs can be used in network troubleshooting, data forensics or security related applications when high capture speed and capacity is required but simple, reliable and fast configuration is important. Data capture is combined with the functionality of a network tap to enable easy access to the traffic stream that has to be analysed. Some applications, specially those related with security, require the capture device presence to remain unnoticed while it is connected to the network. Devices can be designed so that they do not modify the traffic information content or timing in any way even if they are connected in “bridged” mode allowing the traffic to be transmitted through the equipment.

Data capture and protocol analysis are related but different functions. Capture has to be fast to be effective but protocol analysis has no real-time processing requirements. A portable capture device may or may not include protocol analysis. Sometimes it is enough to supply the means to enable the user to identify and download the interesting data within the captured stream and leave protocol analysis to dedicated, usually software-based equipment.

The following sections provide details about the applications, features and architecture of portable capture devices. The focus is on functionality that make hardware-assisted capture functionality in portable devices unique.

APPLICATIONS

Portable capture devices are ideal for enterprises looking to ensure that their networks are robust, scalable and secure. Applications of portable capture devices can be distributed in two large families: troubleshooting of communication networks and security. This is not a revolution when compared with traditional applications of any standard capture device. However, scenarios and applications of portable devices are considerably different and broader due to the ability to be connected and start operation in minutes without any special requirement. Portable capture devices are very well suited to temporary network connections in cases where analysis is required only for a limited period of time of usually a few hours or days. A good example would be analysis carried out in a cellular network through connection to one or several mobile base stations.

- Applications related with network troubleshooting include tracing of difficult to assess, temporary, intermittent problems. Traditional monitoring tools provide permanent information about the network in terms of various Key Performance Indicators (KPIs) but they are unable to deal with issues related with unexpected protocol interactions. Full protocol captures arise as the only way to face these problems.
- Portable capture devices are useful fighting against attacks like phishing linked to malware and other security threats. Event based pre-filtering could be used to detect intrusions. With the help of these tools, investigators will have the capability to reconstruct web sessions, e-mails and 'chat line' conversations in a chronological order to investigate security incidents.
- Finally, portable capture devices could be used in *Lawful Interception* (LI) applications. In case of portable devices the focus is again in non-permanent interception. Both filtering based on fixed patterns and event based filters could be used to built efficient LI based on wire-speed captures.

WIRE-SPEED PRE-FILTERING

Pre-filtering is an important feature even for devices prepared for wire-speed capture. With the help of filters, users make sure that only important data is going to be stored. For example, if only IP telephony signalling is going to be analysed, all other data can be ignored. The effect is a much better usage of the storage capacity. With the help of pre-filtering, it is possible to extend the maximum capture time from a few hours or minutes to days or weeks by constraining the raw volume of data. The second advantage of pre-filtering is that it can be used to mark packets depending on the filtering rule applied to match each of them. This classification can be used later for post-filtering and protocol analysis.

Hardware processing is well suited to filter data based on fixed-length packet fields like IP / MAC addresses or class of service (CoS) marks (see Table 1). As a result, it is possible to match any packet directed to an specific IP address, or directed to a network specified by its network prefix, or packets between two hosts specified by their source and destination addresses.

Port based filtering can be used to match traffic from single applications like web traffic (port 80), e-mail (port 25), VoIP signalling (port 5060) and many others. Filtering based on CoS marks can be used to filter traffic classes subject to controlled performance defined by the Service Level Agreement (SLA).

More advanced filtering is based on fixed alphanumeric patterns. Fixed pattern filters can be used to find any word or sentence within the data stream. There are many applications of this kind of filters. For example, an IP telephony trunk link based on SIP signalling use SIP INVITE messages to establish calls. Filtering the "INVITE" word may be used to get information about IP calls occurring in the link.

Table 1. *Pre-filtering modes*

Filter Type	Details
Ethernet Selection	Selection by source and destination MAC addresses or Ethertype field.
VLAN selection	Selection by VLAN-ID or CoS marks. Matching of C-VLAN or S-VLAN fields in frames with multiple VLAN tags.
IP selection	Matching of source and destination IPv4 / IPv6 addresses, DSCP and protocol (UDP, TCP, ICMP,...).
TCP / UDP selection	Filtering of source and destination TCP / UDP ports. Selection of port ranges.
Fixed offset selection	This filter matches an specific bit pattern in a user configurable position within the packet.
Fixed pattern selection	Matches a fixed patten in a variable position within the frame. The pattern is specified as an alphanumeric string.
Length selection	Matches packets with an specific length or frames within a custom length range.

The third filtering mode is based on user defined events. An event is potentially anything it could happen in the network. It could be an error condition like "an errored packet is received" or it could be that a packet from an uncommon protocol is received or that a packet containing a custom alphanumeric pattern in the payload is found in the traffic stream. The difference between event based filters and basic filters is that events are used to modify the filtering rules. The basic action triggered by an event is the transition from "no frame is filtered" to "all frames are allowed to pass through the filter". However, other more sophisticated actions can be imagined in more advanced capture devices. Filtering based on events find important applications in intrusion detection applications or assessment of difficult to trace problems in communications applications.

An essential feature of filtering blocks is that they can be combined to give more complex filters. For example, users can configure various filters within the same block to get the combined effects of an "AND" filter. In the same way, several filtering rules are combined in different blocks to get the aggregated effect of an "OR" filter.

CASE STUDY: VOIP CALL IDENTIFICATION AND TRACING

Portable capture devices are perfectly suited for capturing VoIP media and signalling in exchanges and cellular telephony base stations. VoIP signalling based on SIP, H.323 or other signalling framework carries information about communicating parties like SIP URIs / telephone numbers and media encoding (ITU-T G.711, G.729, etc.). Media streams contain the voice samples themselves encapsulated in an RTP envelope. Capturing signalling could be useful to collect statistics about network usage and to get information about an specific user or a group of users. Media capturing is required for voice quality benchmarking or lawful interception applications. Connection to the network and capture configuration is different for signalling and media captures (see Figure 1).

- *Capturing signalling*: The capture device could be connected to a SIP trunk to make sure information from all users is available. The capture device could operate in pass-through mode but endpoint operation may be preferred if a mirror port is available for monitoring. Most of the interesting information comes in the SIP header, including telephone numbers. Call duration may be inferred from the timing of different signalling messages generated for the call. Using the TCP / UDP port filter is probably the best choice for global signalling captures. SIP proxies use port UDP 5060 by default. Filter scope can be narrowed down to collect statistics from a single user (IP address filters) or signalling messages from an special type ("INVITE", "REGISTER, "BYE", etc.) with the help of the "pattern" filter (see Listing 1).
- *Capturing media*: Capturing media tends to be more challenging than capturing only signalling. Data rates involved are higher than for the signalling case (about 100 kb/s per call, ITU-T G.711 codec). Moreover, signalling information (the SDP payload) is required to decode the media. The most interesting filter for media capturing is perhaps the IP filter to get information from a single location.

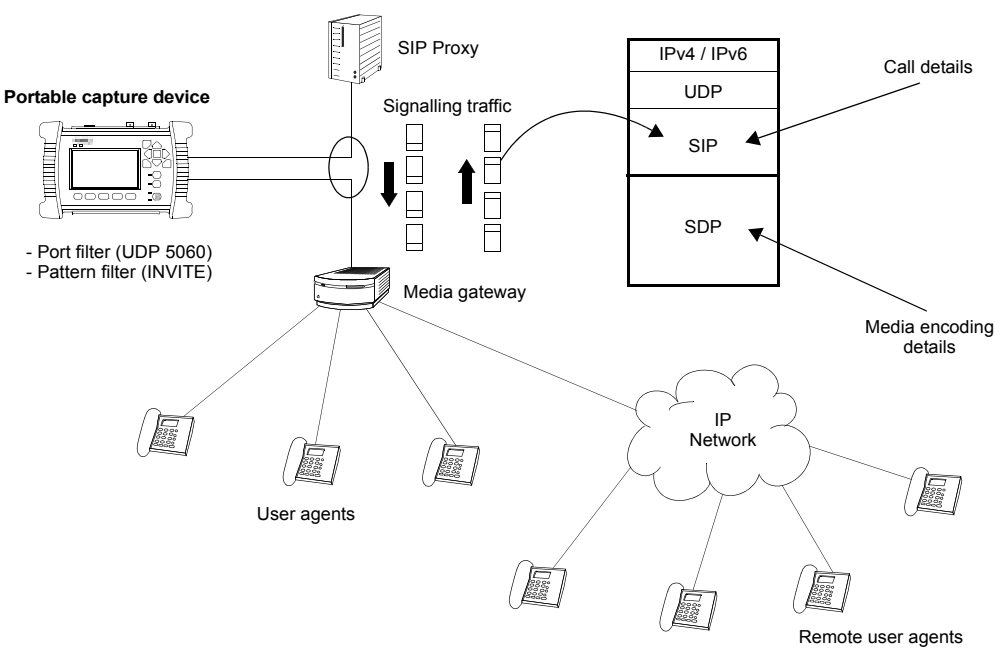

Figure 1. *Configuration of a VoIP signalling capture: The capture device is connected in through mode between the media gateway and the SIP proxy. Applicable filters are the TCP / UDP port filter and the pattern filter. The most interesting header available for analysis is the SIP header*

CAPTURE REPLAY

Traffic replay basically consists in transmitting previously captured data. It is the opposite of traffic capture. Traffic replay is useful to reproduce some network transmission conditions in a controlled environment, usually in a laboratory. Replay is probably not essential in security applications but it is important in network troubleshooting. Portable devices are designed mainly as field tools but their advantages make them suitable for laboratory applications too. For this reason, traffic replay is also a good functionality for them and it should be available at least as an option.

Traffic replay can be *stateless* or *stateful*. Stateless replay can be either timed or not. These are the features, advantages and disadvantages of each:

- *Timed stateless replay*: Packets are transmitted exactly in the same way they were captured. Timestamps are used by the replay equipment to schedule packet transmission at the right time.
- *Not timed stateless replay*: Not timed replay works in a similar fashion than timed replay but information carried by timestamps is ignored. Packet transmission is based on a bandwidth profile statistics configured by the user. Data could be transmitted at the maximum speed allowed by the transmission media but any possible constant, variable or random traffic distribution could be used instead.
- *Stateful replay*: This replay mode is necessary if it is required to keep the interactions between communication parties during transmission. In this case, timing of packet transmission has to be based on events. For example, transmission of the next scheduled packet could wait to the reception of a certain message type from the network. Stateful replay may be used to model interactions between ports of different capture / replay devices. Direct interaction between the capture / replay device and a network entity (server, IP telephone, computer) is the second possibility. Stateful replay is more powerful than stateless replay but the drawback is that is more complex and it is difficult to implement using firmware.

In order to maximize usefulness of replay, the ability to modify the stream while it's been replayed is often required. For example, by replacing source and destination addresses or VLAN tags by user configurable parameters is possible to reuse the same captured data in different test scenarios.

MANAGEMENT AND AUTOMATION

Unlike it happens with large capture appliances, often designed for permanent installation in racks,

Listing 1. *Structure of a typical SIP signalling INVITE message. The message starts with the word "INVITE"*

```
INVITE sip:bob@atsl.com SIP/2.0
Via: SIP/2.0/UDP mkt12.fnetprodoc.es;branch=z9hG4bK776asdhds
Max-Forwards: 6
To: Bob <sip:bob@atsl.com>
From: Alice <sip:alice@netprodoc.biz>;tag=1928301774
Call-ID: a84b4c76e66710@mkt12.netprodoc.biz
CSeq: 314159 INVITE
Contact: <sip:alice@mkt12.netprodoc.biz>
Content-Type: application/sdp
Content-Length: 142
```

Figure 2. *Graphical user interface based on colour screen specifically designed for local management of a portable network capture device*

portable devices are required to be configured locally with the help of an attached keyboard, touch screen or other input device. A dedicated, graphical user interface is very useful for this purpose. (see Figure 2) This solution makes external devices like controlling computers with special management software unnecessary. However, a management interface available through a network interface is still an important feature of portable capture devices. For example, if one or several devices are required to operate in a large system, they need a communication interface enabling coordinated operation between them. Interactions between the management entity and the managed agent could be of three different types:

- *Configuration commands*: These are commands that modify the device configuration in some way. They are required to start / stop captures, configure filters and other operations.
- *Result retrieval commands*: They are necessary to get statistics or other data from the equipment. When the management entity issues a result retrieval command to the managed agent, it is expected that the agent will generate a reply with a response to the previous query.
- *Unsolicited messages*: These messages are generated by the agent without being explicitly requested by the management entity. They are used to signal events occurring in the agent.

SNMP is very well suited to implement all three message types. Unsolicited messages can be implemented by SNMP traps. Most scripting languages like Tcl or Perl have extensions for SNMP. These languages can be used to write scripts implementing complex scenarios involving one or several capture devices and potentially other equipments like network emulators or traffic generators.

ARCHITECTURE

The core of a capture device is a fast FPGA that speeds up critical data process operations. Basically this includes packet forwarding, wire-speed pre-filtering of network packets, hardware timestamping and proper data formatting / storage in a high capacity SSD. Connection between the FPGA and the SSD could be implemented with a mini SATA (mSATA) interface. The minimum SATA speed enabling wire-speed captures in bidirectional 1 Gb/s electrical / optical interfaces is 3 Gb/s. The 6 Gb/s SATA speed is possible as well but is not really required for bidirectional 1 Gb/s captures.

The FPGA / SSD subsystem is controlled by a CPU which is in charge of starting / stopping captures and collecting statistics and status information. The CPU is connected to peripherals that make it possible interaction of the equipment with the real world. Local management could be implemented with keyboards, screens or any other input / output device. Remote management requires the CPU to run the processes necessary to act as an SNMP agent. Finally, the CPU acts as a mediator between the FPGA / SSD subsystems and the external world in capture upload / download processes. Existing captures could be copied to / from

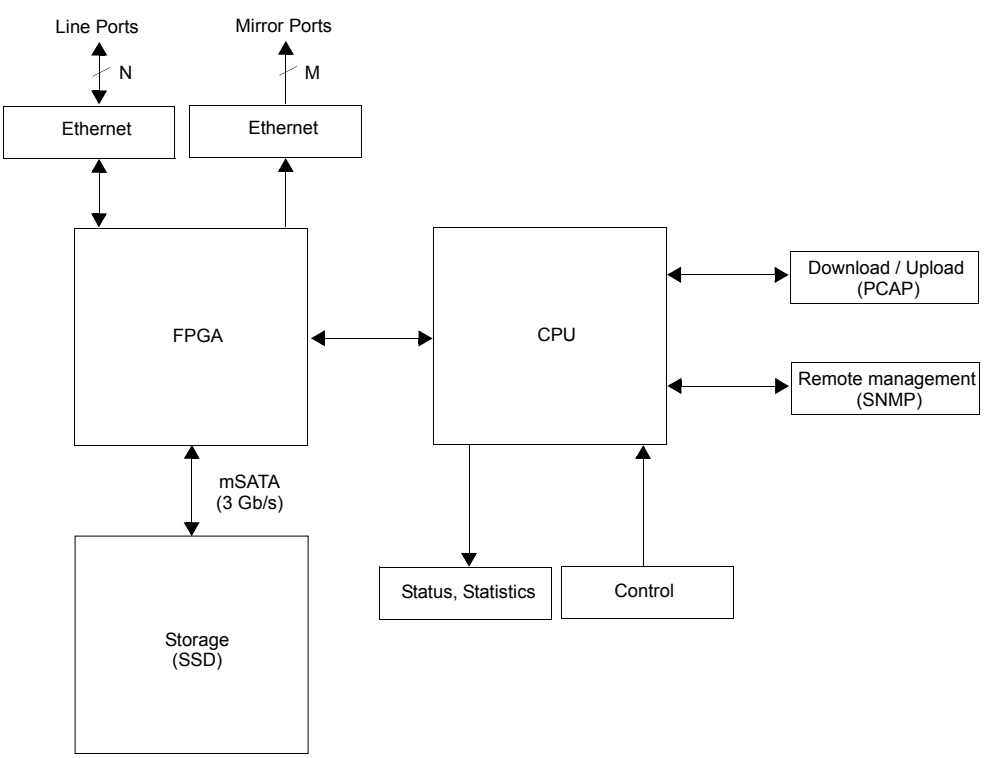

Figure 3. *Simplified block diagram corresponding to a generic portable capture device*

a management Ethernet interface but they can also copied to USB devices like flash memories or disks (see Figure 3).

TAPPING TECHNIQUES

Sometimes, data to be captured is available trough a mirror port but some others is useful to have the capacity to connect the equipment in pass-through mode and forward the network traffic between two line ports (see Figure 4). The port mirroring capability with equipment configured in pass-though mode, is also useful when traffic is going to be forwarded to external analysis devices rather than the internal SSD. In practical terms, port mirroring is used for real-time analysis. Using an external protocol analysis software is possible to display packets as they arrive. The external protocol analysis most likely unable to process packets arriving at high speed. For this reason, in portable capture devices, mirror port usage is often reserved for low speed applications and storage to SSD is used in more general situations.

A feature related with port mirroring is port aggregation which can be configured to aggregate traffic from the forward and backward transmission directions and present them as a single stream. This kind of stream aggregation is useful to check interactions between the communication ends like for example requests and replies in a web application. However, if the aggregated bandwidth is higher than the mirror channel capacity, some frames will be lost.

Network connectivity when operating in pass-through mode is guaranteed by adding batteries to the capture equipment. In this way, the link remains active if the capture equipment suffers from a temporary power shortage. If there is an ongoing capture, no data is lost as long as the internal battery is operative. Current LiPo batteries could last for several hours of continuous usage without strong impact in the overall equipment weight. A second protection against power shortages is accomplished by means power-fail protected tap interfaces which are capable of maintaining a link when they are not powered. Capture data is lost in this case however.

Some applications, specially those related with network security, require the presence of the capturing device to be undetectable by end users. This requirement precludes some of the simpler designs. Specifically, the following are highly desirable features of any network device designed to be undetectable:

- Traffic forwarding based on switching or routing between the line ports is not acceptable. Switching is based on address learning / broadcasting mechanisms and it may also involve special processing of some bridging protocols like the Spanning Tree Protocol (STP) or any of its variations. Routing is even worse because routers require a variable, unpredictable time to process packets. For these reasons, presence of both switches and routers is not difficult to detect.
- It is easy to understand that a network device which must remain undetected does not have to generate any traffic. Even if this is true, some passive devices still generate traffic replying to ARP requests, ICMP messages or TCP handshake packets. Any packet generated by the capture device can be used to detect its presence and it must therefore be avoided.
- Preventing the generation of any data is not enough to guarantee that the device will not be detected when connected in pass-through mode. Store-and-forward processes generate a deterministic delay which depends on the packet length. That means that either store-and-forward has to be replaced with a more advanced forwarding mechanism or the capture device has to carry out some kind of packet delay equalization before forwarding date to the outgoing interface.

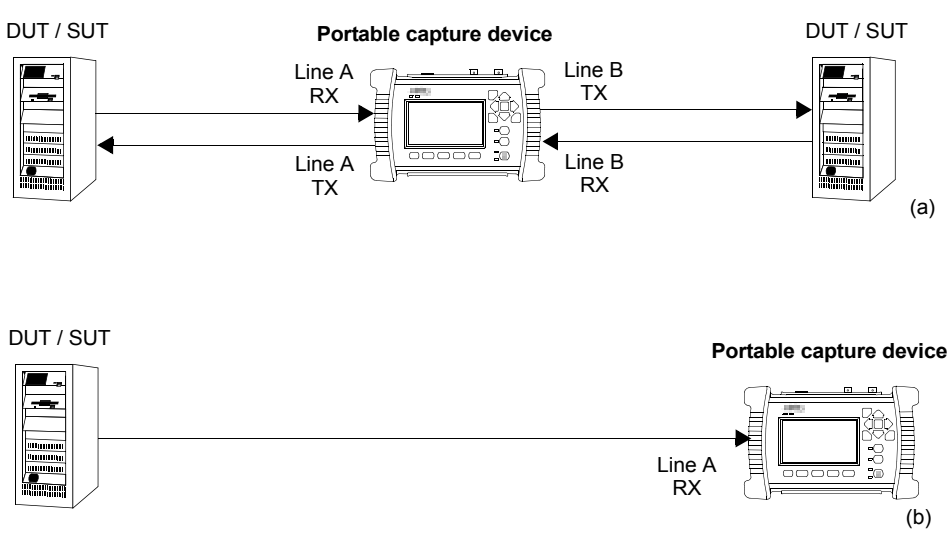

Figure 4. *Connection of a portable capture device. (a) The equipment is connected in pass-through mode. (b) End-point connection*

HIGH PERFORMANCE STORAGE

SSD is currently the best choice for storage of large amounts of data in portable capture devices. SSD storage is faster than traditional HDDs and unlike HDDs, SDDs do not have moving parts which is good in terms of reliability, heating, resistance against vibration and power consumption. Furthermore, SDDs are smaller and lighter than HDDs. SSDs are still more expensive than HDDs but the difference is much smaller today than just three or four years ago. Cost of SSDs is about 0.80 EUR / GB which is roughly twice than for HDDs of standard capacity. Storage capacity of HDDs is also larger than SDD capacity. In case of HDDs, typical capacities range in 1-2 TB while typical capacities of SDDs is still in the range of a few hundreds of GB.

To achieve the required read / write speed in capture applications is important to choose fast storage hardware but is no less critical to properly design the format in which capture data is stored. The most commonly used file systems (FAT32, NTFS, EXT3) provide extended functionality but they are not optimized for maximum read / write speed dealing with potentially very large data blocks.

Finally, being very fast, the SSD is still the slowest component of the capture architecture an thus it is the one that limits the whole system performance. The choice of the remaining equipment elements (and specifically the FPGA) have to be adapted to the SSD specification in terms of speed.

HARDWARE ACCELERATED DATA PROCESSING

Current FPGA technology enables integration of advanced digital signal processing with programmable logic blocks, within a single chip. The ability to concentrate different functions in standard hardware pieces is the key to simpler, more reliable designs and at the same time keep power consumption to the minimum.

In a portable capture device, the FPGA primary task is to provide hardware acceleration to critical operations, including fast read / write operations. FPGA transceivers could potentially work at very high speed of up to about 30 Gb/s but transmission rates much larger than the SSD speed do not add more combined performance to the system. For example in a 2 x 1 Gb/s capture device based on a 4 Gb/s SSD with a 3 Gb/s mSATA bus, a good choice could be a 3.2 Gb/s FPGA transceiver speed which is one of the standard options supplied by FPGA vendors. For this performance grade, components are relatively inexpensive.

A secondary FPGA task is time stamping of received packets. Hardware assisted time stamping is several orders of magnitude more accurate than software time stamping and maximum error could be as low as 10 ns, which is enough to enable analysis of time critical applications. An example is capture of IEEE 1588 synchronization traffic. Speed requirements for IEEE 1588 traffic capture are not strict but timestamps have to be very accurate to be useful.

FINAL REMARKS

Description of portable capture devices has shown that current technology makes feasible to capture full-duplex data in Ethernet interfaces operating at 1 Gb/s using small, highly portable devices weighting no more than 1 kg and operated by batteries

These devices find applications in fighting against security threats, troubleshooting of data and multiplay networks and lawful interception. All these are also applications of traditional appliances but portable devices are cheaper and more versatile. Portable devices are configured locally with the help of a graphical user interface but they may also include interfaces to allow remote management. One attractive possibility is to use SNMP for this purpose.

Availability of cost effective and highly efficient SSD storage devices is the key piece of portable capture device designs. SSD storage is faster and smaller that traditional HDDs and they are perfectly suited for integration into portable devices. Maybe the most exciting fact about SSDs is that there is still a long evolution path for these devices that promise to bring a whole new world of possibilities in capture applications.

Author's Bio

Francisco J. Hens (francisco.hens@albedotelecom.com) is a technology senior specialist at ALBEDO Telecom SL He holds a B.Eng. and an M.A. in Telecommunications from the Universitat Politècnica de Catalunya and 15 years of professional experience in the Test & Measurement sector. He has worked in local and extended area network applications. His areas of interest include most of the technologies currently deployed in the field for triple play, voice, video and data applications: TCP/IP architecture and routing, MPLS, Ethernet and Gigabit Ethernet, SDH, ATM, DSL and Triple Play. He has published articles, white papers and three books about these subjects.

Author's Bio

Vicente J. Bergas (vicens.bergas@albedo.biz) is an electronics engineer working for ALBEDO Telecom SL. He holds a B.Eng in Telecomunications and a B.Eng in Electronics, both from Universitat Politècnica de Catalunya and 7 years of professional experience in electronics design. Most of his experience is in digital logic on FPGAs in the field of high speed serial communications and data processing. He is interested in high end applications demanding hardware accelerated processing. More than seven products in which Vicente has contributed has already hit the market.

Net.Hunter a Tireless Packet Capture

Critical Data
VoIP
IPTV capture
Data Loss
Denial of Service
Threats
Malware
Fatal Errors
Phishing
Protocol Analysis
Hackers
SPAM
Troubleshooting
Forensic Analysis

- Stream-to-disk **packet capture tool**
- Full Duplex **wirespeed** performance
- **No delays**, jitter, or loss to live traffic
- User defined **filters**: MAC, IP, Port...
- Full **Traffic aggregation** (Rx+Tx)
- **Full Tap** to 1000BASE-T funtion
- IDEAL for **Security** and **Forensic**
- GOOD for **Lawful interception**
- **Hand-held,** self contained, batteries
- **Undetectable**: no IP no MAC

ALBEDO

www.albedotelecom.com
info.telecom@albedo.biz