



Frequently Asked Questions on NetFlow

What is NetFlow?

NetFlow is a light-weight data feed developed by Cisco and generated by most of their routers and layer-3 switches. NetFlow summarizes every “flow” (a one-way IP to IP application-specific communication) that traverses the router as a flow-packet data unit (pdu). The flow-pdus are aggregated as NetFlow records and exported to destinations as configured on the router. The destinations receive the NetFlow data feed as a stream of UDP packets.

What information does a NetFlow flow-pdu generally contain?

- Source and destination IP addresses
- Protocol type (TCP, UDP, ICMP, and others)
- Source and destination port numbers (if applicable)
- Byte and packet counts
- Start and end times
- Input and output interface numbers
- Type of service (ToS)
- TCP flags and encapsulated protocol
- Routing information

What advantages does NetFlow have over other data feeds/sources out there?

- Created by existing infrastructure (i.e. routers)—neither agents nor probes needed
- Generated effortlessly by routers/switches
- Produced in a timely and continuous fashion
- Pushed as a data feed to local or remote IPs
- Ease of configuration on routers—simple commands through CLI
- Comprehensive characterization of activity on the network

Do other vendors support NetFlow?

Yes other router/switch vendors support the NetFlow format, The Internet Engineering Task Force (IETF) has created an open standard around the format of NetFlow version 9 which is referred to as IPFIX. In addition, other vendors support distinct “flow data” formats that fall under the de-facto term NetFlow. These formats include sFlow, jFlow and cFlow. Xangati’s solution supports all these formats.



What if my router or switch doesn't support NetFlow and I want to take advantage of this technology?

You can have another device generate NetFlow for you. These devices will connect to a "mirror port" (also known as a SPAN port) on your router and/or switch and convert the traffic they receive into NetFlow-pdus.

What are some general uses of NetFlow?

- Application management
- Business intelligence on application and end-user activity
- Bandwidth utilization monitoring
- Proactive management
- Capacity planning
- Rapid problem identification
- Usage-based billing
- Visual trouble ticketing
- Service Level Agreement (SLA) monitoring
- Security analysis
- End-user experience monitoring

What makes Xangati's Application Management 2.0 solution different than other solutions that consume NetFlow?

The Xangati solution is an Application Management 2.0 solution or next generation application management product. To see how it compares against traditional 1.0 solutions including NetFlow collectors, you can view this grid:

http://www.xangati.com/products_2_dot_0_vs_1_dot_0.php

Although Xangati's solution and NetFlow collectors both support NetFlow as a key data feed, they were designed to do very different things with that fuel. Maseratis and lawn mowers both consume gasoline – but do very different things with it. The Xangati solution is optimized to provide real-time insight into applications and their interactions from the end-user to the end device. NetFlow collectors were designed to provide report data on the capacity utilization, by application, of specific network interfaces. And, because these solutions serve two different purposes, the architectures behind them are very different. The Xangati solution will process the NetFlow records in real-time to present a streaming to-the-second visual of what is occurring on your infrastructure—it does not store the NetFlow record. In contrast, the NetFlow "collectors" will index and store the NetFlow record to allow it to be accessed for report generation. In addition to the streaming views, Xangati does also provide reports (even on network interfaces), but its reports are from a database of summarized information not from the full flow record—meaning much faster report generation.



Do I have to worry about NetFlow affecting the performance of my router?

No. Cisco and the other internetworking vendors have put tremendous effort into optimizing the efficiency of NetFlow (and its flow data equivalents). The NetFlow process has 3 parts to it:

1. Characterization of the flow information in the NetFlow cache
2. Formation of the NetFlow export record
3. The export (push) process

Many of the router/layer 3 switches today have ASIC technology which supports the majority of parts 1 and 2 in hardware. For these ASIC-based solutions, the load on the CPU is entirely inconsequential.

For software-only solutions, Cisco has put together some very clear documentation on best practices for NetFlow. One particularly useful guide is the “*NetFlow Services Solutions Guide*” <http://tinyurl.com/nz9zke> which has a table that documents the incremental CPU utilization of NetFlow. The CPU utilization is dependent on the number of active flows and as you can see from the table below that the incremental CPU load is very minimal even when working with a high flow count.

Number of Active Flow in Cache	Additional CPU Utilization
10000	< 4%
45000	<12%
65000	<16%

(Source: Table 9 Cisco's *NetFlow Services Solutions Guide*)

Do I have to worry about the load that NetFlow will put on the network?

No, according to the Cisco “*NetFlow Services Solutions Guide*” (<http://tinyurl.com/nz9zke>). NetFlow data typically does not account for greater than 1.5% of the bandwidth traversing the network. As a sanity check, it is worth noting that the Xangati solution can see flow data as another application and can show you it is only consuming a very minimal amount of bandwidth.

How do I configure NetFlow on my routers?

Here is a link to various configuration guides produced by Cisco. <http://tinyurl.com/m7wh2m>



Does NetFlow sampling affect the accuracy of the data that I am receiving?

No. Sampling does not affect the accuracy of the flow data. This research paper entitled " *On the Accuracy and Overhead of Cisco Sampled NetFlow*" shows statistically the accuracy of sampled NetFlow. <http://tinyurl.com/m7wh2m>

You keep referring to applications but doesn't NetFlow deliver information in terms of protocol ports? How do you explain that?

There are a number of ways the Xangati system leverages NetFlow data to define applications in addition to just protocol ports.

The first is by leveraging the IANA (Internet Assigned Numbers Authority) which specifically has a registry of protocol ports that have been assigned to explicit applications. Xangati uses this IANA registry to name about 100 different applications out of the box.

The second way applications can be identified is by leveraging the explicit IP addresses of the servers generating the specific application (or the IP address block of a specific content provider like Google). The IPs or IP address block of the servers plus the protocol port (or port range) in effect allow for the definition of an application. An example of where this is particularly helpful is web-based applications (port 80/8080) because many internal IT applications have a web-front end and would be indistinguishable by protocol port alone. The Xangati model allows these various web apps to be segmented and efficiently tracked individually.

If I already have a solution that is collecting flows, can I also get that data to Xangati?

Many devices that collect NetFlow information can also relay the information they receive to another device. All you have to do is configure that device to relay the information to the IP address you have given your Xangati appliance. If the current device does not support flow relay, then you can re-configure your NetFlow sources to push the flow records to Xangati who can then relay them to your existing flow collecting device.