riverbed

SOLUTION BRIEF

Sampled vs. Full-Fidelity Flow: Pros & Cons for Security

Flow-based metadata generated from sampled information leaves significant gaps in data collection which can lead to blind spots that makes it difficult to detect certain security attacks, like command and control and lateral movements.

For this reason and others, it is recommended that you use 100% flow recording (full fidelity) for security monitoring and forensics.

NetFlow

NetFlow is a protocol originally designed to collect IP network traffic as it enters or exits an interface. NetFlow statefully tracks flows (or sessions) and aggregates packets associated with each flow into flow records, which are then exported. NetFlow records can be generated based on capturing every packet (a.k.a. full-fidelity or 1:1 mode) or based on packet sampling. Sampling is typically employed to reduce the volume of flow records exported from each network device. While this practice allows you to deploy cheaper, lower spec'd telemetry solutions, it also effectively cuts corners on providing the complete view that is often needed for fully effective forensics. As such, Riverbed recommends against sampling if you are using flow for security, and instead, collect the raw flows whenever possible.

NetFlow does not forward flows directly. Instead, summaries of flows are cached and then exported based on active and inactive timeouts. Meaning, as flows are collected, the volumetric data builds up in the flow records. After some time has passed—usually about 1 minute—the record is sent to the collector. This means that information about ongoing conversations is exported with a delay. Many newer NetFlow exporters can be tuned to export at higher rights, however.

Wide Variety of NetFlow

There are many varieties and providers of NetFlow, including:

- IPFIX (Internet Protocol Flow Information Export) is based on NetFlow v9 and defines how IP flow information is formatted and transferred from an exporter to a collector.
- JFlow is Juniper's flow protocol, and there are other "xFlows" from a variety of vendors. For the purposes of this discussion, they are all the same, or very similar to NetFlow.
- SteelFlow is Alluvio's proprietary flow version that is based on NetFlow. It's always unsampled, so you never have to worry about whether you

have the right fidelity. Our WAN optimization and packet capture solutions can export SteelFlow with extensions for response time analysis, WAN optimization, and more. (While SteelFlow provides integration benefits as noted, Alluvio also supports a range of flows.)

What is sFlow?

sFlow is a traffic sampling flow technology that scales well for big and busy networks. It provides an industry standard for exporting truncated packets with interface counters.

The sFlow agent runs as part of the network management software within devices, such as routers or switches. It packages the data into sFlow datagrams that are immediately sent on the network to minimize memory and CPU requirements.

Selecting a suitable packet sampling rate is an important part of configuring sFlow. The table below gives suggested values that should work well for general traffic monitoring in most networks. However, if traffic levels are unusually high, the sampling rate may be decreased (e.g., 1 in 5000 instead of 1 in 2000 for 10Gb/s links).

Link Speed	Sampling Rate
10Mb/s	1 in 200
100Mb/s	1 in 500
1Gb/s	1 in 1,000
10Gb/s	1 in 2,000

There are legitimate reasons and use cases for using sFlow. But there are tradeoffs, too, especially when it comes to using sFlow for security or forensics. Flow-based metadata generated from sampled information leaves a big gap when it comes to information security. If we take our 10G link, for example, the flow data is generated by sampling 1 in 2000 packets. That means that 99.95% of traffic is not being viewed or stored. If a cybersecurity professional is looking for command and control activity or lateral movement in the network, they don't stand a chance of finding the vast majority of security issues with a 99.95% blind spot!

For this reason, when using flow metadata for cybersecurity analytics or forensics, it is essential the flow data be full fidelity. When you are looking for an indicator of compromise and come up empty-handed, full-fidelity flow means that you know for sure that what you're looking for isn't there. Sampled flow means you can never be sure—did you not find the IOC because it wasn't there? or because the data was sampled away?

When a breach or an attack occurs, network security professionals need robust investigative capabilities to be able to determine what happened as quickly as possible. Every moment spent in root cause analysis gives the attacker more time to burrow deeper into or across the network and adds to the overall loss from the attack (whether that be financial, customer impact or brand reputation). The net result is that sFlow simply doesn't provide the granularity and accuracy required to perform a full forensic investigation to determine exactly what happened before, during, and after a security breach.

Flow: No One-Size-Fits-All Answer

So, when it comes to network security, can you use sFlow? Or do you really need an unsampled NetFlow/IPFIX solution? Of course, you can. But, with cyberattacks on the rise, do you really want to scrimp when it comes to security? Let's reiterate: use 100% unsampled (full fidelity) NetFlow whenever possible for security purposes.

The bottom line is information security professionals leveraging sFlow for security should understand its significant blind spots, other limitations, and tradeoffs to manage their security tools and processes accordingly.

Combine Flow & Packets for Maximum Security

Whether you are using sFlow, NetFlow, IPFIX, or SteelFlow, you want to make sure that your flow collector gives you the best possible capabilities for performance monitoring and forensic investigations.

Unified NPM monitors network flow traffic in real time and immediately detects unusual behavior and deviations from “normal” patterns that indicate unwanted behavior on the network so you can act fast. Leverage Riverbed always-on, full-fidelity packet, flow and device metrics for cyber threat hunting, incident response and network forensics. After an attack, Unified NPM provides complete digital evidence to understand what happened, how long it was happening, and what other systems may have been affected.

For more information about Flow monitoring, please visit riverbed.com/products/npm/netprofiler.

riverbed

About Riverbed

Riverbed is the only company with the collective richness of telemetry from network to app to end user, that illuminates and then accelerates every interaction, so organizations can deliver a seamless digital experience and drive enterprise performance. Riverbed offers two industry-leading portfolios: Alluvio by Riverbed, a differentiated Unified Observability portfolio that unifies data, insights, and actions across IT, so customers can deliver seamless, secure digital experiences; and Riverbed Acceleration, providing fast, agile, secure acceleration of any app, over any network, to users anywhere. Together with our thousands of partners, and market-leading customers globally – including 95% of the FORTUNE 100 –, we empower every click, every digital experience. Riverbed. Empower the Experience. Learn more at riverbed.com.