Enhancing Flow Based Network Monitoring

Flow-based technologies such as NetFlow, sFlow, J-Flow, and IPFIX are increasingly popular tools used by network operators. The tools leverage the capabilities embedded within existing network switches, routers, and other network devices to obtain data about the information these devices process. However, flow-based approaches alone are often insufficient to solve specific problems experienced by end-users, and they may put additional strain on networks due to the additional data they generate. Flow-level data is great for trending and high-level analysis and some troubleshooting. However, access-to-packet data is required for detailed visibility into application performance, capacity planning, security threats, and end-user experience.

Flow Technology

Flow-based solutions consist of network equipment with embedded flow agents and flow collectors. Flow agents gather information about traffic on interfaces and forward it to flow collectors using the UDP protocol. A flow collector is a software application, running on a workstation or server, that collects the traffic data from a number of flow agents, stores the data, analyzes it, and presents the analysis to the network administrator in a variety of ways such as charts, dashboards, and thresholds. Many agents can send data to the same collector. Flow collectors may be incorporated with other systems which provide congestion control and troubleshooting, route profiling, audit trail security analysis, and accounting for billing.

What is a Flow?

Flow analysis examines each packet forwarded within a router or switch for a set of IP packet attributes. Traditionally, a flow is based on a set of 5 to 7 IP-packet attributes such as:

- IP source address
- IP destination address
- Source port
- Destination port
- Layer 3 protocol type
- Class of service
- Router or switch interface

All packets with the same source/destination IP address, source/destination ports, protocol type, interface, and class of service are grouped into a flow. This information is condensed into what is called a flow cache.

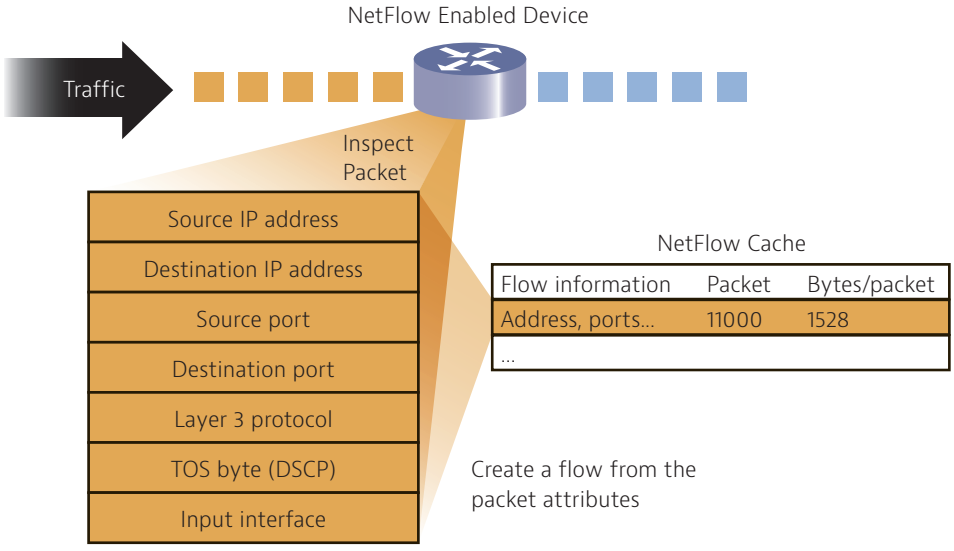

Figure 1. Attributes of a flow cache

Flow information is useful for understanding network behavior by denoting:

- Source address and destination addresses to understand traffic origin, including subnet masks
- Destination addresses
- Port information to characterize application traffic utilization
- Class of service to examine traffic priorities
- Device interface identification to indicate traffic utilization for specific network devices
- Tallied packets and bytes to show the amount of traffic
- Flow timestamps to understand the life of a flow and to calculate packets and bytes per second
- Next-hop IP addresses to understand routing topology

Flows are sent to collector servers from which collector software creates real-time or historical reports from the data. There are many commercial and freeware flow-reporting software applications available from a variety of suppliers.

Currently, there are a number of common flow implementations. NetFlow is the most common. Supported by Cisco® on its routers and switches, NetFlow sends information about completed traffic flows to a central collector. The device decodes every IP packet, maintains tables of active flows, and forwards flow records periodically or when they complete to a network management application. NetFlow has recently become adopted by the IETF Flow Information Export (IPFIX) standard as an approach that allows non-Cisco devices to send data to a NetFlow collector in a NetFlow-recognized format.

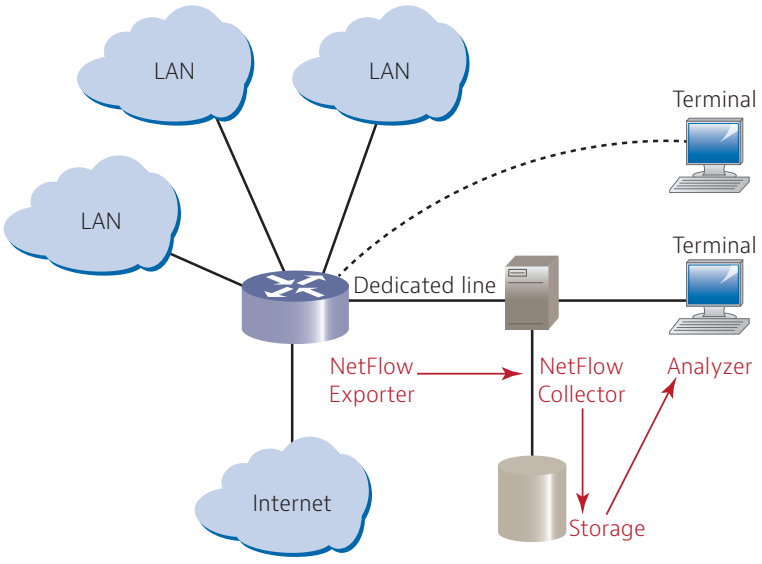

Figure 2. Types of flow-based reporting implementations

NetFlow v5 is the most popular and basic implementation. It provides basic information such as source and destination IP addresses, source and destination ports, and the byte counts of transferred data. Most Cisco devices running IOS 11.1 and above support NetFlow v5. NetFlow v9 adds support for different technologies such as multi-cast, IPSec, and multi-protocol label switching (MPLS). It is supported in IOS 12.4 and above. Flexible NetFlow added the ability to filter collected data. Supporting deep packet inspection (DPI) adds network-based application reporting (NBAR).

sFlow is an open standard that is based on RFC 3176. It was created by InMon, which provides sFlow collectors. sFlow uses statistical sampling of the state of the routing and bridging tables used by the switch to forward randomly-selected packets, where 1 of n (for example, n=1000) packets are copied from the network stream and sent to a collector/analyzer. Usually, only the first 128 or 256 bytes of the packets are sampled, since this is where the most important information resides, such as IP addresses and TCP/UDP sockets. Further information is available at: www.sflow.org.

NetFlow	IPFIX	sFlow	JFlow
· Developed by Cisco · Proprietary · Transit traffic and terminated traffic · Detailed information for each flow · Header only, no payload visibility · Sampled option · Flexible NetFlow introduces DPI · Vendors: 3Com, Adtran, Cisco. Enterasys, Expand, Juniper, nProbe, Riverbed, VMWare, Vyatta, and others	· Internet Protocol Flow Information eXchange · IETF standard · Based on Netflow · Header only, no payload visibility	· RFC 3176 – not a standard · Maintained by InMon · Statistical time-based sampling · Much less common than Netflow · Detailed information for each flow · Header only, no payload visibility · Sampled · Less expensive for vendors to implement · Vendors: 3Com, AlaxalA, Alcatel-Lucent, Allied Telesis, Brocade, D-Link, Extreme Networks, Enterasys, Force10 Networks, H3C, Hewlett-Packard, Hitachi, Juniper Networks, NEC, and others	· Developed by Juniper · Proprietary · Similar to NetFlow · Detaled information for each flow · Header only, no payload visibility · Sampled

Table 1. Various flow-based reporting implementations

Flow-Based Reporting Considerations

Flow-based approaches create certain considerations for network planners and administrators.

Network Blind Spots

NetFlow is the most pervasive flow technology, but it is typically implemented only in higher-end Cisco switches and routers. More cost-effective enterprise remote branch office routers may not include the capability or may require costly, specialized blades. For instance, most NetFlow implementations are not able to handle more than 10,000 packets per second unless a specialized module such as a Cisco multilayer switching feature card (MSFC) is used. In many areas of the network, particularly at the edge (such as enterprise remote branch offices), it is cost-prohibitive to implement. Similarly, service-provider networks are non-heterogeneous using non-NetFlow-capable infrastructure. Most if not all flow-enabled devices reside in the network core. With many reporting applications depending upon NetFlow for monitoring performance, compliance, or security, non-NetFlow capable devices create large blind spots on the network.

Overtaxing Infrastructure

Sending flow data may add overhead to already overtaxed routers and switches. A high volume of packets and high packet rates can tax processing performance on switches and routers. The risk of overloading infrastructure often prevents network engineers from enabling flow reporting on their network. They fear it may introduce symptoms of jitter and increased delay, possibly affecting services traversing these devices. Moreover, fine-tuning flow parameters adds complexity. For example, the techniques for tuning sampling rates in order to not impact device parameters such as memory buffers varies from device to device and increases in complexity depending on network size.

Limited Content Intelligence

Flow reporting is typically limited to routed traffic. For instance, in most cases flow data does not provide TCP timing and application response times. Flow monitoring provides visibility on network utilization: who/what/where/when. However, applications can no longer be identified by just L3/L4 information, and DPI capability is necessary. Application-affecting parameters such as window problems, application calls, response codes, and DNS and DHCP response times require packet-level visibility. For instance, many applications may share TCP Port 80. Without application-level visibility, it is impossible to determine which applications are traversing the network, to identify rogue or suspicious applications, or to ascertain application-specific traffic levels. While many flow-capable devices are introducing DPI, this only adds to the cost and complexity of implementing a holistic solution.

Costly and Complex External Appliances

Standalone probes are alternatives to flow collection from routers and switches. This approach can overcome some of the limitations of router-based flow monitoring. Probes are typically connected to a monitored link as a passive appliance using the TAP or SPAN port of the appliance. Typically, DPI is easier to implement in a dedicated probe than in a router—a purpose-built DPI appliance does not tax the router and impact packet-processing performance. However, probes must be deployed on every link that must be observed. This causes additional hardware, setup, and maintenance costs, while potentially introducing new points of failure.

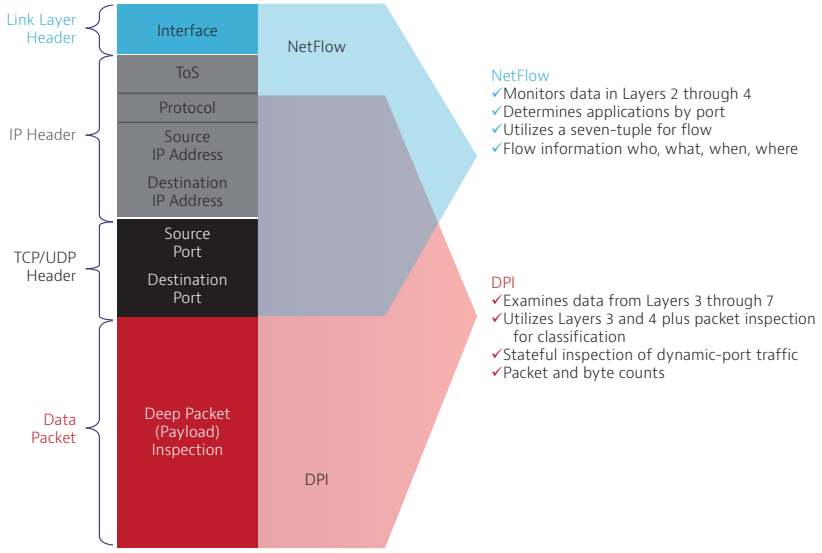

Figure 3. Characteristics of NetFlow and deep packet inspection packet processing

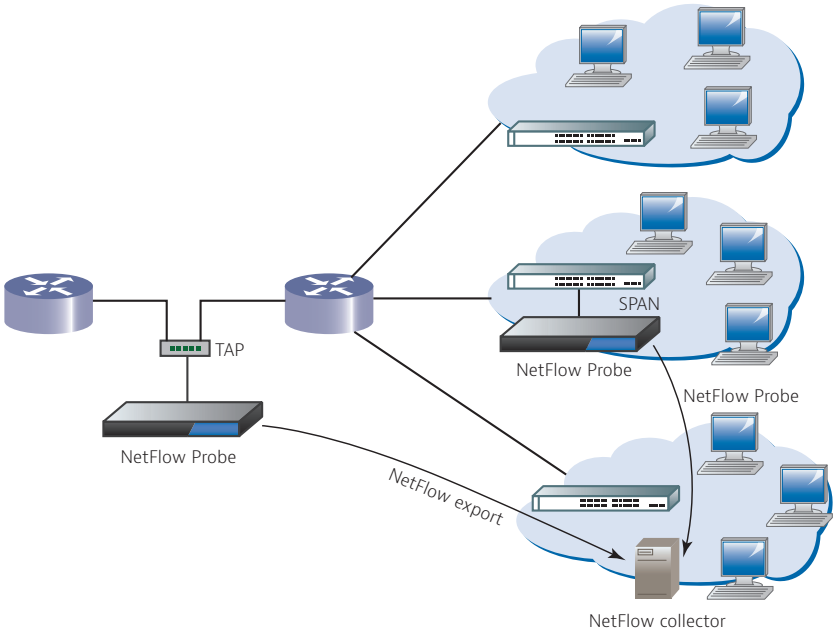

Figure 4. Typical NetFlow architecture

PacketPortal™

Viavi Solutions™ PacketPortal is a new software solution that decouples data capture from management and analysis. Distributed microprobes embedded in existing network elements such as switches, routers, and edge devices communicate over a self-forming, secure cloud to a centralized, open-software platform, providing unprecedented visibility into the network. This new visibility enhances existing troubleshooting tools, optimizing the customer experience for voice, video, and data services.

PacketPortal complements flow-based reporting strategies, providing packet-level visibility, extending visibility to blind spots at the network edges such as remote offices where implementing flow-reporting may not be possible. PacketPortal overcomes the cost, reliability, and complexity issues of using dedicated network probes while providing in-line content level visibility. PacketPortal adds deep-packet and content intelligence to any network device, providing consistent access to key information for monitoring and troubleshooting. It adds consistent packet capture to non-flow-capable devices.

PacketPortal also augments embedded flow capability available on routers or switches. It provides packet-level visibility, offloads performance taxes, and forwards critical data to flow-analysis applications. By offloading flow sampling, PacketPortal frees the network device to do what it does best: process and forward packets. PacketPortal, in combination with flow, can provide a holistic view from Layers 2-7 to reduce the time it takes to identify, diagnose, and resolve complex performance issues. PacketPortal minimizes the traffic load on the network by only sending filtered data and adds no additional processing tax on network devices that support PacketPortal-enabled SFProbes™.

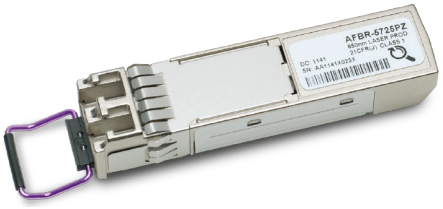

The PacketPortal SFProbe

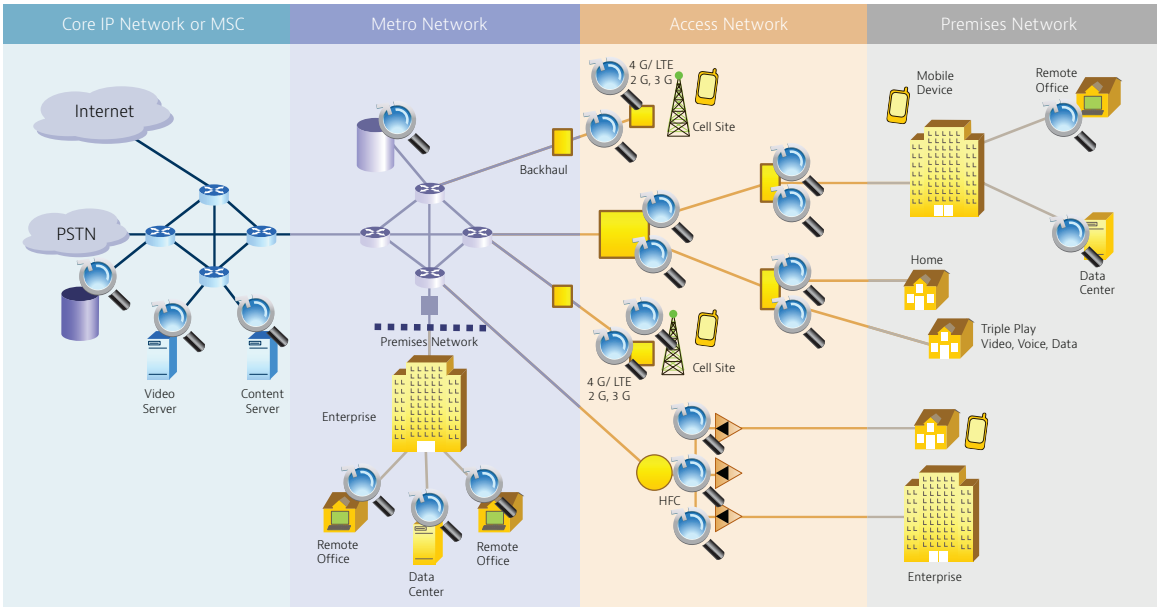

Intelligent Data Collectors Deployed Throughout the Network

PacketPortal Enabled Flow

PacketPortal complements flow capabilities embedded in network devices, and in combination with flow-analysis software tools, can deliver a number of key benefits. These include but are not limited to:

- **Congestion Control** — By monitoring traffic flows on all ports continuously, flow can be used to instantly highlight congested links and identify the source of the traffic and the associated application-level conversations. PacketPortal can be used to filter data used on aggregated TCP ports, such as port 80, forwarding this data to network- and flow-analysis tools. Armed with information about both links and applications, network operators can determine effective controls. For example, they can determine which application traffic to rate-control or prioritize or where to provision more bandwidth.
- **Security and Audit-Trail Analysis —** A comprehensive security strategy involves protecting the network from external and internal misuse and protects information assets from theft. Accomplishing this requires complete network surveillance with alerts to suspicious activity. PacketPortal can enhance basic flow-provided information to enable continuous, network-wide surveillance and route-tracing information. Upon detection of anomalies using flow, PacketPortal can drill down to isolate specific issues.
- **Route Profiling and Capacity Planning —** Since flow contains forwarding information, it can be used to profile the most active routes and the specific flows carried by these routes. PacketPortal helps isolate application details within flows. Understanding routes, flows, and details within flows makes it possible to optimize route performance—improving connectivity and performance, and choosing the most cost-effective peering partners based on a highly granular basis.
- **Accounting and Billing for Usage —** Detailed network usage information is needed to fairly charge for network services and to recover the costs of providing value-added services. Flow data can be used to account and bill for network usage by customer. PacketPortal adds transaction, content, and message-level billing which can also be used to provide customers with an itemized breakdown of their total traffic, highlighting top users and applications.

The combination of flow and PacketPortal lets network operators implement network-wide monitoring with deep content intelligence, extending capability to remote locations where previously no detailed monitoring may have been possible. PacketPortal augments the capability of flow-capable network routers and switches, providing critical information for capacity planning, historic data collection and traffic analysis, network performance analysis, and unified visibility across networks.

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the Viavi office nearest you,
visit viavisolutions.com/contacts.

packetportalflow-wp-nsd-tm-ae
30173384 900 0213