

Silicom SPDK-1.1 Traffic Recorder

Overview

Among Silicom's recent SPDK-1.1 release' features (released on Dec. 2014), wire speed packet recorder is one of the comprehensive solutions. Based on Intel® DPDK, the SPDK-1.1 is Recorder presents an efficient and scalable solution for wire speed packets capture recording.

Traffic recording at line rate from network to disk, with indexing capabilities is a requirement raised every so often by operators, data centers, ISPs, and even enterprise. The goal for traffic recording is to collect raw data for later offline scrutiny, on one hand, or for online read or replay, on the other.

The de-facto standard for captured packets, the *pcap* format, attaches a time stamp label for each frame. This time stamp enables time based indexing of the captured frames. The finer the resolution of the time stamp, the better it may serve for more and more use cases, such as financial trading, where real time events could mean change of prices.

Generating a time stamp is more complex task as resolution gets finer. Standard OS can cope with microsecond's resolution of packet arrival time, but such stamping is at best, an approximation for true arrival time. NIC based time stamping is far more true to the actual arrival time of a frame. Silicom's time stamping solution does just that. It stamps the packet upon arrival to the NIC, at 8 nanoseconds (eight billionth, 0.000000008 of a second) resolution.

Disk access speed to uphold up to 20Gbps traffic rate is a mandatory requirement from such system. The reason is simple enough. Recording live 10Gbps link is often done by tapping the link, and sending a copy of it to a recorder. The total worth of bandwidth to be recorded in this case is the total of the uplink 10Gbps **plus** the downlink of 10Gbps, both accumulates to 20Gbps.

Two requirements, thereof, are to be fulfilled. One, is an efficient storage access logic. Storage managing software has to be efficient, and robust. Second, storage hardware has to be fast.




Figure 1- Tapped bi-directional traffic accumulates to doubled total bandwidth

SPDK-1.1 recorder follows the philosophy of the SPDK package as a whole. It is based on Intel® DPDK and uses standard COTS hardware components. Silicom's line of Intel® based NICs, with time stamping hardware or without, all work natively with Intel® DPDK, and on the same token with SPDK.



Figure 2 - Silicom PE310G2TSi9P Time Stamp Network Adapter Employs Intel® 82599 MAC

SmartSilc

Solutions by Silicom Ltd.

The storage requirements to sustain line rate of traffic recording are met straightforward with available COTS hardware such as **RAID 0** or **RAID 5** systems, or with solid state storage devices.

Packet streams recording with SPDK-1.1 is done using *pcap* formatted files, with scalable log rotation mechanism, assisting with the ease of management of the resulting data collection.

SPDK-1.1 Recorder

Running on standard COTS x86 serve, the SPDK-1.1 Recorder brings wire speed recording capability as a simple straight forward implementation over COTS hardware.

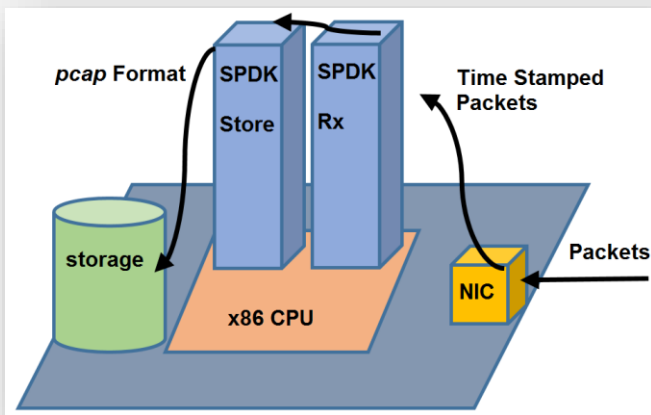


Figure 3 - SPDK-1.1 Recorder - Simple and Efficient using COTS Hardware

SPDK-1.1 Recorder smartly uses DPDK based processing threads to absorb wire speed traffic and to transfer it neatly to disk.

The use of SPDK based utility easily enables further system architecture around the core capture capability. Taking advantage of the stable SPDK framework, it is possible to add, for instance, management mechanism, or to scale up with more CPU cores to the optimum, to develop a comprehensive solution.

Benchmark Results

Testing the SPDK-1.1 Recorder capability was performed with standard off the shelf hardware. The platform was based on Intel® Xeon E5 v2 (Ivy Bridge)

system, with Silicom's Intel® based time stamp adapter and RAID 5 setup.

Two sets of tests were done. One using single ingress port, and a second, where two ingress ports were used to capture traffic in. Each test included varying packet sizes.




Table 1 – Single Port SPDK-1.1 Recorder Benchmark

Single port capture reaches 10Gbps line rate. CPU usage in the case that presented in Table 1 leaves a lot of headroom for additional processing.

Summary

Integrating SPDK-1.1 with commercial off the shelf hardware enables quick and easy assembly of a capable packet recording device. Intel® DPDK kernel bypass approach proves once again the superiority of efficient polling mode, in order to reach line rate processing, at 10Gbps and beyond.