

March 2014, HAPPIEST MINDS TECHNOLOGIES

Network Packet Monitoring Optimizations in Data Centre

Author

Dharmraj B Jhatakia



Copyright Information

This document is an exclusive property of Happiest Minds Technologies Pvt. Ltd. It is intended for limited circulation.

Abstract

Network Packet Monitoring is a critical function in a data center. From doing recording of VOIP calls, performance audits, and applying forensics – the use of Networking Packet Monitoring are varied.

The modern day data center throws newer challenges for monitoring systems – like virtualized workflow, emergence of 10, 40 and 100G fabrics, and a very high East West traffic flow and thus forcing investment of expensive Network Packet Brokers in each segment in a distributed approach or compromising the available bandwidth in an aggregated approach.

That Software Defined Networking has revolutionized the way the Data Centers are managed is a known fact. Bandwidth of Demand, Network Slicing, and Rapid Service Chaining are some of the use cases where SDN has emerged powerful. This paper attempts at examining possible solutions and techniques, which can be used to face and overcome those challenges with dynamic programming of network and leveraging the SDN controller intelligence and cost effective bare metal switches to offer a robust and cost effective solution for Monitoring. Not only that, the power of SDN offers additional features which enable reducing in the replicated traffic, distributed monitoring functionality and very granular access control.

Further, the availability of Big Data platforms and Analytics further help in applying predictive orchestration of the monitoring solutions, thus optimizing on the scaling needs and at the same time giving accurate visibility to the administrator.

This white paper would give a brief overview of the monitoring technology, the challenges of the modern day data center and possible solutions / enhanced features that can be achieved through SDN applications and analytics.

Network Packet Monitoring Overview

There are varied business cases for monitoring a network. There are certain functional, security and law requirements which necessitate the need of monitoring solutions in the network.

Some of the typical functional use cases include recording of conversations. For e.g. recording of an executives conversation with the customer to provide feedback / improvement suggestions OR for providing training to new joiners.

On the other hand, a typical security requirement is Intrusion Detection and prevention systems which are deployed to protect the network from inside / outside security threats and breaches.

A telecom service provider would have to comply with the lawful interception rules wherein the conversation between suspects is tapped and the information sent to the law enforcement agency.

All of the above use cases would need the deployment of a specialized Intelligent Network Packet Monitoring solution (or as Gartner has coined the word), a Network Packet Broker.

Deployment of a Network Packet Broker in Data Centers

A network packet broker can be deployed inline in the network or it can be deployed out of band leveraging the SPAN ports. The NPB can either be centralized or distributed with each NPB performing a different function. The figure below attempts to show case these deployment scenarios.

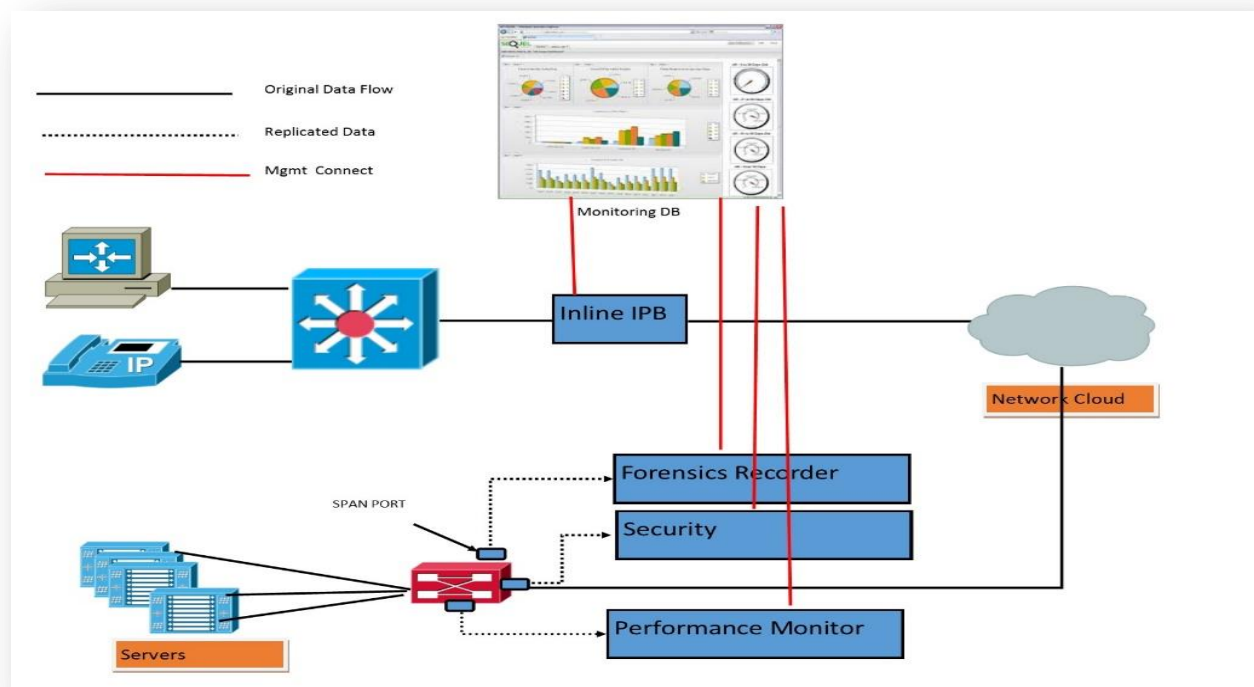


Figure 1 NPB Deployment View

There are multiple ways in which the NPB can be deployed. One method of classification is inline, wherein the NPB sits in the path of the traffic and performs certain functions. This method is suitable in deployment where the throughput needs are not very high and the application are not latency sensitive. However, in deployments where there is high throughput and low latency requirement (like trading systems) an offline method is chosen, wherein the data packets are mirrored on the SPAN ports and sent to the device which is sitting on the side rather than inline.

Functions of a typical Network Packet Broker

- Packet Filtering – Analyze and store only those packets which are needed by applying packet matching rules.
- Packet De-duplication – Remove the duplicate packets that are being monitored
- Web Security – Prevention of attacks as well as monitoring of data leaks to the internet.
- Forensics – for recording of transactions.
- VOIP Recorder.

Types of TAPs / SPAN Ports

Different businesses and deployment scenarios have different needs. Today's switches are programmable and offer the flexibility of having mirroring configured in various ways. Two prominent modes are:

Aggregation Mode:

Consider a small enterprise which just has a handful of devices connecting to the network. The only need that the enterprise might have is single Packet Broker. In such a case, this enterprise can choose to have a configuration where most of the ports on the switch are used for normal packet processing and there is 1 port designation as the SPAN port / Mirror port for all the other ports. Such a configuration is called the Aggregation mode where traffic flowing through multiple ports is aggregated and sent to the SPAN port.

Regeneration Mode:

The other extreme end of the Aggregation mode would be the regeneration mode. Consider that there is a large enterprise which has very high throughput needs which also needs to have very robust set of features including Performance Monitoring, security, forensics, intercepts etc. In such a scenario, the mode of deployment would be Regeneration Mode wherein each of the packet processing port would have multiple mirrors, each mirror leading to a separate NPB.

As one would have guessed, today's programmable switches offer the flexibility of having Aggregation / Regeneration and hybrid configurations. However, they still offer certain challenges which are described in subsequent sections.

Challenges of today's data centers

Trade Off between comprehensive solution and cost

Essentially, there are two choices one has. Either deploy a NPB in each of the Segment (or for each of the Top of Rack Switches) and achieve comprehensive monitoring -this of course gives complete visibility but results in increased cost. The magnitude is much higher if there are multiple functions that the broker has to perform. Or selectively place the NPB in some of the segment and thus risk the lack of comprehensiveness.

Virtualized Work Flow

More and more of the work loads are virtualized. The VMs from which the traffic is being monitored move across the racks and hence get the connectivity to a different switch. The configuration updates to change the monitoring post such movements are complex when monitored and done manually.

Contention of SPAN ports

The number of SPAN ports on a switch are limited. Flexible switches can allow additional ports being converted to Mirror Ports, however, this results in reduced through put. The applications of packet monitoring are increasing day by day and hence this creates a larger demand than the supply of SPAN ports.

Identification of Correct Mirror Point

Perhaps one of the significant challenges that today's data center pose is of identifying the correct mirroring point in the scenario of EAST – WEST Traffic, i.e. the traffic that flows within the data center. For North South Traffic, i.e. the traffic coming in and going out of the Data Centre, this challenge is not there, as we can enable the SPAN at the data center entry / exit point, since that would be a single point through which all North – South traffic would flow. Day by day, the amount of East West Traffic is increasing and hence optimization of correct mirroring point can reduce the duplicate traffic flowing in the data center network.

Reducing the “mirrored” traffic

All said and done, the traffic that is actually processed for performing monitoring functions would essentially be the mirrored or duplicate traffic. Every passing of this traffic in a particular segment thus reduces the available bandwidth for the “real” traffic. At the same time, the systems performing advanced monitoring functions have to process loads of unnecessary traffic thus reducing the efficiency of such systems.

Deep Packet Inspection (DPI) and Software Defined Networking (SDN) to the rescue

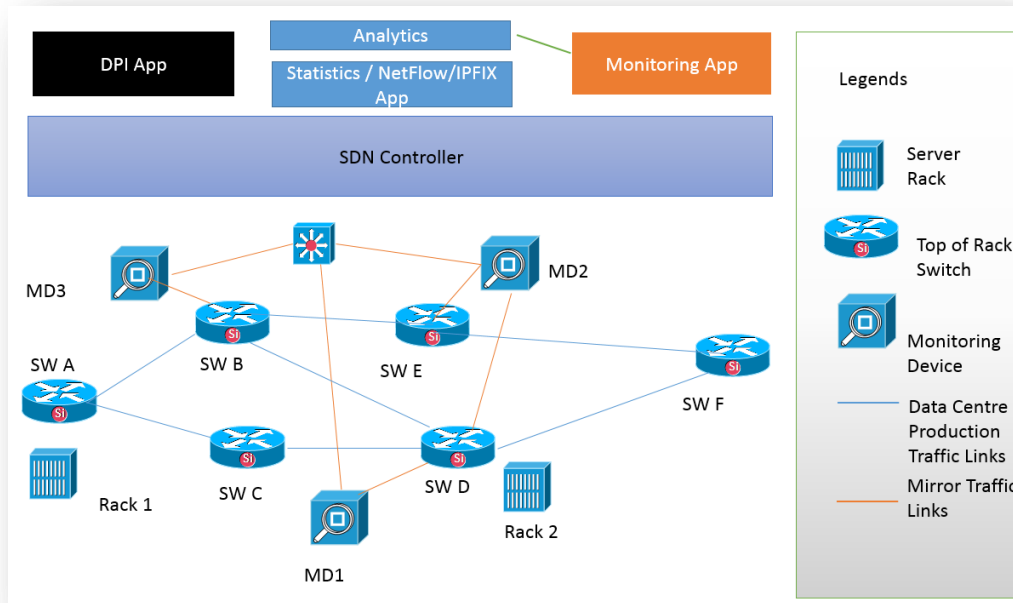


Figure 2 SDN Monitoring Solution

Consider a sample deployment topology as shown in Figure 2. Here, all the data center switches are connected and the blue links indicate the production traffic link. The orange links indicate the Mirrored traffic links that are connected to the monitoring devices. All the monitoring devices are connected to the monitoring switch, which is also controlled by SDN controller.

In a traditional network scenario, such a network would have resulted in a homogenous network with no difference of the production traffic and the monitored traffic and some of the network links would have been disabled to avoid loops by virtue of the protocols like STP / MSTP etc. However in an SDN network, we can create multiple separate slices of the network which do not impact each other and thus allow to perform experimentation as well as optimizations as discussed below. This also gives a key benefit of having a distributed monitoring network and at the same time, not resulting into multiple SPAN ports and hence can leverage the data center switches for doing maximum of production traffic whereas the mirrored traffic is switched in a separate slice.

Also, the SDN controller has the complete view of the topology and hence can install the switching paths on the switches on a per flow basis, unlike in traditional networks, where the switches would not be aware about the complete network and hence the paths are learned and converged rather than programmed.

Deep Packet Inspection (DPI) is a technology by which a deeper examination of the packet, up to the Layer 7 can be performed. DPI has been used for variety of functions like lawful intercept, targeted advertising etc.

Now, in the context of the same, we would discuss some of the use cases where SDN / DPI and analytics together help provide richer monitoring results and optimizations.

Use Case 1 – Movement of VM, the traffic from which is being monitored.

Consider that there is a VM, the traffic originating from which and going towards Switch F is to be monitored. The controller knows that path taken would be SW A--- SW B—SWE---- SWF. It also knows that there is a monitoring device MD3 attached to switch B. Hence it would configure the port mirroring on Switch B.

Now assume that the VM migrates from Rack 1 to Rack 2 and hence the traffic emerging from VM would be first switched on SW D which is the Top of the rack switch for Rack 2. In a traditional network scenario, it would have been complex to provision the tap links and making configuration changes, apart from being time consuming. However, in the SDN scenario, the Monitoring App would receive the information about the change in the movement of the VM from the Controller and hence can instruct the controller to disable the mirroring on Switch B and enable the mirroring on Switch D. Apart from that, the controller can also either enable monitoring on MD2 or The controller can also program the Mirror network switch to route the specific traffic of the VM to MD3 , while MD2 continue to process other traffic which it was previously doing.

Let us examine how the above behavior is achieved with SDN enabled switches.

Let us assume that the VM, VM A has a certain IP address VMAIP. Usually, when the VM migration happens, the IP address is retained and IP tunnels are created to ensure the continuity of the sessions.

The moment the migration happens from the Switch A to Switch D, the next immediate packet that originates from the VM would now ingress into Switch D. Now for Switch D, it is a new flow and hence it would send the incoming packet to the SDN controller.

Now the controller would pass the packet information to the Monitoring App to see if the flow has to be monitored. Since as per the monitoring requirements, the flow has to be monitored, the monitoring App would look into the packet details and find that the source IP of the flow matches with the rule that has been previously set up for monitoring. In parallel, the SDN controller would also be setting up the data path based on the packet properties.

The mirroring App would then perform the following actions:

1. The mirroring App would query the SDN controller for the path that is being set up for the packet.
2. Based on the path, the monitoring App would determine the optimal switch at which the packets should be mirrored in the production network. (The above simple topology is not indicative of the humongous complexity of the Datacenter network where we would typically have 100s or 1000s of links for production network and only 10s of links for monitoring network).
3. The monitoring app would also determine which port should be configured as the SPAN port on Switch D, considering the monitoring network topology and which services to enable at the new monitoring device – which would be the MD2
4. The monitoring App would provide these information to the SDN controller, which in turn would make the necessary configuration changes on the switch and on the monitoring device.

5. The monitoring App can then disable mirroring on Switch B and monitoring services on MD3, if there are no other flows that are to be mirrored on Switch B and no service required to be performed on MD3.

Some of the advanced Monitoring Apps can proactively perform the above steps by registering for VM migration events while other can take a reactive approach based on the information sent by the Controller.

Now let us consider some other use cases in brief.

Use Case 2 – Making the distributed monitoring system failsafe.

Consider in the above topology that MD1 and MD2 are independently doing their monitoring functions. The traffic coming from Switch B is mirrored and sent to MD 1 on Switch D and the traffic coming from Switch C on Switch D is mirrored towards MD2. Now in a situation that MD1 fails, in a non SDN scenario, this would have resulted in the monitoring information getting lost till the time the admin notices the malfunction and makes the necessary configuration change. However, since the SDN controller would know about the state of the devices, it can quickly make the configuration change wherein, the mirror port towards MD2 becomes an aggregator port, sending both the traffic from Switch B and Switch C to MD2 and thus preventing any loss of monitoring information.

Use Case 3 – Conservation

An extension of Scenario 2 could be the Scenario 3, where MD1 and MD2 are performing similar services for different traffic sources. Such a configuration would have been done in situations where the traffic is huge and a single monitoring device cannot handle the same. Now if the monitoring app determines leveraging the flow statistics that the actual traffic flowing is way below the capacity (say during non-peak hours / weekends etc.), it can dynamically make changes such that MD2 can shut down and MD1 performs the entire processing. When the load increases, the app can bring MD2 to service. Such an approach can result in energy conservation and thus help in improving bottom line as well as making the data center greener.

Use Case 4 – On demand mirroring

Consider that mirroring is used for recording of VOIP conversation. It is known that whenever a new stream of flow starts, the controller would receive the notification from the 1st open flow enabled switch and controller would be programming the path on the switch. With emergence of DPI and DPI offload systems, it is possible to create a DPI App, which can interact with the controller. The DPI App can examine the packet and inform the controller whether the packet is indicating start of a conversation. If so, only at that moment, mirroring can be enabled and the same is continued till the controller is notified from the DPI App that it has seen the packet information which indicates the conversation is concluded. This feature can help reducing the load on the monitoring devices as they are now processing the packets only when they have to, based on very granular criteria and thus help optimize capacity planning and the investments for the same.

Use Case 5 – Audit

Some of the mission critical monitoring requirements demand the need for a constant health check and proactive actions to ascertain that the monitoring devices perform the task without interruption. The monitoring app can trigger such audits leveraging the controller and tests the monitoring devices in terms of function and performance.

Use Case 6 – Monitoring Access Control

By automating the configuration of mirroring / monitoring through the monitoring app and DPI, there is a possibility of having additional feature of allowing the configuration changes based on the role of the person performing the same and thus restrict the privileges to monitor only certain applications / VMs rather than entire system. This is a huge security compliance benefit which allows the Data center admin to provide restricted privileges/flexibility to the admins of their customers and at the same time protecting their customers from data theft.

Business Benefits of SDN aware / SDN enabled Mirroring Solutions

For Enterprises / IT Admins / CIOs

- Capex savings by using bare metal switches / Virtual Switches in case of lower scale. For e.g. one can set up a monitoring system at no additional cost by having Open VSwitch VM for switching, Wireshark VM for monitoring and Open Source SDN controller for orchestration.
- Capex savings by achieving distribution of monitoring functions without replication in each segment due to advanced network slicing capabilities of SDN networks.
- Opex savings by dynamic programming of the monitoring device.

For Existing Monitoring Solutions Vendors

- Ability to offer variety of services with scale as you go / pay as you go models.
- Ability to break into price sensitive customers
 - Smaller customers with virtualized solutions.
 - Larger customers with Service chaining.

For Start Ups

- Ability to focus on differentiated application, which can help market penetration, by leveraging low cost open source solutions and reduced Time to Market.

For ISPs

- Agility and flexibility of offering monitoring services.
- Newer business models like monitoring infrastructure sharing, which is enabled by combination of DPI to differentiate the traffic from different enterprises and SDN for rapid configuration updates.

Conclusion and call for action

As noted above, there are multiple use cases where an SDN orchestrated monitoring system can help in various ways. The above list of use case is a very small indicative list and there are numerous functions which can be made possible using the power of SDN, DPI and analytics.

As the CEO of cPackets, one of the prominent monitoring solution vendor has rightly noted, Network Monitoring is dead unless the solutions are agile to keep up with current cloud and virtualized data center needs.

The market leaders in monitoring systems are adopting SDN and it is imperative to develop SDN monitoring applications which help make the monitoring systems agile and provide scale and flexibility.

The monitoring application should have a pluggable architecture which allows it to interact and thus leverage the power of DPI, Analytics and other applications for an efficient operation. It should naturally be vendor independent, giving the customer, the choice and flexibility to choose the monitoring device.

References

- Material on Vendor websites like VSSmonitoring.com / cpackets.com
- <http://www.youtube.com/watch?v=cYe2aWebHZo&feature=youtu.be>
- <http://www.slideshare.net/SarmadMakhdoom/challenges-in-cloud-computing-vm-migration>
- <https://www.opennetworking.org/>

About Happiest Minds

Happiest Minds is a next-generation IT services company helping clients differentiate and win with a unique blend of innovative solutions and services based on the core technology pillars of **cloud computing, social computing, mobility and analytics**. We combine an unparalleled experience, comprehensive capabilities in the following industries: **Retail, Media, CPG, Manufacturing, Banking and Financial services, Travel and Hospitality and Hi-Tech** with pragmatic, forward-thinking advisory capabilities for the world's top businesses, governments and organizations. Founded in 2011, Happiest Minds is privately held with headquarters in Bangalore, India and offices in the USA and UK.

For further information, please contact business@happiestminds.com

Corporate Office

Happiest Minds Technologies Pvt. Ltd.
Block II, Velankani Tech Park
43 Electronics City
Hosur Road, Bangalore 560100, INDIA
Phone: +91 80 332 03333
Fax: +91 80 332 03000

United States

116 Village Boulevard, Suite 200
Princeton, New Jersey, 08540
Phone: +1 609 951 2296
2018 156th Avenue NE #224
Bellevue, WA 98007

United Kingdom

200 Brook Drive, Green Park, Reading
Berkshire, RG2 6UB
Phone: +44 11892 56072
Fax: + 44 11892 56073

About the author

Dharmraj B Jhatakia

General Manager, Product Engineering Services, Happiest Minds Technologies Pvt. Ltd.

Dharmraj is part of Happiest Minds Product Engineering Services team and leads the Telecom & Networking practice and manages engagements with the clients in Telecom & Networking space. SDN is a key focus area for Happiest Minds and passion for Dharmraj. Dharmraj and his team have built show case solutions which leverage the power of SDN including a field engineer monitoring application and a dynamic flow based bandwidth on demand application. They are currently focused on developing a multilayer networking platform, which uses Open Stack as the base cloud management platform. The platform would have reference applications like Network Visualization, Bandwidth on Demand and Service Orchestration.