Augtera The First Network AI for the Modern Enterprise

Solution Brief

Proactive Detection and Isolation of NVMe/TCP Congestion using Machine Learning.

Introduction

Architects designing the next generation data centers are looking for lower-cost IP SAN-based solutions as an alternative to Fibre Channel (FC)-based Storage Area Network (SAN) infrastructures. NVMe/TCP is an extension of the NVMe base specification that defines the binding of the NVMe protocol to message-based fabrics using TCP. NVMe/TCP enables architects to build highly scalable storage environments that allow large-scale deployments and operations over distances.

Benefits of using NVMe/TCP are that it works with regular TCP/IP based networks, is highly scalable, supports disaggregated infrastructure, and reduces cost.

As organizations seek to take advantage of these benefits by migrating to NVMe/TCP it becomes critical to make sure the underlying IP infrastructure is performing as expected. Specifically, that the underlay has no congestion or errors. Augtera platform provides a simple workflow to enable this. This solution brief will detail how:

- Proactive detection and notification of congestion can be achieved using native sFlow telemetry from switches coupled with machine learning capabilities from Augtera.
- Augtera analytics and visualizations can be used to determine the TCP flows that are impacted by the congestion
- Operators can gain an end-to-end view including the servers and the Data Center fabric to isolate whether the congestion is caused by the fabric, the servers or the interconnect between the servers and the fabric.
- Operators can be notified of congestions events using collaboration tools such as Slack or enterprise ticketing tools such as ServiceNow

In this solution brief we will not be reviewing the benefits of selecting a disaggregated NVME/TCP storage architecture, more details on that can be found here.

TCP session establishment and recovery

TCP is a connection-orientated transport. Before data can be transmitted between a host and storage subsystem over a TCP fabric, a TCP connection must be made. Source IP address, source TCP port, destination IP address, and the destination TCP port define the TCP connection. By default, NVMe/TCP uses port 8009 for discovery controllers and port 4420 for I/O. When a TCP connection is fully established, the host and controller can begin communication.

The TCP protocol supports a method to request a retransmission of lost packets between a sender and a receiver. Specifically, if there is loss in the path, TCP selective acknowledgements (SACK) from the receiver will inform the sender which packets were lost. This native feature of TCP minimizes the number of packets that a sender must retransmit and improve overall throughput.

A byproduct of this capability is that it enables the detection of congestion. A well-functioning TCP/IP infrastructure between the sender and the receiver should have a minimal number of packets with SACK data present in them. What makes SACK especially useful is that the presence of packets with SACK can be detected using sFlow, a feature most switches support requiring no additional hardware. In the picture below the gap between the “left edge” and “right edge” represent the “missing” packets and what the sender should retransmit.

Image 1: Wireshark Capture of a packet with SACK

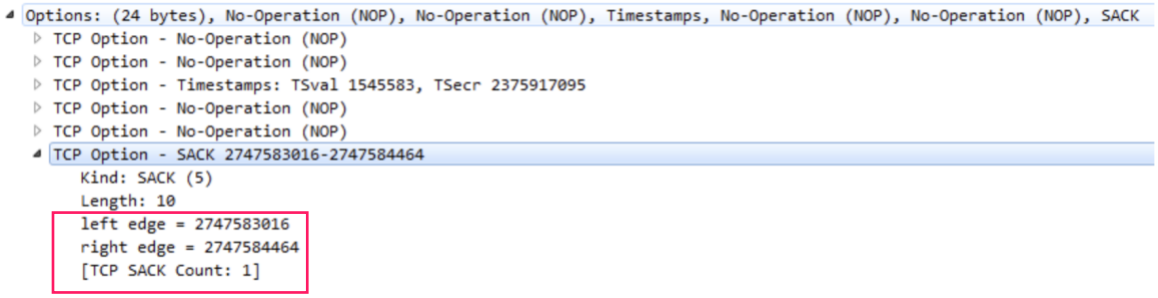

Enabling congestion detection using sFlow

Augtera is a purpose-built network AI platform, designed to enable proactive network operations. It collects network telemetry from every data source and applies Machine Learning (ML) algorithms to that data to identify actionable anomalies or “AI insights”. The operators do

not need to configure thresholds or specify normal behavior. Augtera automatically learns normal patterns and identifies misbehaviors with very high fidelity. For an overview of the Augtera platform please read the Data Sheet and visit the Augtera Platform page.

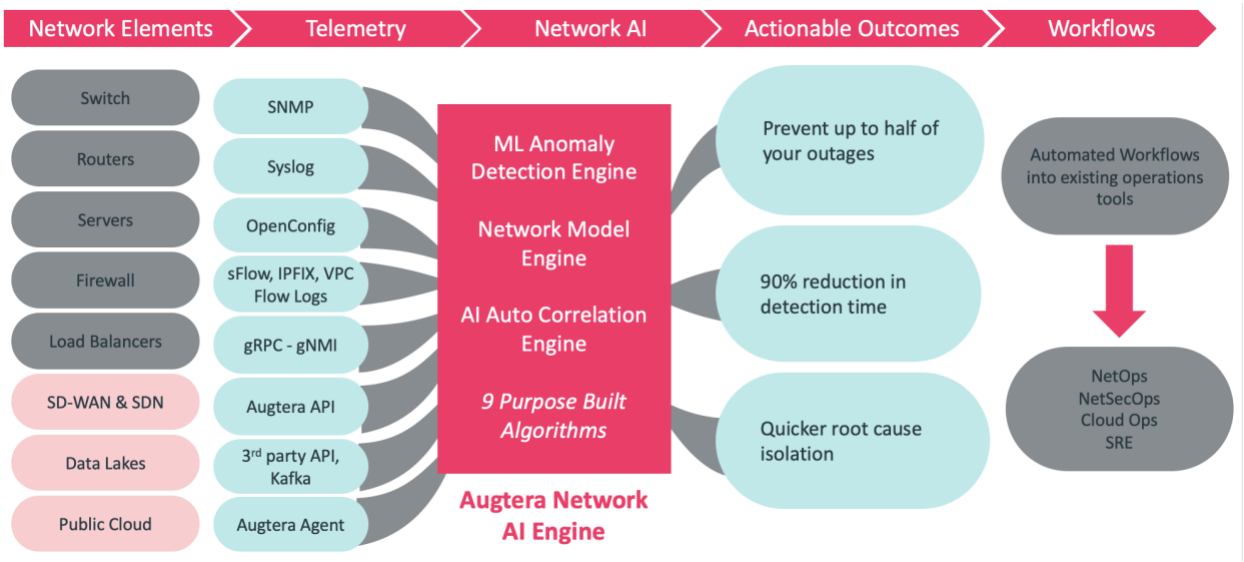

Augtera is a software only solution that can be deployed on a collection of VMs in your premises network or using Augtera SaaS. To detect congestion Augtera will be configured by the support team as part of the installation to collect sFlow data from the Data Center fabric switches. As the Augtera platform ingests the sFlow packets it will extract the components of the packets such as the source IP/Ports, Destination IPs/Ports, TCP flags and identify packets that have the SACK option in them.

Image 3: Augtera anomaly for packets containing SACK

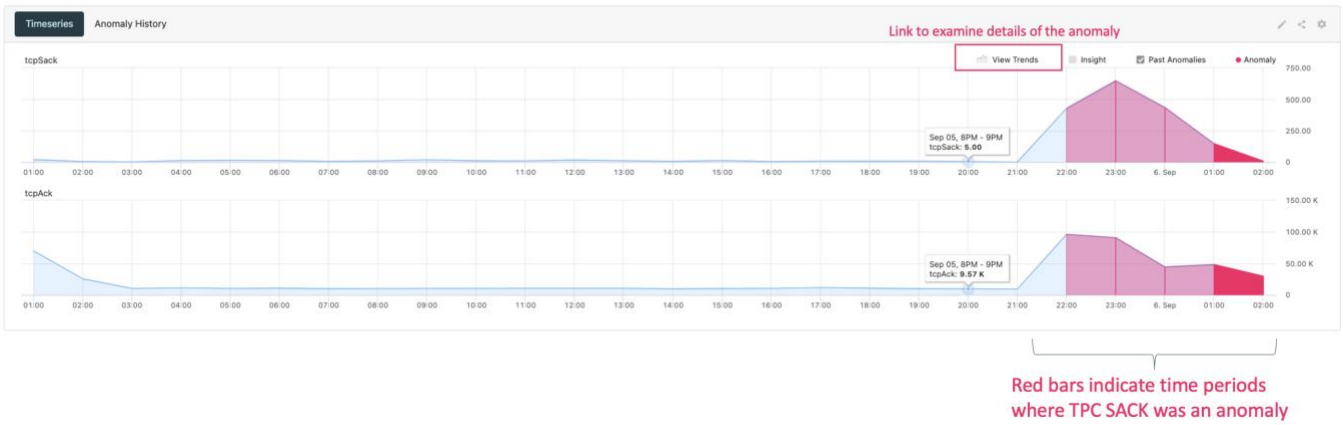

Going one step further, it learns the normal rate, if any, of packets with SACK in them per switch, port, and even IP tuple. Augtera ML will identify and automatically notify when there is an unusual number of packets with SACK in them, indicating congestion. Below is an example of a SACK anomaly as detected by Augtera. In the example below, the fabric was running very clean with almost no retransmissions. In this network, sFlow was configured to sample 40,000:1 packets, so the anomaly below really represents millions of lost packets.

Identifying the TCP flows impacted

The initial identification and notification of congestion is a key first step. However, more is needed. The operator will next want to identify which flows are having congestion and identify the root cause of the congestion.

Image 4: Flows with SACK options present in them that align to anomaly

As mentioned earlier, Augtera identifies the IP tuples inside the sFlow data and provides deep analytics and visualizations to the operator to leverage this data both in real-time and historically. This enables the operator to quickly see what applications are impacted. With

Augtera, the operator clicks the “View Trends” link from the anomaly to see the underlying flows in the image below, the operator can now see the source IPs (servers) and destination IPs (storage arrays) and the destination port of 4420 which is NVMe/TPC I/O.

Identifying the cause of the congestion

Identifying where in the network packet are being dropped is the logical next step. In addition, to collecting sFlow data from the switches, Augtera collects many other metrics such as In/Out Octets, In/Out Discards, Frame Errors, Optical levels and more, not just from the switches but from the Dell ESXI servers as well. This end-to-end observability enables the operator to quickly figure out whether it is the servers or the fabric that is misbehaving. Further, the operator can quickly identify the specific server(s) or the specific fabric switch(es) and the interfaces that are misbehaving. Augtera comes with out of the box dashboards, or the operator can build custom ones if desired. For example, an operator may want a dashboard for different availability zones, regions, etc., and these can be built rapidly with the help of the Augtera support team.

In addition to data plane metrics, Augtera can collect control plane metrics such as route counts, advertised prefixes, BGP sessions, and log data such as Syslog. It is often the case that a control plane misbehavior results in data plane degradations.

Image 5: Consolidated view of traffic & discard ratio of the servers. Notice discard ratio goes down when traffic increases which indicates servers are not dropping packets.

Image 6: Consolidated view of the traffic and discards of the switch fabric. Notice the spike in discards that aligns to the change in traffic rate.

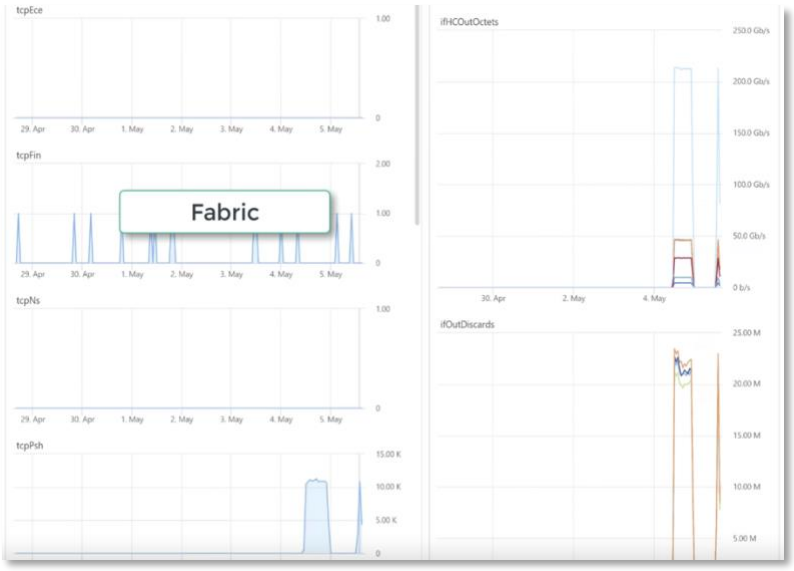

Correlation and Notification

As network degradations are detected it is important that they can be sent to the right teams and tools without creating duplication and noise. Augtera solves this by using auto-correlation of anomalies and events to a single “Incident”. Auto-correlation is powered by an auto-discovered network model. By understanding how servers and switches are connected, it can know when anomalies or events are related to each other and combine them into Incidents. These Incidents can be delivered to collaboration tools such as Slack or Teams or Enterprise workflow platforms like ServiceNow.

Image 6: Sample Augtera Proactive notification of congestion to Slack

Conclusion

While NVMe/TCP has many benefits, organizations should consider proactive operations of the underlay that it will utilize. It is critical to ensure the TCP/IP infrastructure is congestion free. Augtera machine learning along with the native switch telemetry, such as sFlow provide the means to do this, enabling pro-active congestion detection and using SNMP or streaming telemetry the ability to isolate the offending server(s), switch(es) and interface(s).

Related Links

- Dell NVME/TCP Overview
- Augtera Data Sheet
- Augtera Data Center Solution Brief