Enabling Pervasive Network Observability with DANZ Monitoring Fabric (DMF)

Document Version 1.0

Pervasive **Network Observability** empowers a network operator to determine connectivity and performance-related issues across endpoints (clients, devices, applications, etc) on the network, with contextual awareness and predictive analytics regarding any packet, any flow, any endpoint at any time.

Table of contents

Introduction

This guide serves Network and Security Architects with a solution for not only pervasively monitoring the network but bringing in context-aware visibility to the networks they manage and operate. This Network Observability helps the Network and Security architects not just to swiftly pinpoint the problem but to take immediate prescriptive actions and reduce the meantime to repair (MTTR).

To efficiently operate the network and provide end-users with the optimum resources to be productive, network administrators need more than just a packet broker that filters and delivers traffic to the centralized tools. They need to have visibility about flows in the network, application services being used, and network resources usage patterns. Additionally, they need to deepen this visibility by mapping the flows to the actual end-users, servers, custom services within the organization, and production switches these flows traverse through. Network and security architects then need to carry this context to probe raw packet storage and perform a deep application performance analysis that enables them to remediate the problem and restore business productivity.

Today, network administrators are forced to pivot from a packet monitoring tool to a network flow collector to a device metric monitoring tool, and back again - accessing separate dashboards with different parameters, widgets, and reports. There is no context that can be shared between these various tools. Scouring through performance reports and checking into dozens of dashboards to identify network issues or security breaches is inefficient, and the siloes of network data created by this approach limit visibility into overall network performance and connected devices.

DANZ Monitoring Fabric (DMF) solution provides all these components of network observability with a single pane of glass management. The DMF solution combines the Network Packet Broker functionality with Packet storage and Analytics to enable pervasive Network Observability across data centers and campus networks.

This solution guide discusses:

DMF network observability components

- Hardware and Software options
- Connection Diagram
- Installation and Bringup
- Configuration
- Use cases
  - Slowness in database response reported by the end-user
  - User not able to connect to a Service
  - Find SSL sessions using expired certificates

DMF Network Observability Components

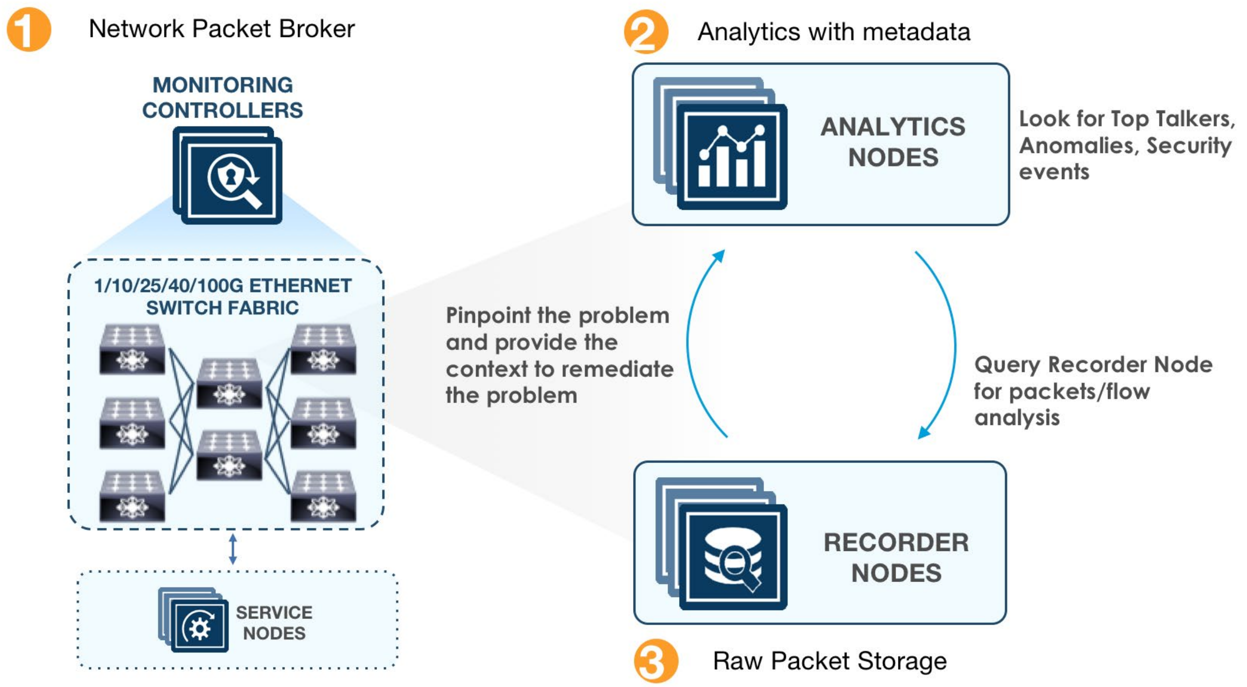

Hardware and Software Options

Following are the hardware and software components used in this solution

Component	Hardware/Software
DMF Controller	Hardware Appliance or Virtual Software
DMF Switch	Ethernet switch (1G/10G/25G/40G/100G) with DMF software license
Service Node	Hardware Appliance
Analytics Node	Hardware Appliance
Recorder Node	Hardware Appliance

Connectivity Diagram

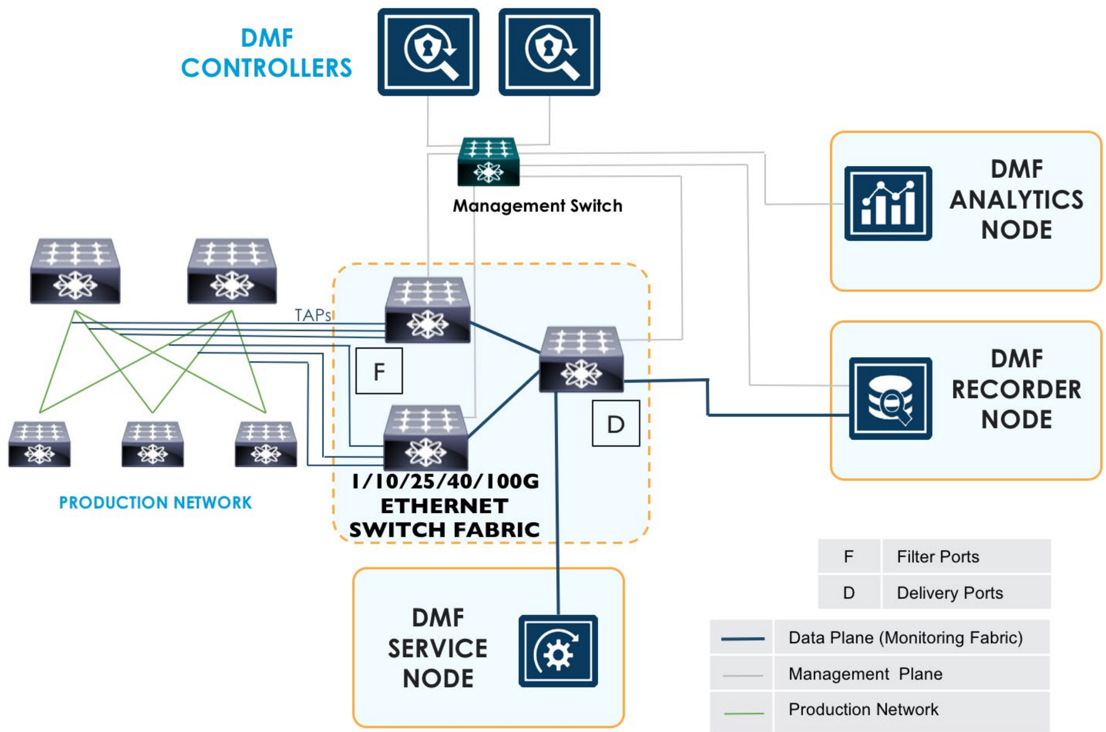

Installation and Bring up

1. Install the DMF controller software.
2. Access the Controller GUI to add the switches one by one to the fabric. Zero touch networking displays all the switches by listing them with their MAC addresses. The network admin adds the authorized switches to the monitoring fabric. The controller automatically provisions the switches with appropriate software and configuration.
3. The Service Node is automatically provisioned by the DMF controller.
4. Similarly, authorize the Recorder Node with its MAC address. The Recorder node is also provisioned with the correct software and configuration.
5. The correct analytics ode software is then installed as the last step in the process.

Please refer to the Deployment Guide for detailed steps for the installation process.

Configuration

Configuration for IPFIX Service

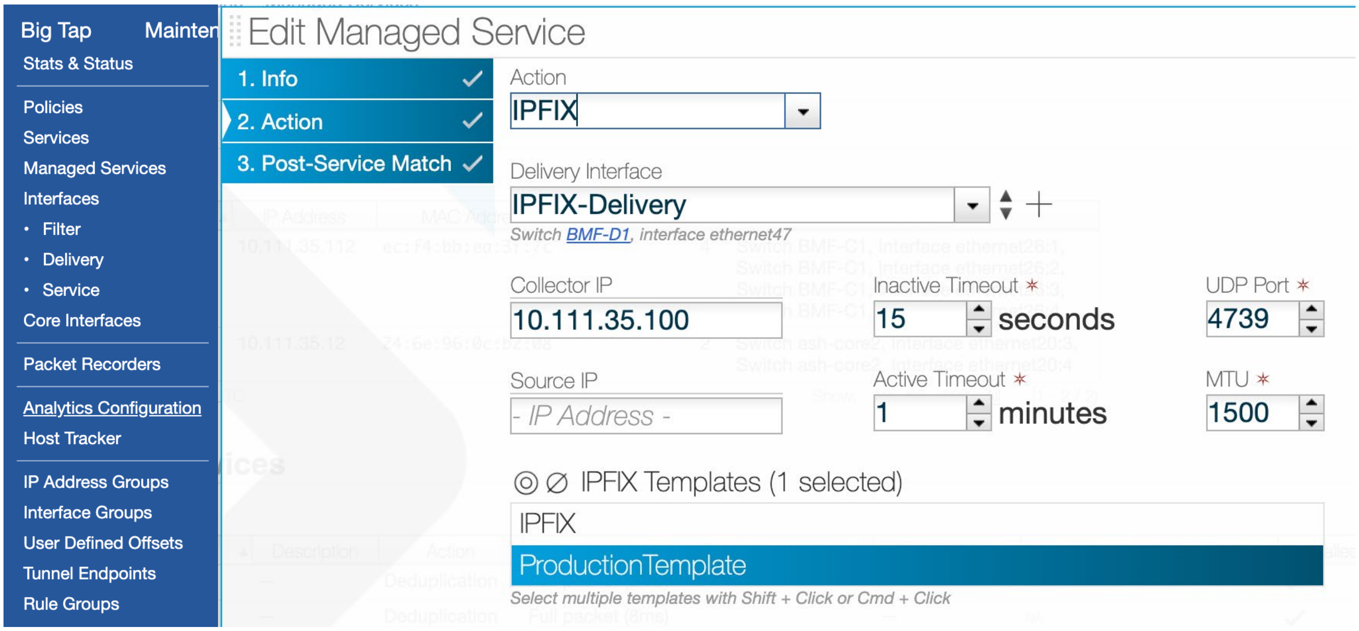

Configuration for Recorder Node

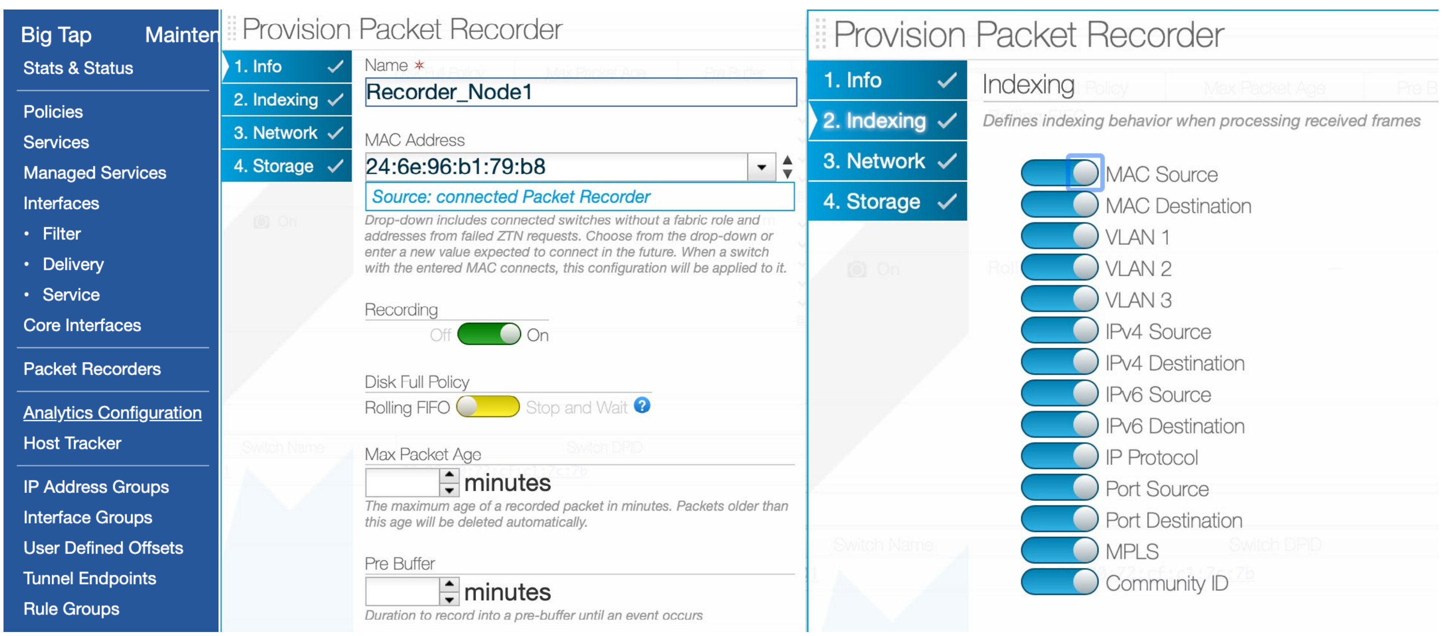

Configuration for Analytics Node

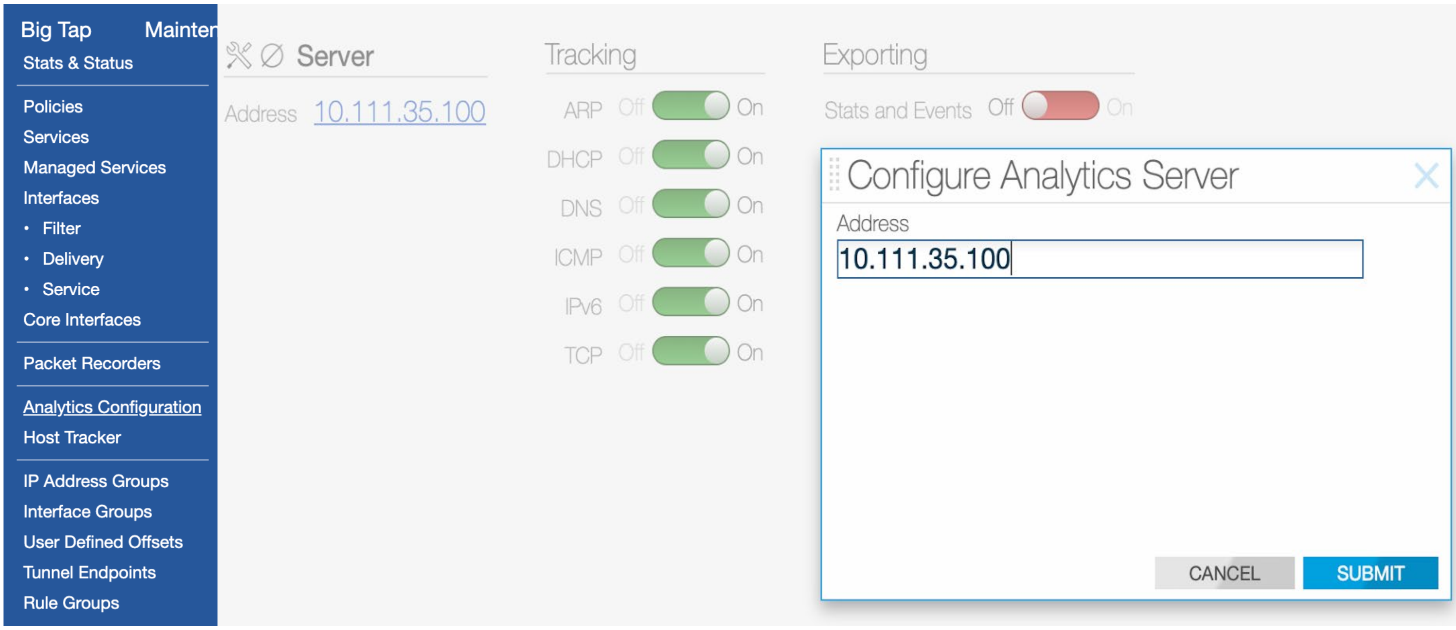

Configuration for enabling sFlow generation

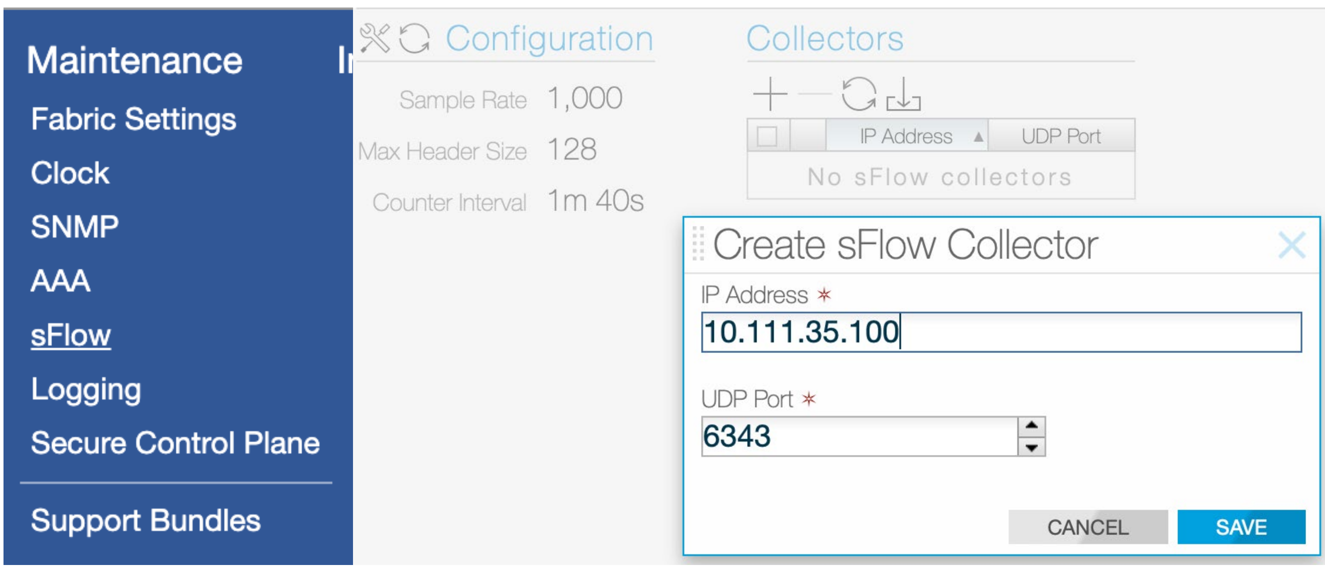

Policy Configuration for Packet Storage

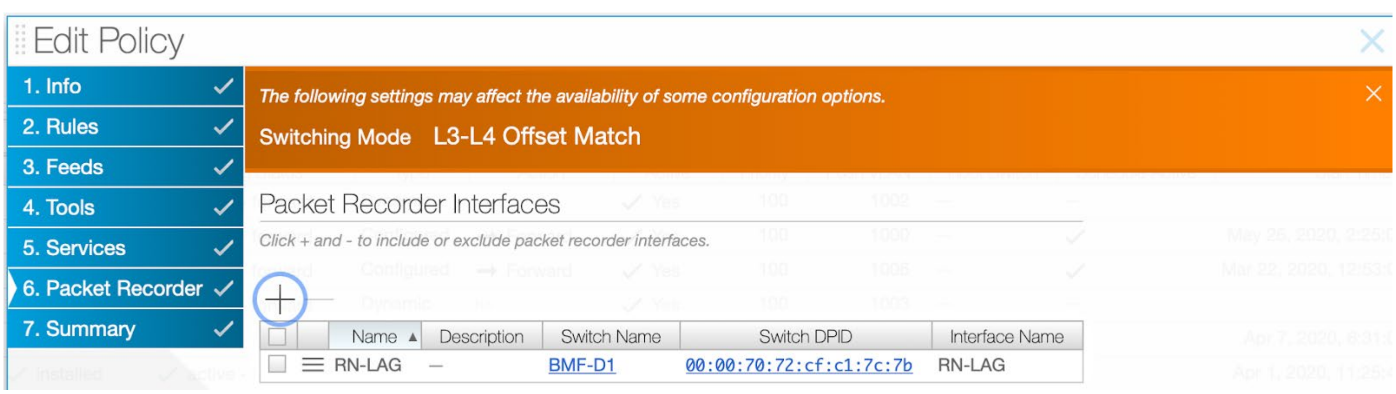

Use Case

1. Slowness in database response reported by the end-user

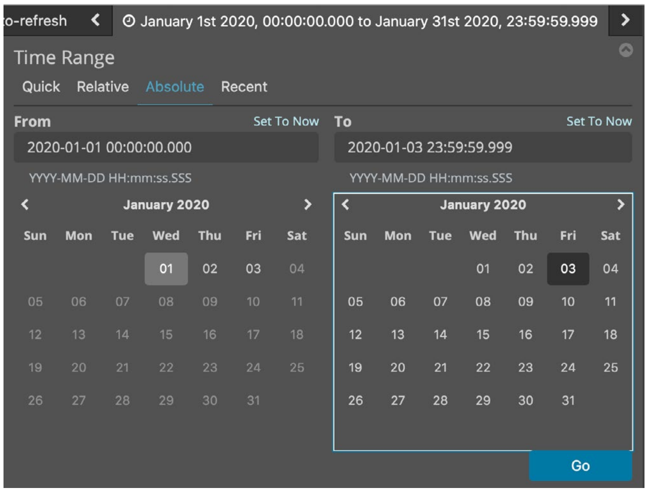

A user complains about a slowness he/she is experiencing consistently for over three days. With pervasive visibility in place, the network administrator is assured that the sessions initiated by this particular user are captured. The admin can access any flow-based dashboard on the Analytics Node (predefined dashboards like sFlow, Netflow, or custom dashboard) and select the time interval (Figure 1) when the user faces slowness.

Figure 1: Select time interval

Now filter the flows by selecting the end-user, either by source IP address (Figure 2) or by usernames. The Analytics Node enriches the flow data collected by integrating with OpenVPN or Active Directory. This integration learns the mapping of usernames to their assigned IP addresses as they log in to the network. This mapping is then used by the Analytics Node to enrich the existing flow data.

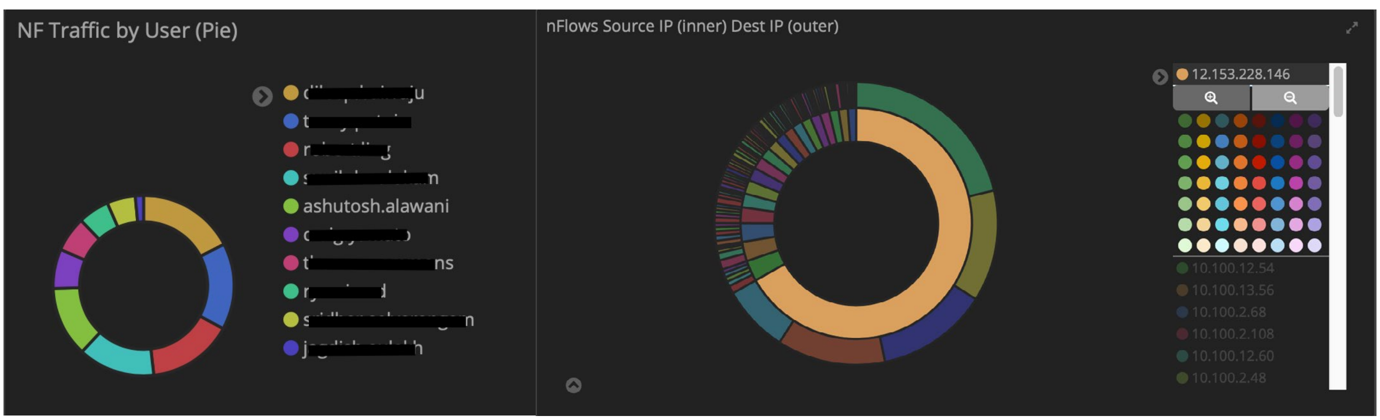

Figure 2: Select the end-user

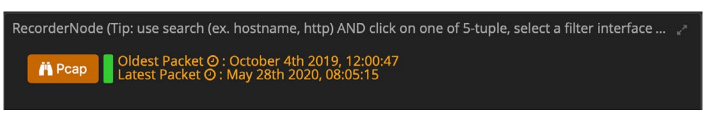

Figure 3: Recorder Node query for flow analysis

With the filters in place for the user described problem, we now query the Recorder Node (Figure 3) for further flow analysis. The recorder node is closely integrated with the Analytics Node and provides a single seamless workflow.

The context created with the help of previous steps is now carried over to the Recorder Node. This context (Figure 4) can include the time interval and source or destination IP addresses or port numbers. Among the various query types, the user has the option of generating application identification, download the raw packets in pcap format, replay the packets to a tool for a detailed security analysis or perform flow analysis for protocols like HTTP, DNS, RTP, or TCP and all of this can be achieved within the context of the use case.

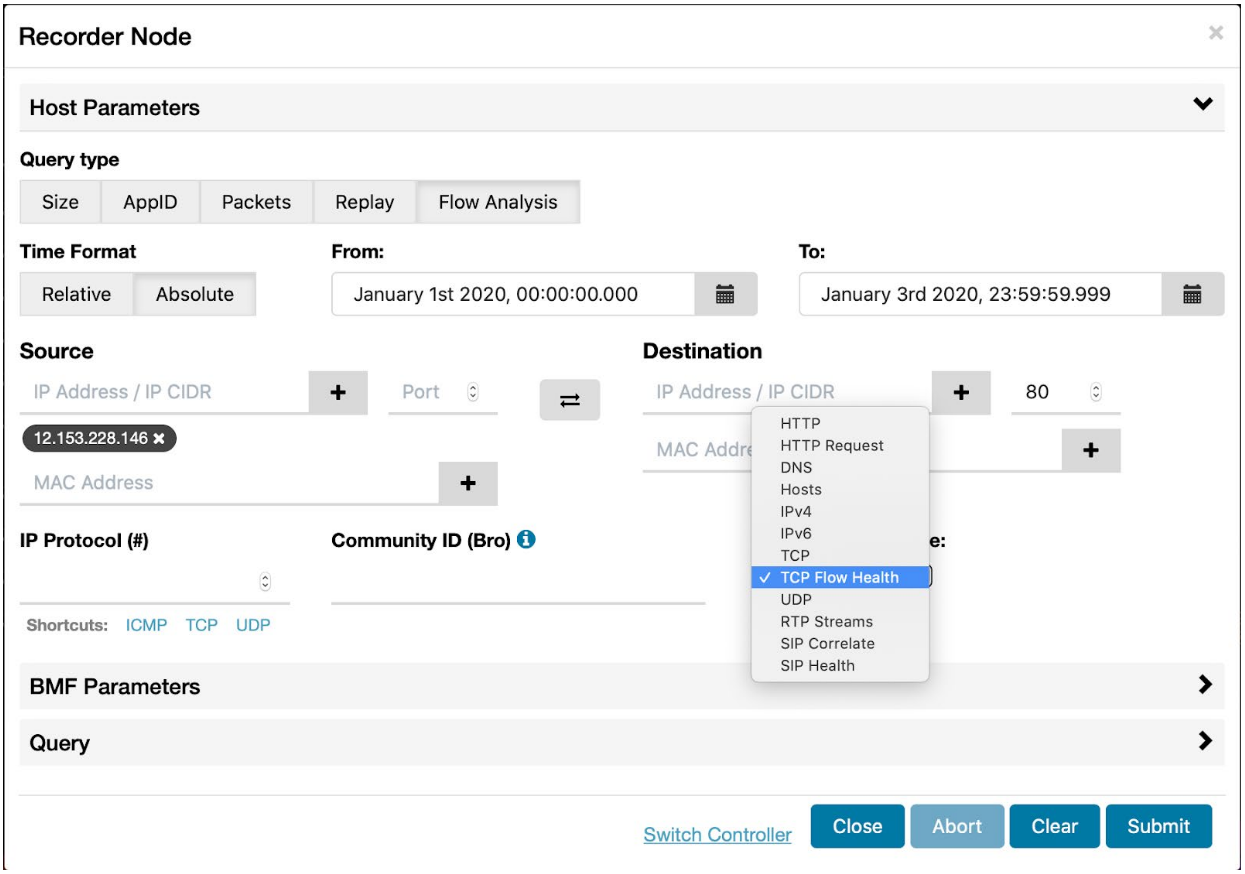

Figure 4: Context-aware Flow Analysis of raw packets

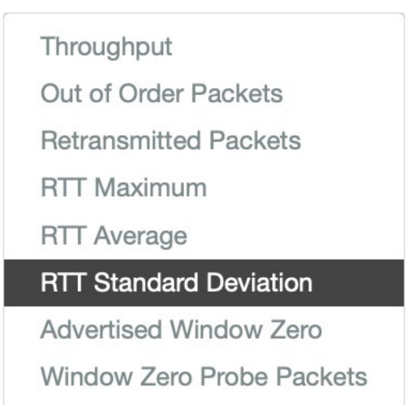

Figure 5: TCP Flow Analysis metrics

For the slowness that the user complained, one can perform a TCP flow analysis This analysis helps derive metrics associated with the TCP sessions and include Round Trip Time (RTT) encountered by the sessions, retransmissions or out of order packets, and zero window size advertisements (Figure 5). The network administrator can use statistics of any of these metrics or can derive custom metrics based on these basic TCP metrics

On performing the analysis based on standard deviation in round trip times, we do see that the stdev in RTT for one particular flow (Figure 6) is very high compared to the other direction of the flow and compared to other flows as well.

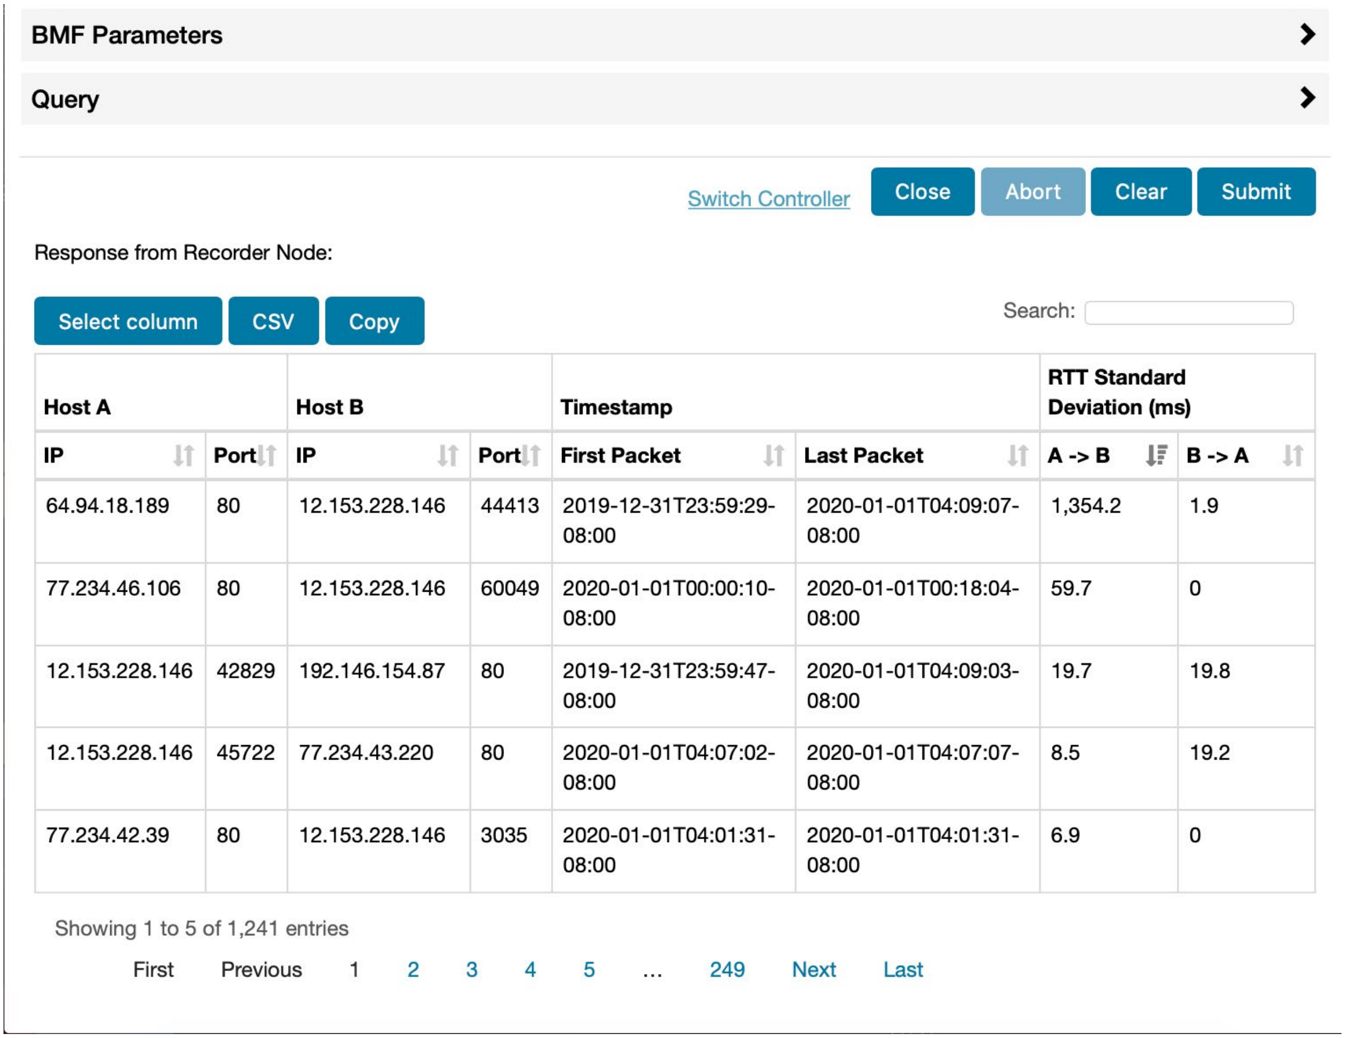

Host A		Host B		Timestamp		RTT Standard Deviation (ms)	
IP	Port	IP	Port	First Packet	Last Packet	A -> B	B -> A
64.94.18.189	80	12.153.228.146	44413	2019-12-31T23:59:29-08:00	2020-01-01T04:09:07-08:00	1,354.2	1.9
77.234.46.106	80	12.153.228.146	60049	2020-01-01T00:00:10-08:00	2020-01-01T00:18:04-08:00	59.7	0
12.153.228.146	42829	192.146.154.87	80	2019-12-31T23:59:47-08:00	2020-01-01T04:09:03-08:00	19.7	19.8
12.153.228.146	45722	77.234.43.220	80	2020-01-01T04:07:02-08:00	2020-01-01T04:07:07-08:00	8.5	19.2
77.234.42.39	80	12.153.228.146	3035	2020-01-01T04:01:31-08:00	2020-01-01T04:01:31-08:00	6.9	0

Figure 6: Standard deviation in RTT analysis

This helps the network administrator that there have been some routing changes that occurred while this session was active. This high value would be seen for all the sessions that followed the same routing path as the first session did. If the network admin notices that the stdev in RTT is high only for a particular destination, then the admin can conclude that the network is behaving normally whereas it might the application that is not behaving normally. This helps the network administrator isolate the problem to the network or the application service.

2. **User not able to connect to a Service**

   A user cannot simply connect to the service. When this user complains about the issue, it would greatly help the network administrator if this issue could be quickly identified as a network issue or application issue. With the close integration of the Recorder Node with Analytics Node, the administrator can filter the flows in Analytics Node for the particular service that the user is facing issues with. The administrator then carries this context to the Recorder Node to perform a network health analysis for this flow. Amongst the various metrics of a TCP flow, the administrator can find if the server that the user is trying to connect to is terminating the connection immediately by sending a "reset". The administrator has the flexibility to define an expression that exactly achieves this.

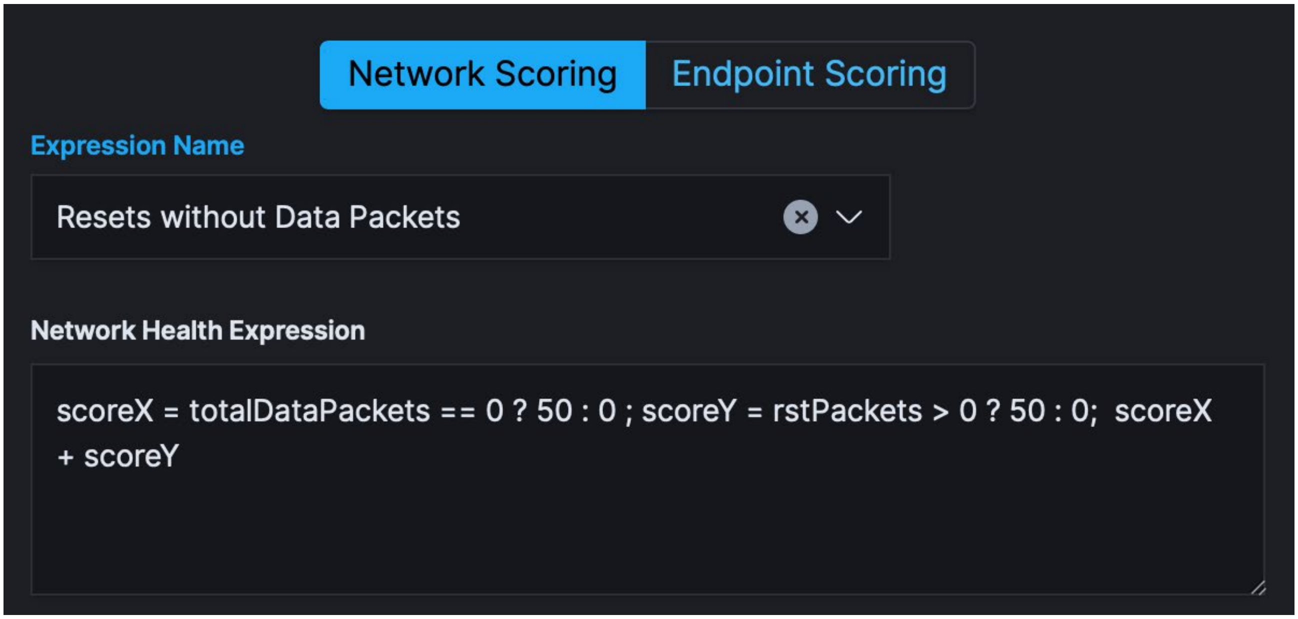

$$scoreX := \begin{cases} 50, & \text{if } totalDataPackets = 0 \\ 0, & \text{otherwise} \end{cases} ; \quad scoreY := \begin{cases} 50, & \text{if } rstPackets > 0 \\ 0, & \text{otherwise} \end{cases} ; \quad scoreX + scoreY$$

Figure 7: Customized expression using TCP Flow Analysis Metrics

As an example, the expression shown in Figure 7, finds sessions where the server has immediately disconnected the connection initiated by the client. A score equaling 100 for the session being analyzed, will help the administrator confirm that the server is disconnecting the connection before it is being set up. This reduces the time taken to identify if the issue is network-related or application-related.

In a similar manner, one can use the expression to analyze the RTT observed for the TCP connections.

$$score1 := \begin{cases} 50, & \text{if } rttAvgMs > 15 \\ 0, & \text{otherwise} \end{cases} ; \quad score2 := \begin{cases} 50, & \text{if } rttStdev > rttAvgMs \\ 0, & \text{otherwise} \end{cases} ; \quad 100 - (score1 + score2)$$

Figure 8: RTT Analysis for TCP Flows

Based on the average RTT and standard deviation in RTT, once can measure the health of the TCP flows and can identify issues like packet loss, unintended routing changes, and congestion.

3. Find SSL sessions using expired SSL certificates

One of the security use cases is to manage the SSL certificates used by servers within an organization or identifying the conversation initiated with servers that use either expiring or expired certificates. Analytics Node integrates with Zeek, a network security monitoring platform to enable security through observability. Analytics Node consumes the logs from the Zeek platform and correlates the analysis done by Zeek with the flow data it collects.

Figure 9: Advanced Query Bar

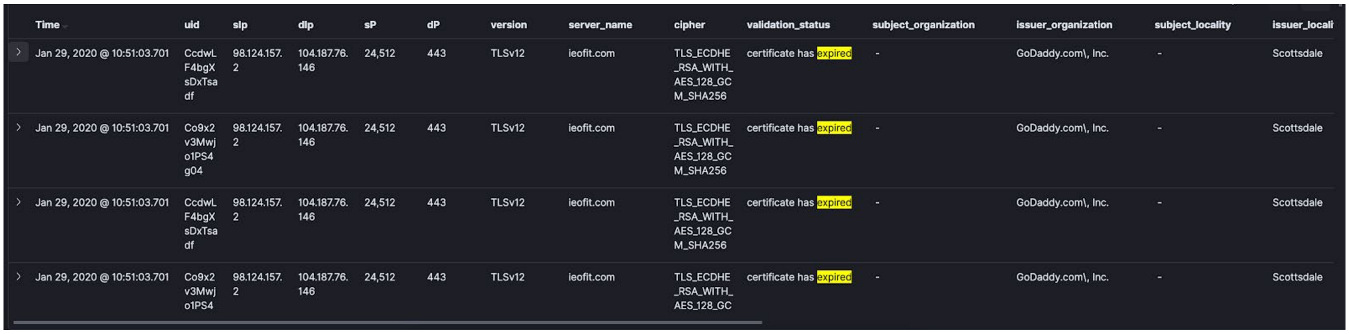

Time	uid	sIp	dIp	sP	dP	version	server_name	cipher	validation_status	subject_organization	issuer_organization	subject_locality	issuer_locali
Jan 29, 2020 @ 10:51:03.701	CcdwLF4bgXsDxTsadf	98.124.157.2	104.187.76.146	24,512	443	TLSv12	ieofit.com	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	certificate has expired	-	GoDaddy.com\, Inc.	-	Scottsdale
Jan 29, 2020 @ 10:51:03.701	Co9x2v3Mwjo1PS4g04	98.124.157.2	104.187.76.146	24,512	443	TLSv12	ieofit.com	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	certificate has expired	-	GoDaddy.com\, Inc.	-	Scottsdale
Jan 29, 2020 @ 10:51:03.701	CcdwLF4bgXsDxTsadf	98.124.157.2	104.187.76.146	24,512	443	TLSv12	ieofit.com	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	certificate has expired	-	GoDaddy.com\, Inc.	-	Scottsdale
Jan 29, 2020 @ 10:51:03.701	Co9x2v3Mwjo1PS4	98.124.157.2	104.187.76.146	24,512	443	TLSv12	ieofit.com	TLS_ECDHE_RSA_WITH_AES_128_GC	certificate has expired	-	GoDaddy.com\, Inc.	-	Scottsdale

Figure 10: SSL sessions with expired SSL Certificates

One can use the query bar (Figure 8) to search for SSL sessions that are using expired SSL certificates. This search gives the user a list of SSL sessions with expired SSL certificates presented by the server and also the details of the sessions. These details include source and destination IP addresses, which can be correlated to the flow information and help the Network administrator locate the specific servers and update the SSL certificates.

Technical Resources

DANZ Monitoring Fabric (DMF) - https://www.arista.com/en/products/danz-monitoring-fabric

Hands-on Labs - https://dmf-labs.arista.com

Summary

Production networks generate invaluable data for understanding application performance problems, security issues, day-to-day troubleshooting, and reducing Mean Time to Resolution (MTTR). But as the volume of data flowing through the network continues to surge, capturing and analyzing that data is becoming a daunting task. Scouring through performance reports and checking into dozens of dashboards to identify network issues or security breaches is inefficient, and the siloes of network data created by this approach limit visibility into overall network performance and connected devices.

DANZ Monitoring Fabric out-of-band solution scales with the production network with a single pane of glass management. The solution also integrates Network Packet Broker (NPB) functionality with packet recording (Recorder Node) functionality and Network Analytics (Analytics Node). This integration helps bring in context-aware visibility into the Network which in turn increases the network observability and helps Network architects administrate the network more productively.

References

[1] Observations on Round-Trip Times of TCP Connections, Phillipa Sessini and Anirban Mahanti University of Calgary 2500 University Drive NW Calgary, AB, Canada {sessinip, mahanti}@cpsc.ucalgary.ca https://people.cs.umass.edu/~phillipa/papers/SPECTS.pdf

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062