Plixer

WHITE PAPER

Avoid legacy NetFlow traffic analyzer woes

Gain detailed, accurate, end-to-end visibility into critical NetOps and SecOps data with Plixer Scrutinizer

Summary

Network and security operations teams use network traffic flow analysis technology like NetFlow and IP Flow Information Export (IPFIX) to monitor application performance, troubleshoot problems, and mitigate threats.

But when it comes to data accuracy and richness, not all NetFlow traffic analysis solutions are the same. Legacy solutions can misreport key statistics or conceal critical data when they deduplicate and stitch flow records, making it difficult for NetOps and SecOps professionals to pinpoint and resolve issues.

In fact, legacy NetFlow traffic analysis solutions suffer from three fundamental limitations related to how they deduplicate and stitch flow records:

1. Legacy solutions only consider the source IP address/port number, destination IP address/port number, and protocol ID when examining flow records. Lacking detailed visibility into traffic flows, they can inadvertently deduplicate dissimilar records, discarding or falsely characterizing critical troubleshooting and forensics data in the process.
2. Legacy solutions deduplicate and stitch flow records at collection time, squandering CPU cycles and storage capacity. The vast majority of flows are deduplicated and stitched, but are never actually analyzed by a network or security administrator.
3. Legacy solutions do not support deduplication and stitching across multiple collectors. Instead, each collector deduplicates and stitches flows independently. Lacking a holistic view of an end-to-end session, NetOps and SecOps professionals are forced to analyze issues and troubleshoot problems one collector at a time—a manually intensive, error-prone, and time-consuming proposition.

Plixer Scrutinizer is designed from the ground up to

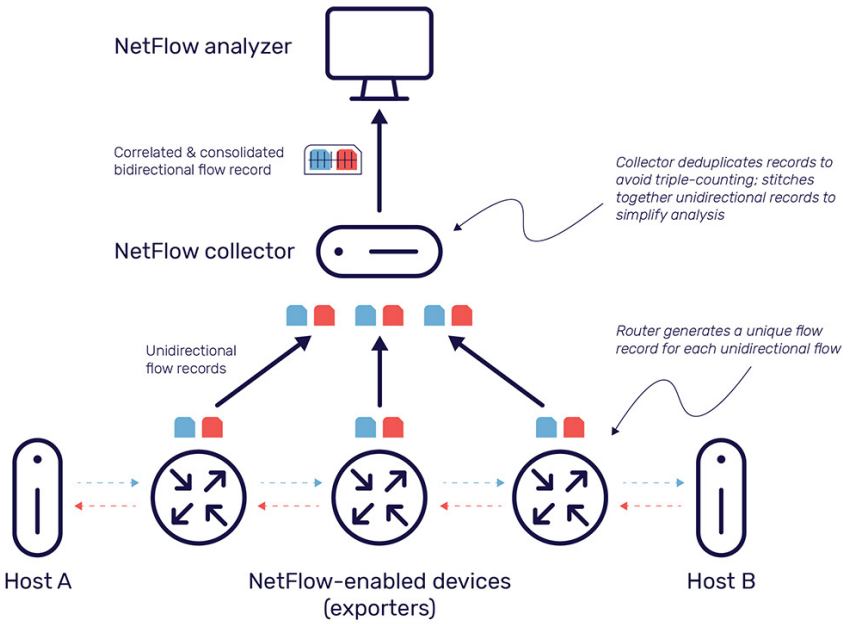

provide full and accurate visibility into critical network and security data. Unlike legacy NetFlow traffic analysis solutions, Scrutinizer:

1. Supports "dynamic key field" deduplication, helping network and security operations teams improve visibility, with detailed and reliable statistics.
2. Deduplicates and stitches flow records on-demand for ultimate efficiency and scalability.
3. Deduplicates and stitches flow records across collectors, giving NetOps and SecOps teams holistic, end-to-end visibility across the entire extended enterprise network, streamlining analysis and forensics.

This white paper provides a brief introduction to NetFlow and explains how Plixer's unique approach to deduplication and stitching eliminates the inherent architectural constraints of legacy NetFlow traffic analyzer solutions, helping IT organizations gain actionable insights into key network and security data so they can detect and resolve issues more quickly and effectively.

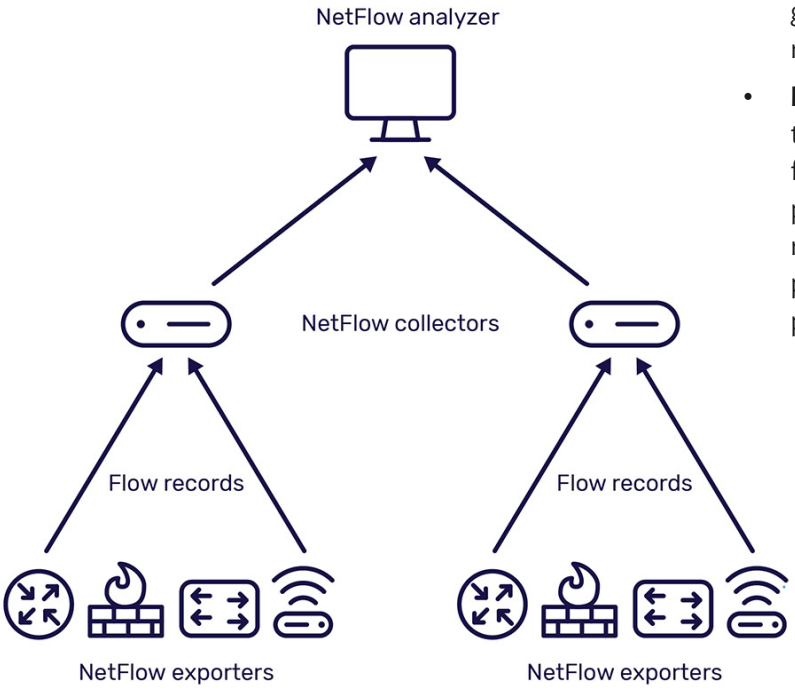

What is NetFlow?

NetFlow is a popular protocol used to capture and analyze traffic in IP networks. NetFlow data can be used for a variety of purposes, including:

- **Network operations**—performance monitoring, network optimization, and capacity planning.
- **Security operations**—threat and vulnerability assessment, detection, analytics, forensics, and remediation.
- **Compliance and risk management**—improving compliance with government and industry regulations, and corporate policies.

NetFlow is based on a hierarchical architecture, which includes:

- **NetFlow exporters**—network elements such as routers, switches, firewalls, probes and packet brokers that capture and export flow data. An exporter tracks key traffic statistics, generates flow records, and sends them upstream for processing and analysis.
- **NetFlow collectors**—specialized upstream network appliances or software applications that gather and pre-process flow records from one or more exporters.
- **NetFlow analyzers**—applications that transform raw data into meaningful information for NetOps and SecOps professionals. Most provide graphical user interfaces to help network administrators and engineers manage performance and capacity, troubleshoot problems and mitigate threats.

NetFlow exporters capture traffic statistics for each unique IP traffic flow. (NetFlow tracks traffic metadata only, not the actual data payload). All packets with the same source/destination IP address, source/destination ports, and L3 protocol are grouped into a NetFlow flow record.

Originally developed by Cisco and embedded with Cisco IOS, NetFlow is now a de facto industry standard, supported by a wide range of equipment manufacturers and software vendors.

Historically, two versions of NetFlow have been widely deployed in production networks: version 5 and version 9. NetFlow v5 supports IPv4 networks and unidirectional flow records only, and captures data in a fixed-format record that cannot be augmented or customized by vendors. NetFlow v9 adds support for stateful flow records, IPv6, MPLS and BGP, and introduces templates that allow vendors to extend flow records to support custom applications and services.

IPFIX is an IETF standard protocol based on NetFlow v9. NetFlow v9 and IPFIX significantly expanded the size, scope and volume of flow records, introducing new scalability and performance challenges for network traffic analyzer solutions.

Other flow technologies include sFlow (sampled flow), an alternative packet sampling technology that is supported by a number of network equipment and software vendors; J-Flow, an IP traffic flow sampling technology from Juniper Networks; and IxFlow, enriched flow data generated by Ixia network packet brokers.

NetFlow record deduplication and stitching

All NetFlow collectors perform two important functions to aggregate and correlate flow records from multiple exporters: deduplication and stitching.

Deduplication

NetFlow collectors deduplicate flow records to avoid duplicate counting when multiple exporters capture the same flow. The diagram below shows a client-server session traversing two routers in one direction and one router on the return path. Note that routers X and Y capture the same flow. As part of pre-processing, the collector discards one of the duplicate flow records to avoid double-counting statistics.[1]

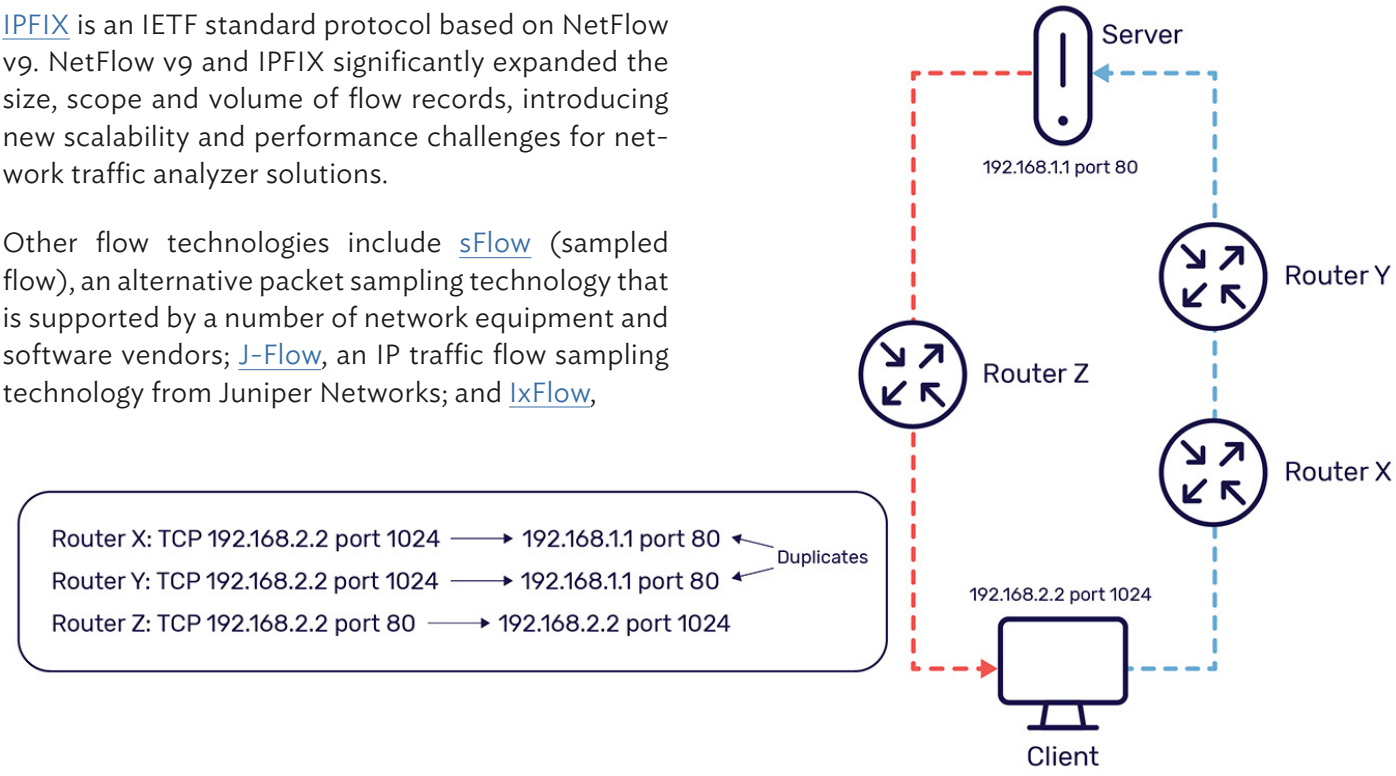

[1] Some legacy NetFlow traffic collector/analyzer vendors claim that deduplication helps conserve collector storage capacity. In fact, most legacy collectors store complete copies of both the original flow records and deduplicated flow records, which actually increases storage consumption.

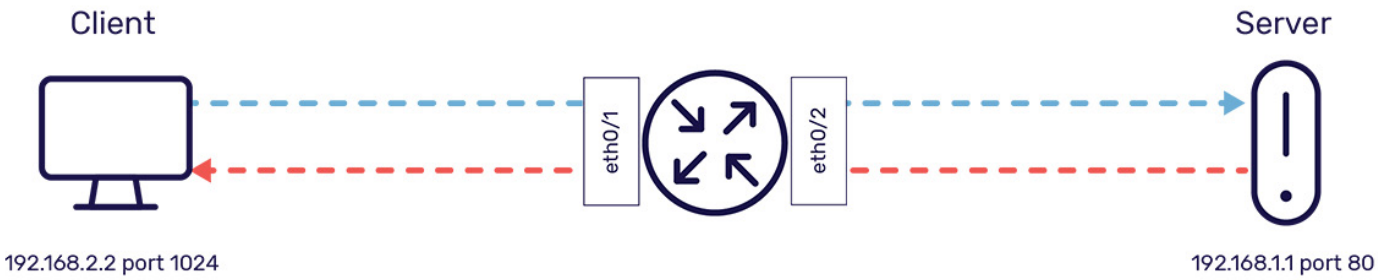

Unidirectional flow records

Start time	Interface	Source IP	Source Port	Destination IP	Destination port	Protocol	Packets sent	Bytes sent
09:00:00:001	eth0/1	192.168.2.2	1024	192.168.1.1	80	TCP	10	1005
09:00:00:301	eth0/2	192.168.1.1	80	192.168.2.2	1024	TCP	20	20005

Bidirectional flow records

Start time	Client IP	Client port	Server IP	Server port	Protocol	Client bytes	Client packets	Server bytes	Server packets
09:00:00:001	192.168.2.2	80	192.168.1.1	80	TCP	1005	10	20005	20

Stitching

NetFlow exporters generate two unidirectional records; one for each path of a session. In the diagram above, for example, the router generates one record for the client-to-server traffic, and another record for the server-to-client traffic. Network and security operations personnel often need to correlate the upstream and downstream flows to analyze issues and diagnose problems. To that end, a NetFlow collector "stitches" the upstream and downstream flows into a single bidirectional flow record.

Plixer eliminates legacy NetFlow deduplication and stitching constraints

Legacy network traffic analyzer solutions can misreport or discard data because of the way they deduplicate and stitch flow records. Plixer Scrutinizer overcomes the architectural constraints and design limitations of legacy solutions, helping NetOps and SecOps teams resolve issues and respond to threats faster and more effectively.

Only Scrutinizer supports dynamic key field deduplication

Legacy solutions deduplicate records based on a TCP/IP 5-tuple (source IP address/port number, destination IP address/port number, and protocol ID). As a result, they can inadvertently discard dissimilar records, drop valuable data, or mischaracterize traffic statistics. Plixer can dive deeper into a flow and deduplicate records and report on metadata beyond the standard 5-tuple, for more granular and precise reporting.

For example, say a unified communications solution is using QoS to optimize service quality for latency-sensitive voice and video traffic. As shown in the diagram below, the UC client tags packets with a certain DSCP value N, and router X re-tags and forward them upstream. During deduplication, a legacy collector might discard the router X flow record, dropping detailed troubleshooting information such as the original code point value set by the client. Plixer, in contrast, can dive deeper into the

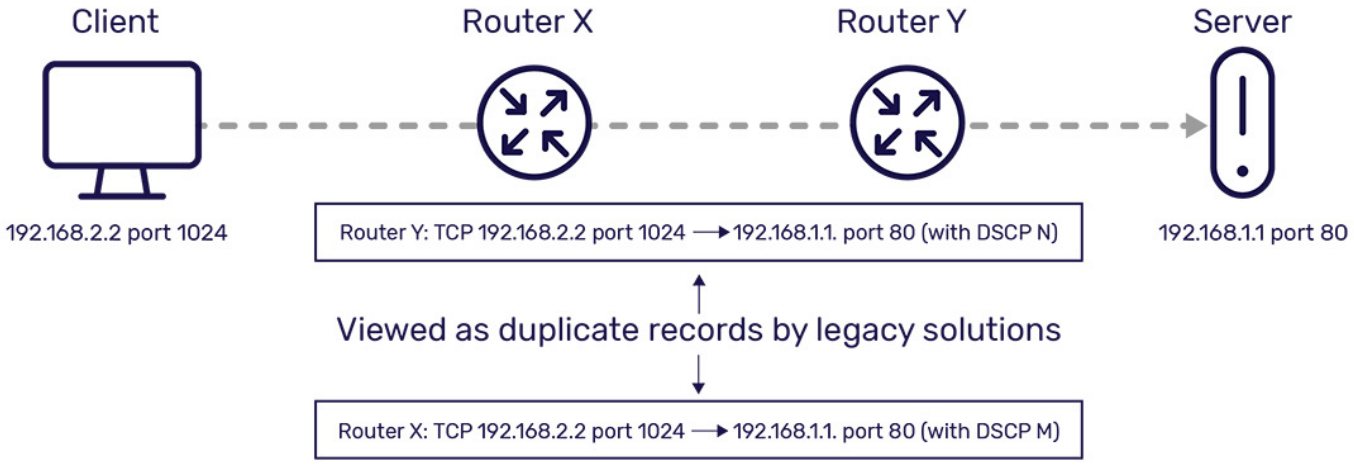

flow, determine that the records are in fact unique, and preserve the detailed information.

Here's another example. Say that a customer is using Virtual Route Forwarding (VRF) in a cloud or multitenant application. With VRF, a single router can be partitioned into multiple virtual routers with potentially overlapping IP addresses. As shown below, the router could generate two identical 5-tuple records for what are, in fact, two distinct IP flows.

A legacy network traffic analyzer solution might deduplicate the records, average their counters, and create a single record with a byte count of 200. Plixer, on the other hand, can dive deeper into the flow, determine the flows are unique, and provide distinct metadata and byte counts for each individual virtual router.

Dynamic key field deduplication also helps SecOps teams detect and respond to incidents more quickly and efficiently. In today's world, security platforms must analyze metadata elements beyond the 5-tuple so administrators can effectively monitor, detect, and respond to advanced persistent threats (APTs).

Plixer's method of deduplication and stitching maintains granular metadata within every flow, extending all the way to layer seven and providing a wealth of detailed information such as username, application

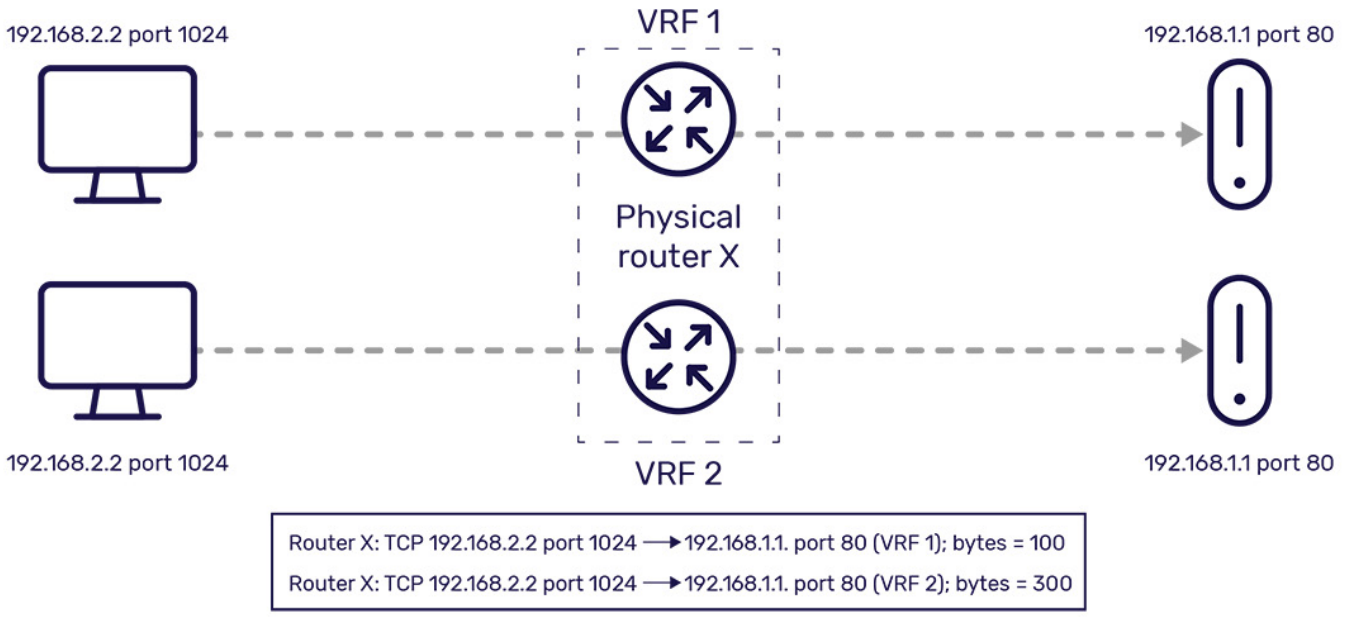

specifics, fully qualified domain name and other DNS information, URL/URI, and SSL certificate information to support forensics. SecOps teams can improve APT detection and accelerate incident response by leveraging Plixer's advanced security algorithms, and filtering and reporting on any metadata element.

Only Scrutinizer deduplicates and stitches flow records on-demand

Legacy network traffic analyzer solutions deduplicate and stitch all flow records at collection time. This approach wastes storage capacity (collectors maintain copies of both original flow records and deduplicated flow records) and squanders collector CPU cycles, which can impair performance and drive up collector hardware costs. Worst of all, the vast majority of records are deduplicated and stitched for naught; in practice only a small fraction of records are ever examined by NetOps and SecOps staff.

Unlike legacy solutions, Scrutinizer deduplicates and stitches flow records on the fly for greater efficiency and scalability. By only deduplicating and stitching flows when needed, Scrutinizer reduces collector storage and compute requirements, and enables customers to keep pace with ever-increasing NetFlow traffic densities and volumes as NetFlow support is added to more and more network elements, and vendor extensions proliferate.

Only Scrutinizer deduplicates and stitches records across collectors

With legacy network traffic analyzer solutions, each collector deduplicates and stitches flow records independently, which can complicate problem isolation and resolution efforts. Consider the asymmetric traffic flow in the diagram below, for example. Routers X and Y export records to collector B, while router Z exports records to collector A. With a legacy solution, network and security administrators must manually examine and correlate records from each collector individually—an inefficient and error-prone approach that hinders troubleshooting, root-cause analysis, and incident response.

Scrutinizer, in stark contrast, deduplicates and stitches flow records across collectors, providing full, end-to-end visibility across the entire extended

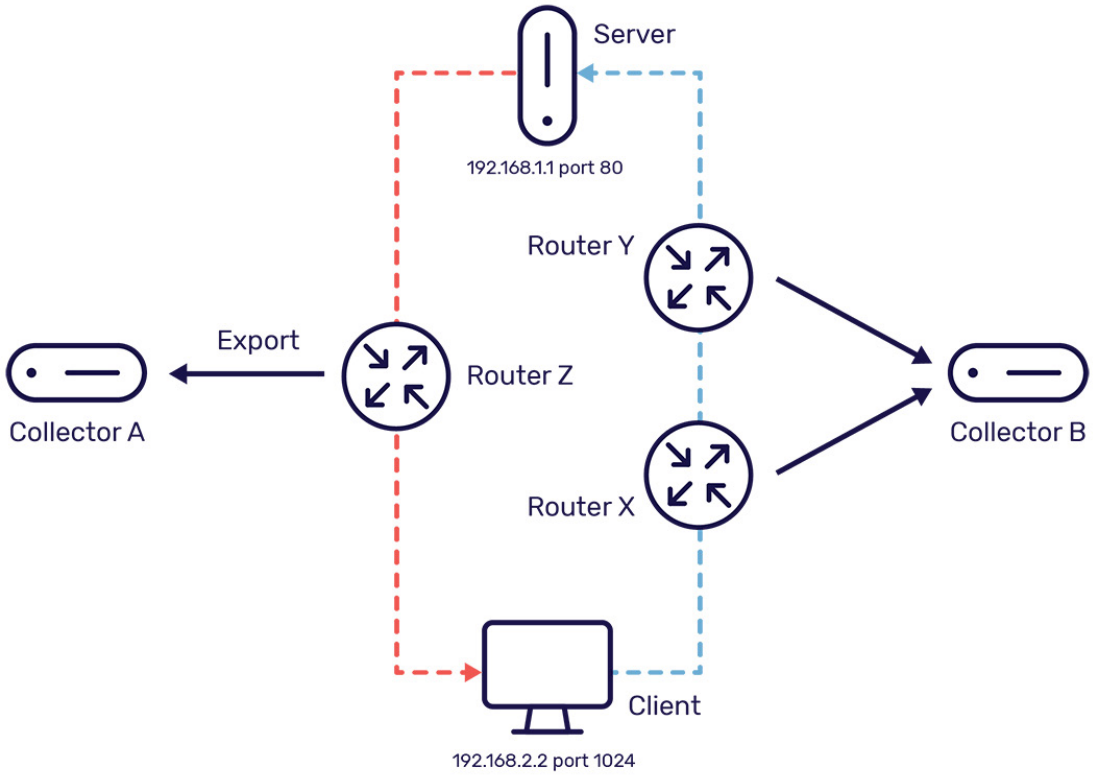

enterprise network through a single pane-of-glass. NetOps teams get a high-level, hop-by-hop view of a session and then drill down on an individual switch or router to get detailed information to diagnose and remediate an issue quickly and efficiently. And SecOps teams get a centralized view of any flow across the entire network for faster remediation. When security events occur, mean-time-to-know (MTTK) is paramount. Scrutinizer eliminates manually intensive, error-prone SecOps tasks, providing a holistic view of a conversation across the entire network, accelerating root-cause analysis and incident response.

Conclusion

Network and security operations professionals count on NetFlow traffic analysis solutions to monitor network health, diagnose issues, and mitigate threats. But when it comes to product usability and data accuracy and richness, not all NetFlow traffic analysis solutions are the same. Only Plixer Scrutinizer supports dynamic key field deduplication, and stitches and deduplicates records on-demand, across collectors for ultimate efficiency and visibility.

Avoid legacy NetFlow traffic analyzer woes. Gain accurate, actionable insights into critical NetOps and SecOps data, and accelerate remediation efforts with Scrutinizer.

To learn more about how Scrutinizer can help your company improve visibility, optimize performance, and mitigate risk, contact Plixer today.

About Plixer

Plixer provides a security and network intelligence platform that supports fast and efficient incident response. The solution allows you to gain visibility into cloud applications, security events, and network traffic. It delivers actionable data to guide you from the detection of security and network events all the way to root-cause analysis and mitigation. Network and security incidents are inevitable. When they occur, Plixer is there to help you quickly return to normal and minimize business disruption. Thousands of organizations rely on Plixer solutions to keep their IT infrastructure running efficiently. Learn more at plixer.com, stay connected with the Plixer blog, and follow us @Plixer.