WHITE PAPER

Dynamic Flow Mirroring with the Pensando Distributed Services Platform

Introduction

For data center administrators, being able to analyze traffic to help quickly resolve issues is critical to maintaining business applications. Technologies such as port mirroring (SPAN/RSPAN/ERSPAN) and inline network TAPs have various limitations depending on the implementation. Port mirroring technologies that are often implemented at the switch level have the following limitations:

- **Performance**: Activation of SPAN can cause production traffic degradation (jitter, latency, throughput)
- **Scale**: Limitation on the number of taps within the particular switch
- **Visibility**: Many switch vendors only support uni-directional traffic capture

As a result, the network administrator has a limited ability to capture and SPAN all traffic on the server and switch ports, leading to more difficult network diagnostics and troubleshooting.

To mitigate some of these limitations, some vendors provide hardware inline taps with a "bump in the wire" solution that will capture all of the data. Such solutions provide partial relief from some of the limitations but are very costly – TAPs typically costing 100s of dollars per link. Furthermore, such a solution doesn't guarantee that all of the data is captured due to potential buffer overflow within the vendor's traffic monitoring solution.

Today's Solutions

Network switches are designed for efficient forwarding; they are not built to capture traffic at line rate on all of the ports. More so, the resources including buffers, and queues allocated for mirroring constitute only a fraction of the overall switch capacity. Consequently, they cannot perform a bi-directional mirroring function at full scale and at line-rate performance. This is to say that switch vendors have designed the packet capture feature primarily to troubleshoot a flow or specific traffic type, *not* to capture the entire conversation that can be used for a deeper analysis. Even for debugging, captured traffic is given lower priority so as to not disturb the production traffic, thus further compromising the ability to get reliable packet capture from the network.

The problem compounds when this traffic is sent over the same production network, making production and mirrored traffic share the same infrastructure and thus become indistinguishable.

To summarize, the packet capture capability in the switching infrastructure doesn't scale, doesn't guarantee isolation, and doesn't ensure reliable delivery of the mirrored traffic to a designated collector.

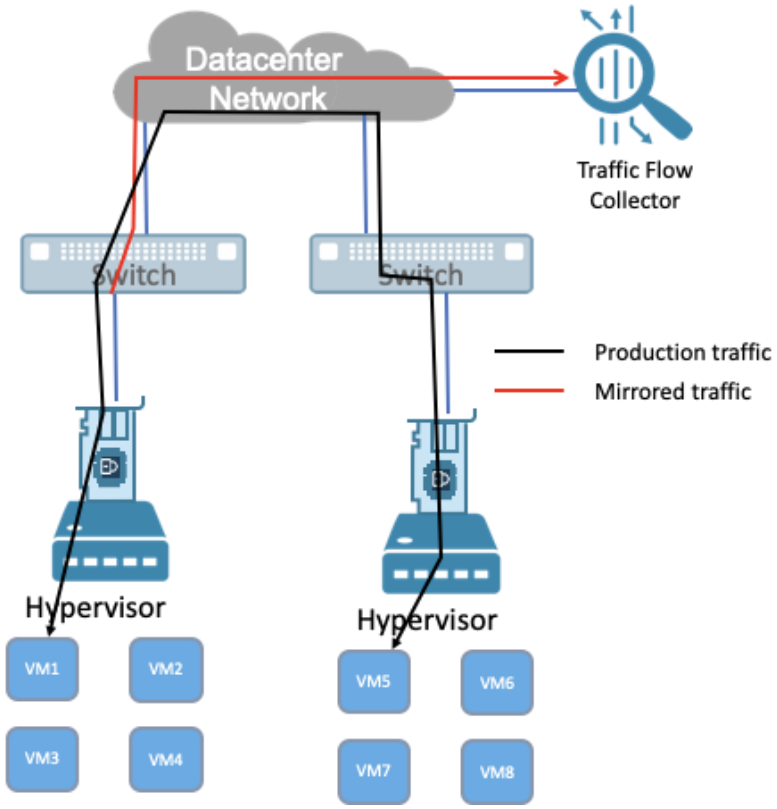

Figure 1: Today's Network Tap Solutions

The scalability of the mirroring solution from the switch vendors also limits the number of mirror sessions that can be established on a per-switch basis that is typically no more than a few 10's of concurrent sessions. This hinders the visibility needed to analyze a typical data center environment today, where multiple source and destination traffic flows are spread throughout the data center.

When analyzing application traffic for a specific workload, it is critical to see **all** incoming and outgoing traffic (Tx/Rx) from that workload. Switch vendors have a very difficult task in capturing both Tx/Rx for a specific workload. Other vendors that provide an inline tap between the workload and the switch do provide that capability, but it comes at a very high cost and with limited scalability. In addition, because of potential buffer overflow, the inline tap cannot guarantee that all traffic will make it to the destination.

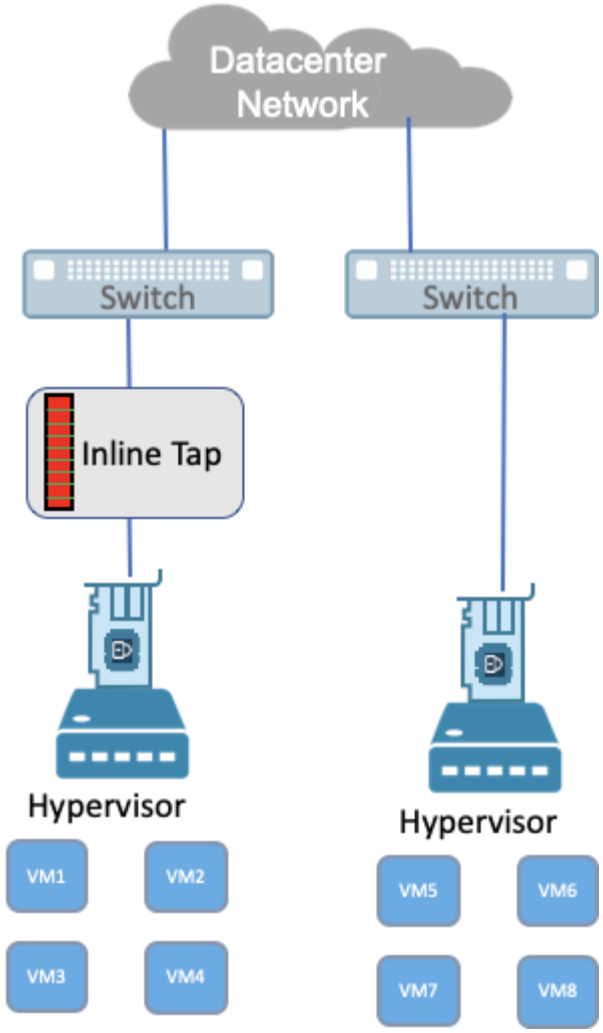

Figure 2: Inline Tap Solution

The Pensando Solution

With Pensando's Distributed Services Cards (DSC's) highly scalable architecture, mirroring of the traffic flow is made possible without affecting the performance of the production data flow. The Pensando DSC allows production traffic to forward as usual while providing the necessary internal bandwidth within the DSC to mirror that traffic flow to any remote collector.

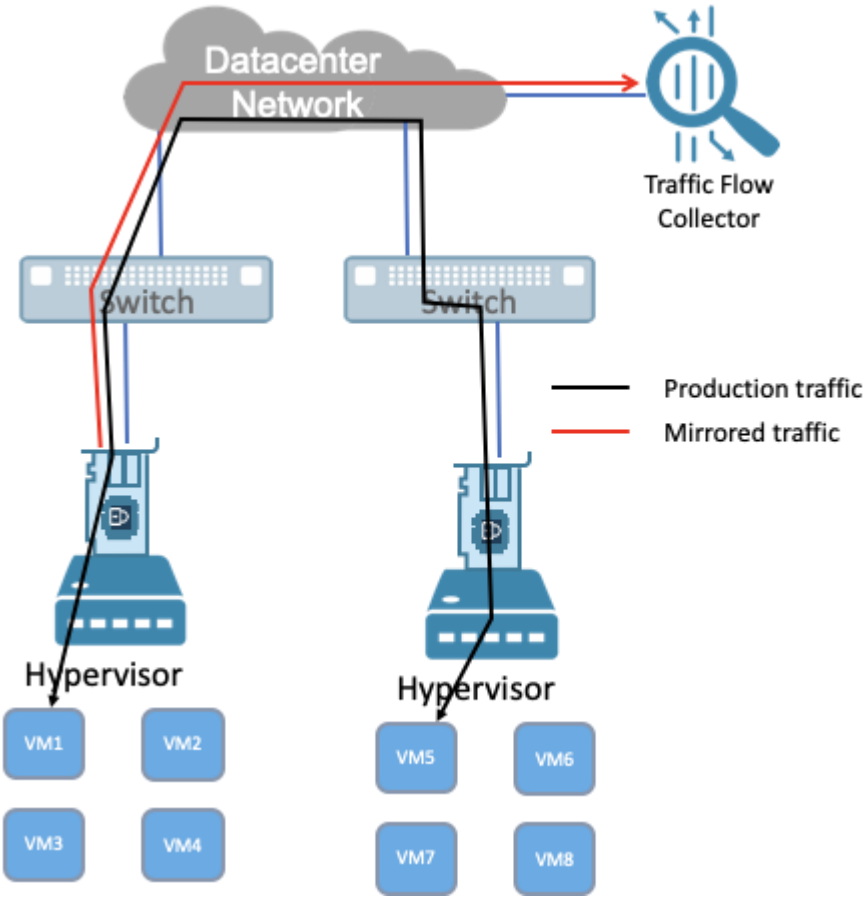

Figure 3: Pensando's Dynamic Flow Mirroring Solution

The mirroring sessions are confined within the DSC infrastructure, enabling a very scalable design with the support for 1,000s of DSCs in the solution. Complete and reliable capture of traffic flows regardless of their location within the data center can provide much deeper and much more scalable analysis of application data.

Pensando's DSC offers up to 8GB of memory buffer to ensure reliable packet capture and reliable delivery, handling traffic bursts or mega flows without affecting production traffic. With bi-directional (Tx/Rx) mirroring support for every workload, Pensando's Distributed Services Platform is uniquely positioned to assist in quick analysis of application data flows across the enterprise.

Pensando's Advantage

One of the biggest operational challenges during troubleshooting using traditional ERSPAN is to locate and turn on the packet capture accurately at the right place in the network. This normally requires the cumbersome process of manually mapping application identifiers to IP addresses, then identifying the VM to switch connectivity, enabling ERSPAN on switches, setting up a collector, visualizing the captured data, and finally turning off the packet capture session. Often it takes 10s of minutes to hours to do this process.

Pensando's Distributed Services Platform with the Policy and Services Manager (PSM) provides visibility of the location of all of the workloads (bare metal and virtualized) within the network, eliminating the cumbersome process of locating the workload before relevant flows can be captured.

With the PSM's integration with VMware's vCenter® and Kubernetes, the flow mirroring session can be initiated based on the VM name, VM Tags, or a Kubernetes Pod/Service name regardless of the workload location. Pensando's PSM also exposes the ability to initiate mirroring sessions based on traditional attributes such as 5-tuple or MAC address within the packet. This reduces the time to capture and visualize the traffic of interest from hours to a few seconds, significantly improving the operational efficiency.

With the ability to assign labels to each of the workload, users can now label (or tag) workloads that belong to a particular application. This provides the dependency of the workloads for the particular application. Once that is assigned, the PSM can automatically capture all of the workloads that belong to that application (label) and enable the mirroring sessions. This makes the process of production application analysis much more operationally simple.

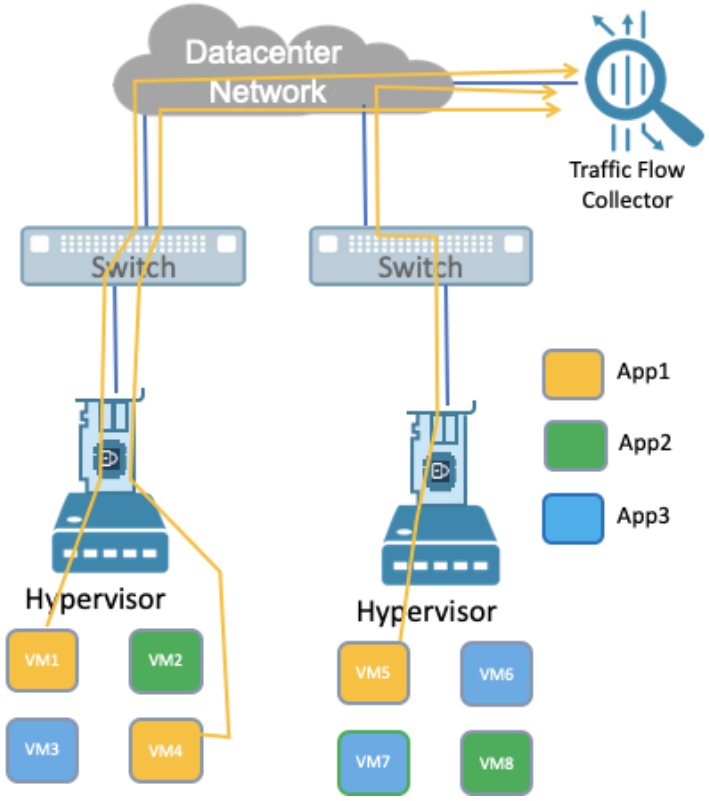

Figure 4: Pensando's Dynamic Application-Based Flow Mirroring

With a simple single configuration, Pensando will be able to mirror only workloads that are tagged as "App1" regardless of location to quickly isolate issues as shown on Fig. 4 above. A sample operational troubleshooting process is shown below.

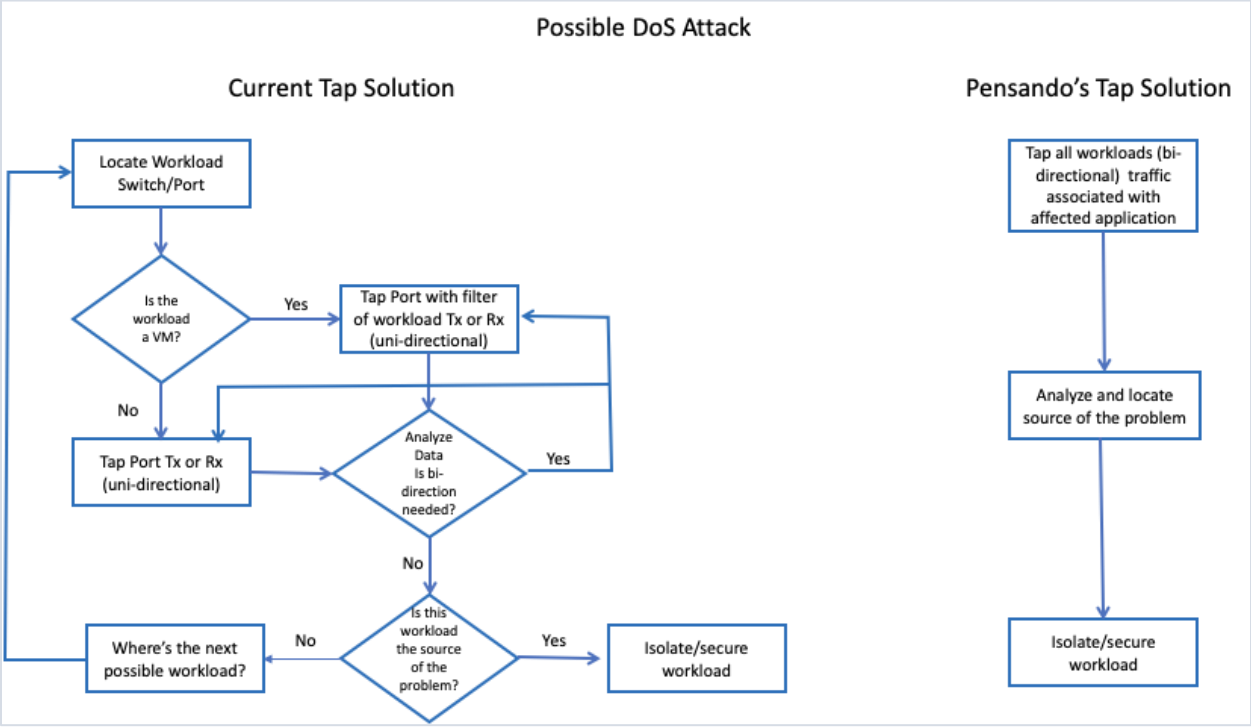

Figure 5: Troubleshooting Operational Tasks Flow

Note: Due to the serial nature of the traffic capturing in existing solutions where first Tx, then Rx traffic belonging to particular workload needs to be tapped and analyzed before moving to the next one, issue resolution may take hours.

To further help analyze traffic flows, Pensando provides additional standard capabilities within the mirroring feature set, including:

- Timestamp capabilities for every packet for consistent and precise time analytics
- Further filtering of traffic flows based upon protocol and other attributes
- Ability to only capture certain portion of the packets belonging to a particular traffic flow, where capturing of the full payload is not required.

Summary

Pensando's Dynamic Flow Mirroring solution allows a customer to quickly and easily resolve problems in their data center. The traffic mirroring sessions are created instantly at source, without the need for special cables and without disrupting traffic. The mirroring rules are specified using logical names (e.g. a VM name or an application tag), freeing the user from tedious physical addresses. In addition, system-wide analytics provide insights to the application behavior across the whole data center. The key benefits are:

- Zero performance hit on production traffic
- Application location awareness
- Operational Simplicity
- Complete Visibility, capturing bi-directional data (Tx/Rx)

By design, Pensando's architecture is distributed and scalable. A federation of DSCs managed under a "single pane of glass" of the Policy and Services Manager (PSM), helps to eliminate the cost, scaling and

administration challenges found in prior solutions. In particular, Pensando's Dynamic Flow Mirroring solution will reduce the need for costly network TAP appliances in data-center infrastructure, and eliminate the need to reconfigure top-of-rack switches and apply traffic spanning sessions. Pensando's Dynamic Flow Mirroring is the simplest, easiest and most efficient way to analyze traffic in the enterprise data center.

About Pensando

Founded in 2017, Pensando Systems is the company pioneering distributed computing designed for the New Edge, powering software-defined cloud, compute, networking, storage and security services to transform existing architectures into the secure, ultra-fast environments demanded by next generation applications. The Pensando platform, a first of its kind, was developed in collaboration with the world's largest cloud, enterprise, storage, and telecommunications leaders and is supported by partnerships with Hewlett Packard Enterprise, NetApp, Oracle, IBM, Equinix, and multiple Fortune 500 customers. Pensando is led by Silicon Valley's legendary "MPLS" team – Mario Mazzola, Prem Jain, Luca Cafiero, Soni Jiandani and Randy Pond – who have an unmatched track record of disruptive innovation having already built eight $Bn+/Year businesses across storage, switching, routing, wireless, voice/video/data, & software-defined networking. The company is backed by investors that include Lightspeed Venture Partners, Goldman Sachs and JC2 Ventures.

For more information, please visit www.pensando.io