

Application Tap™ - Overview



Author Adam Woolhouse
Date 24th December 2009
Version 1.2

Contents

Contents	2
Revision History	2
1 Introduction.....	3
1.1 The Application Tap™	3
1.2 TipOff™	4
2 Application Tap™ Integration.....	5
2.1 Application Tap™ Support.....	5
2.2 Application Tap™ Transmission Addressing	5
3 References	6
4 Contacting TS-Associates.....	7

Revision History

Version	Date	Name	Description
1.0	22/12/2009	Adam Woolhouse	Document created
1.1	23/12/2009	Henry Young	Review amendments
1.2	24/12/2009	Henry Young	Development & product management feedback

1 Introduction

This white paper assumes that you are new to the TS-Associates (TS-A) Application Tap™. It introduces the key concepts of the Application Tap™, and provides initial guidance on how to use it. The paper is intended as a high level overview. More in-depth technical documentation is available to facilitate planning, installation, development and operation – see Section 3 for details.

The Application Tap™ is normally used in conjunction with TS-A's monitoring appliance – TipOff™. Both products are introduced in this section.

1.1 The Application Tap™

The Application Tap™ is a PCI Express (PCIe) card that enables the implementation of precisely time stamped software instrumentation with negligible performance overhead. The product is suitable both as an enabler of production systems monitoring, and for a wide range of software performance optimisation tasks such as deterministic profiling.

The traditional approach to high precision, non-invasive monitoring of distributed systems, such as market data and trading systems, has been to capture and precisely time stamp network traffic, decode protocols, and analyse data flows, yielding latency and performance metrics. However as architectures evolve from highly distributed systems towards consolidation onto multi-core servers, less inter-process communication is physically available to capture on the network. At the same time, it is becoming increasingly important to be able to monitor software events occurring in the depths of applications, such as order matching engines, and to emit data associated with events, precisely time stamped, in a manner that imposes negligible performance overhead.

The Application Tap™ solves the challenge of minimally invasive software instrumentation by providing a user mode API that enables software events within applications to be instrumented with minimal overhead. This is done with the aid of an FPGA based co-processor and precise hardware clock. Instrumentation code added to applications is able to pass instrumentation metadata to the co-processor which time stamps each event to 10ns resolution, and forwards the instrumentation metadata to an external monitoring application or appliance, such as TipOff™, using an on board gigabit network interface. Clock synchronisation is supported using either PPS or PTPv2.

The Application Tap™ can be applied to a wide range of use cases, including:

- Software latency monitoring – in conjunction with TipOff
- Event logging with low overhead, high precision
- Deterministic code profiling
- Device driver instrumentation
- CPU process/thread stats (requires instrumented kernel scheduler)

The first use case above is the exclusive focus of the remainder of this white paper.

1.2 TipOff™

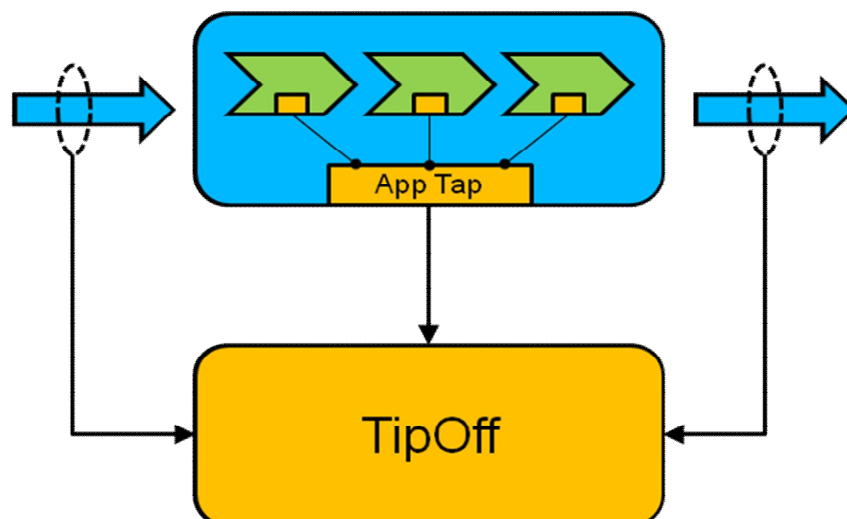
TipOff™ is a real time network monitoring, middleware analysis and passive latency monitoring appliance that supports all standards based and vendor proprietary middleware stacks together with a growing repertoire of direct exchange price and order entry feeds.

TipOff™ supports passive latency monitoring of market data and transaction flows. By decoding packet streams in real time through all layers from network up to application message content, TipOff™ can monitor both packet and message latency, accounting for retransmission latency in the reliability layer. A single multi-hop flow may comprise several different middleware stacks or protocols. TipOff™ supports flows that undergo transformations through order aggregation, trading engines and execution venues.

TipOff™ supports the leading edge approach to high precision, non-invasive monitoring of distributed systems, such as market data and trading systems. TipOff™ captures and precisely time stamps network traffic, decodes protocols, and analyses data flows, yielding latency and performance metrics. However as architectures evolve from highly distributed systems towards consolidation onto multi-core servers, less inter-process communication is physically available to capture on the network. At the same time, it is becoming increasingly important to be able to monitor software events occurring in the depths of applications, such as order matching engines, and emit data associated with events, precisely time stamped, in a manner that imposes negligible performance overhead.

While TipOff™ is able to monitor the latency using network packet capture techniques, the Application Tap™ is able to monitor events in processes, threads and even function calls, transmitting this information *out of band* to TipOff™ for correlation, analysis, storage and presentation.

The schematic below shows TipOff™ and the Application Tap™ used in conjunction to monitor both data flow latency across a server, and also hop by hop software component latency (process, thread or straight line code latency).



2 Application Tap™ Integration

The following is the typical procedure for Application Tap™ integration with TipOff:

- Install Application Tap™ card in an eight lane PCIe slot, and fit appropriate fibre or copper SFP transceiver into port cage 0.
- Confirm operation of Application Tap™ with test application, and packet capture device connected to port 0. Note that TipOff™ can be used for this purpose, or any other capture device.
- Add suitable instrumentation code to target application, making calls to the Application Tap™ API to emit instrumentation messages.
- The instrumentation message flow generated by the Application Tap™ may be captured and processed in any way by the user. However a standardized instrumentation message format that can be decoded by TipOff™ will offer flexibility and enhanced support from TS-A. Example instrumentation message formats can be obtained from TS-A.

2.1 Application Tap™ Support

The Application Tap™ is currently supported for 32 and 64 bit Linux 2.6 kernels only. Support for Solaris and Windows may be added in the future.

2.2 Application Tap™ Transmission Addressing

The Application Tap™ has 4 x 1GE SFP transceiver ports. These ports are actually used by the same hardware running different firmware, for capture and precise time stamping. In this mode, the card may be used as a 4 port full line rate 1GE capture and precise time stamping card. However, with the Application Tap firmware, port 0 only is used for transmitting time stamped instrumentation data in UDP packet format.

With the current release of the Application Tap, there is no default MAC address associated with the transmission port, and no ARP capability implemented. Although a MAC address can be set via the Application Tap API, the default is 00:00:00:00:00:00. This has an impact on the possible deployment modes in terms of the addressing of transmitted packets.

If the Application Tap™ is directly connected to a TipOff capture port then there are no constraints on addressing. However, if the user wishes to connect the Application Tap™ to a management network via a switch port, and then on to TipOff or another monitoring appliance/application, currently it is only possible to use broadcast or multicast addressing.

3 References

This section to be done ...

4 Contacting TS-Associates

TS-Associates Customer Support

www.ts-a.com

help@ts-a.com

+44 (0) 20 8020 2460

TS-A may also be contacted during UK office hours:

info@ts-a.com

+44 (0) 20 7415 7028

TS-Associates plc
40 Bowling Green Lane
London EC1R 0NE
UK