Cisco NetFlow™ Briefing Paper

Release 2.2

Monday, 02 August 2004

Apoapsis, The Bridge, 12-16 Clerkenwell Road, London, EC1M 5PQ, UK
Tel: +44 20 7324 6050 Fax +44 20 7324 6051 Email info@netusage.net

Contents

Executive Summary

This document is aimed at network managers and engineers responsible for enterprise scale IP networks, and will review how Cisco NetFlow technology functions, how it can be deployed to solve management problems, and how it compares with other more traditional traffic monitoring and analysis techniques.

Cisco NetFlow is an extremely cost effective and efficient IP network monitoring and traffic analysis technology. It provides the type of granular data required by enterprise network professionals for network monitoring, capacity planning, and chargeback.

The problem

Enterprise Network Managers have a challenging job. They are faced with increasingly complex demands on their network infrastructure, and need to find a balance between satisfying the expectations of users, meeting strategic business objectives, and complying with the financial restrictions of their allocated budget.

These demands are especially challenging in the WAN where lower bandwidths and higher costs combine to create headaches and magnify user perceptions of problems, whether with performance or cost effectiveness.

To perform their job competently the network manager needs to be able to answer questions such as:

- What applications are running on the network, and how much bandwidth is each using?
- Can a new application be rolled out safely, without impacting the rest of the network?
- When applications run slow is it the network that is at fault?
- Does a WAN circuit need to be upgraded, or can traffic flows be optimised to make do with existing bandwidth?
- How much bandwidth each business unit within the enterprise is consuming, and is the correct information available to enable chargeback based on usage?
- How do you justify upgrades to business managers - what hard evidence is available?

And over-riding all of this is the issue of how the information necessary to answer these questions can be obtained without spending vast amounts of money on monitoring equipment and staff, both now and in the future as the network scales.

The traditional solutions

Answering questions like those listed above hinges on the availability of accurate and comprehensive management information and statistics. Five techniques have traditionally been used to obtain this vital management data:

- SNMP based monitoring: Most enterprises will have some sort of SNMP based management system deployed. This will poll network devices for selected information from their SNMP MIBs, and provide methods for graphing or otherwise reporting on this data. Key advantages include universal support across multiple vendors networking hardware and access to "inside the box" data such as the number of packets an interface has dropped. Disadvantages include the traffic and CPU loading generated by frequent polling, and the lack of any ability to break traffic down for accounting purposes, such as by application or by destination or source IP address

- Integrated RMON: Generally only RMON1 is supported in most network devices. This captures traffic statistics at the MAC layer, and provides useful standard reports such as "top-talkers" and "conversations". The standard applies to only Ethernet and token ring LANs, and is widely supported on Ethernet switches. However support on routers is very limited – for example most Cisco routers support only the alarm and event groups, and only on Ethernet interfaces, so this method cannot be used to monitor WAN traffic. Operation at the MAC layer also means traffic cannot be broken down by protocol or application, or by source or destination IP address. This data requires use of RMON2, for which external probes are required.

- External Probes: Provide RMON2 analysis and often additional vendor proprietary techniques. As RMON2 gathers statistics for Layers 2 through to 7 of the OSI model an external probe can break traffic reporting down by protocols, applications, and source and destination IP address. RMON2 also supports capture and decode of individual packets and can monitor other protocols besides IP, for example IPX.
  As an external hardware solution however, external probes are unaware of "in the box" data such as interface packet drops, and add substantially to the cost of the overall network monitoring solution. This cost element means a probe solution will not scale to general deployment across large networks – it is only feasible to deploy a small number of probes and move them to whichever point in the network is causing most problems at any given time. In addition each probe will support only 1 or 2 physical interface types, so if multiple WAN technologies are in use then multiple probes will be required, whereas Cisco NetFlow is supported internally on all router interfaces. In addition inserting a WAN probe into a circuit causes an outage, whereas Cisco NetFlow can be enabled without disrupting service.

- Network Analysers or "Sniffers": Some organisations try to make use of tools such as network "Sniffers". These are designed for troubleshooting, and as such they are optimised for detailed analysis of short-term "snapshot" captures. They cannot supply the long-term history required for trend and capacity analysis. They are also manpower intensive, and do not lend themselves to centralisation of management reporting and data storage. They are not viable for a true enterprise solution.

- IP Accounting: This is a Cisco router specific feature, which provides the number of packets and bytes switched through the router on a source and destination IP address basis. Only transit

traffic passing through the router is recorded – traffic generated or terminated by the router itself, for example routing updates or SNMP poll replies, will not be recorded. There is no notion of flows, and data cannot be broken down by application or protocol.

In addition to these five traditional techniques, a further option for Enterprises with Cisco devices in their network is the use of a monitoring technology known as Cisco NetFlow. The rest of this document will discuss this technology in detail.

Comparison with other techniques

Table 1 compares the key monitoring technologies.

Feature	SNMP	Integrated RMON1	External probes	Portable Analyser or "Sniffers"	IP Accounting	Cisco NetFlow
Router CPU Impact	Low	Low	Zero	Zero	Low	Low To Medium
Network Traffic Impact	Low To Medium[1]	Low	Low	Zero	Low	Low
Configurable Data Aggregation at Source	No	Yes	Yes	Yes	No	Yes
Real-Time Reporting	Near-Realtime[2]	Near-Realtime[2]	Near-Realtime[2]	Yes	Near-Realtime	Near-Realtime[3]
Extra Hardware required	No	No	Yes	Yes	No	No
Standard or Proprietary Technology	Standard	Standard	Proprietary & Standard (RMON2)	Proprietary	Proprietary	Proprietary But Published
Access to "in-the-box" Data	Yes	Yes	No	No	No	Yes
MAC Layer Data	No	Yes	Yes	Yes	Yes[4]	No
Layer 3 and Above Data	Very Limited	No	Yes	Yes	Layer3, No Higher	Yes
Support for Non-IP Traffic	Very Limited	Yes	Yes	Yes	No	No
Multicast IP Support	Limited	Limited, Not IP Layer Aware	Yes	Yes	No	Future (In IOS 12.3)
Data Suitable for Usage Based Chargeback and Billing	No	No	Yes	No	No	Yes
Information for Individual IP Flows	No	No	Yes	Yes	No	Yes
Support for Wide Range of Interfaces (Serial, Ethernet, ISDN, Frame Relay Sub-interfaces etc.)	Yes	No – Ethernet And Token Ring Only	Yes, But Must Buy Multiple Probes	Yes, But Must Buy Multiple Units	Yes	Yes
Push or Pull Access to Traffic Data	Pull	Pull	Pull	N/A	Pull	Push
Centralised Control	Yes	Yes	Yes	No	Yes	Yes
Scalable to Large Networks	Yes	Yes	No	No	Yes	Yes

1. Depends on polling interval and volume
2. Depends on polling interval
3. Depends on cache timeouts configured on router
4. If IP MAC accounting enabled

Cisco NetFlow overview

Cisco NetFlow is a Cisco originated technology, however the interface has been published to encourage development of third party applications in what Cisco terms a "Partner Ecosystem". On Cisco routers Cisco NetFlow is built into the IOS software. It allows a router to capture data about IP traffic as it enters the router on selected "ingress" interfaces. This data is extremely granular and accurate, and allows IP traffic flows to be analysed to a detailed level without the expense of external probes.

Data is captured on a 'per-flow' basis, where a flow is defined as a unidirectional stream of IP data with the same values for the following seven fields:

- Source and destination IP address
- Source and destination port numbers
- IP protocol type (e.g. 6=TCP)
- Type of service (ToS)
- Input interface

A cache entry is created for each new flow, and maintained as long as the flow remains active. In addition to the fields listed above, each cache entry includes information such as packet and byte counts, start and end timestamps, output interface, and routing information such as next-hop address, and source and destination AS numbers. There are a number of Cisco NetFlow record formats to choose from, which determine the fields that are included in the cache. As each packet passes through the router the accounting parameters for the relevant flow are updated.

The Cisco NetFlow architecture consists of three elements:

- Flow Caching and Data Export:
  Occurs in the router, which stores data about each IP flow in the flow cache, and prepares data for export to external analysis systems. Data is exported when a flow expires, or when a configurable timer expires (30 minutes by default), whichever comes first. Data can also be aggregated before export to reduce the amount of data export traffic, using a range of user selectable aggregation schemes.

- Flow Collection:
  An external server or 'mediation device' receives the flow data exported from multiple routers, filters and aggregates it as defined by the user, and then stores the processed data for later retrieval and analysis. This allows Cisco NetFlow data collection to scale to suit large networks.

- Data Analysis:
  A tool that allows users to view and analyse the data has been stored on the Flow Collector. Tools are available for network traffic analysis, capacity planning, and accounting and billing. Typically these provide a graphical user interface, and can perform near-real-time traffic visualisation and long-term trend analysis.

The Apoapsis NetUsage™ Console appliance fulfils both the Flow Collection and Data Analysis functions, in a 1U rack mounted unit. No other software or hardware is required – all reports are web based and available from any browser. Refer to Appendix A for further information.

Cisco NetFlow Record Format Versions

Flow data being exported is grouped together so that multiple flow-records are sent in a single UDP Cisco NetFlow Export datagram. The format of these export packets, including the number of flow records contained in each export packet, is determined by the record format version configured. The user can select this from a choice of five options:

- Version 1: the original format, which should only be used with legacy flow collector that doesn't support later formats.
- Version 5: improved format including BGP AS details and a sequence number to allow detection of lost datagrams.
- Version 7: only for Catalyst switches when running MLS – 5500 with NFFC or 6500 with MSFC.
- Version 8: used for export of aggregated records.
- Version 9: recently announced flexible format, which adds the concept of templates, allowing further addition of new fields without the need to release new format versions. As new Cisco NetFlow features are added applications do not need to be recompiled to support them – rather an external data file can be accessed which lists the templates for the new features.

In a typical installation Cisco NetFlow Version 5 or 7 will be configured, depending on the hardware platform, and then Version 8 added later if router based aggregation is required.

Router Based Aggregation

Cisco NetFlow aggregation enables the routers to export only a subset of the total data by consolidating it locally. There are a number of different aggregation schemes that can be configured. The aggregated data is exported using the Version 8 record format. This technique reduces network traffic generated, the CPU load on the router, and the workload on the Flow Collector.

Advantages and Disadvantages of Cisco NetFlow

Advantages of Cisco NetFlow:

- A very low cost solution:
  - Cisco NetFlow is integral to IOS (although an appropriate image and/or feature licence is required)
  - No costly external probe hardware required
- Statistics are gathered before compression and/or encryption, thus providing data on flows that would not be accessible to external probes
- Supported on most router interfaces, including WAN, MAN, LAN and tunnel interfaces
- Works on sub-interfaces e.g. frame relay
- The best solution for granular detail on IP packet flows
- No polling required, information is "pushed" when available
- Traffic can be aggregated in the router before export reducing traffic volumes

- Provides duration and absolute timestamps for each flow
- Wide support on Cisco router platforms, and on 5500 with NFFC and 6000 with MSFC
- Makes usage based chargeback and billing possible
- Widely deployed (Cisco state that it has been deployed by over 15,000 customers)

Disadvantages of Cisco NetFlow:

- IP only – does not monitor IPX or other protocols
- Monitors ingress traffic only
- Does have an impact on router CPU, which imposes a limit on the total number of flows that can be monitored on any single device
- Only supports unicast traffic until IOS release 12.3

Impact on Network Infrastructure

Impact on "no-drop" packet switching rates:
The no-drop packet rate is the maximum rate at which packets can be switched through the router before packets start to be dropped.

Cisco has published some results for the impact of enabling Cisco NetFlow on an RSP2 with 128Mbytes memory. Using the worst-case packet size of 64bytes the total rate was 216kpps (kilo packets per second) using the optimum switching path, and 120kpps after enabling Cisco NetFlow with data export. Thus the no-drop rate was degraded by approximately 44%.

CPU impact:
Cisco has also published a white paper on Cisco NetFlow Performance Analysis (check www.cisco.com), detailing the impact on CPU levels across a range of routing platforms. A range of tests was performed on each platform and the results compared. One of the main conclusions was that for any given platform the primary factor is the number of flows being processed, rather than which Cisco NetFlow version or features are configured. For example it made little difference whether Version 5 or 8 was used, or how many destinations where configured for exported Cisco NetFlow data. Table 1 shows a summary of the CPU impact for each platform at various flow volumes.

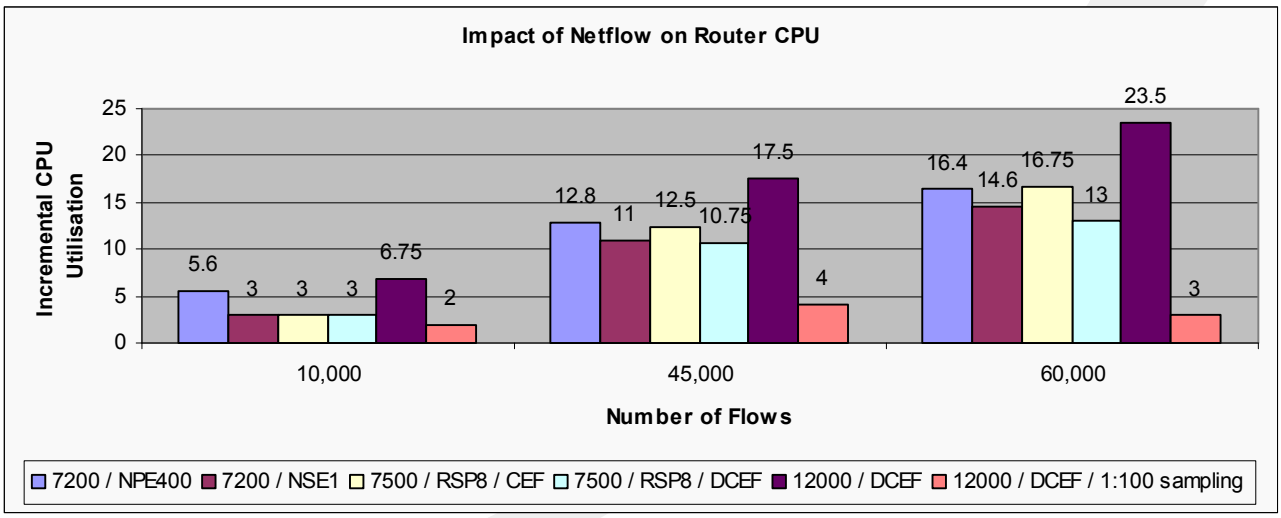

Table 2 - Impact of Cisco NetFlow on Router CPU

It should be noted that all platforms tested were configured with high-end processors cards, for example the RSP8 on 7500, and the NPE400 on the 7200. In practice many enterprise networks will have an installed base of routers with lower specification processors, and should thus expect to see increased CPU impact. For this reason these figures should be used only as a guide to identify which platforms will require lab testing or pilot deployments. If CPU levels are a concern then Cisco NetFlow can be enabled in a phased manner, one interface at a time, whilst monitoring the impact.

Cisco recommend running with a CPU level at or below 60% after enabling Cisco NetFlow. This allows for growth and traffic surges.

Traffic impact

The volume of traffic generated by enabling Cisco NetFlow can be estimated after enabling Cisco NetFlow caching but prior to enabling data export. This is a four-step process:

- Use the "*show interface switching*" command to determine the average throughput in packets per second, for all relevant interfaces
- Use the "*show ip cache flow*" command to verify the average number of packets per flow
- Calculate the total number of flow per second = average packets per second / average packets per flow
- Use the number of flows per export packet and export packet length (these vary from 24 to 51 fpp and from 1200 to 1500 bytes respectively, depending on record format version) to calculate total exported traffic:
  total traffic in bytes = (flows per second / flows per packet) * packet size

To illustrate with an example

- Number of packets per second switched on an interface (from *show interface switching*) = 100pps
- Average packets per flow (from *show ip cache flow*) = 20ppflow
- Then flows per second on the interface = 100/20 = 5fps
- If Version 5 format is used, then there are 30 flow records per packet, and a packet size of approx 1500 bytes, so traffic volume generated = 5/30 * 1500 = 250Bps (bytes per second), or 2kbps (kilobits per second)

This calculation can be repeated for each Cisco NetFlow enabled interface to obtain the total impact.

Conclusions

Cisco NetFlow is an extremely effective IP network monitoring and traffic analysis technology. It provides the granular flow based data required for applications such as billing and chargeback, network monitoring, and capacity planning. The raw data can be processed using tools from Cisco or any one of a range of 3rd parties, including the Apoapsis NetUsage solution. The fact that the technology is built into Cisco's IOS software and is supported across a wide range of physical WAN and LAN interfaces makes it easy to deploy without requiring network infrastructure upgrades.

However no single monitoring technology will provides a complete solution. A full management solution should include a combination of elements. A realistic approach would involve the following:

- SNMP for non-traffic statistics such as router CPU, memory, and interface errors
- Cisco NetFlow for flow based traffic statistics for network monitoring and support, planning and analysis, and chargeback
- One or more portable analysers or external probes for detailed packet capture and troubleshooting

Sources of Further Information

NetUsage™ products:

Web Site

- NetUsage™ www.netusage.net
- Cisco NetFlow™ www.cisco.com/go/netflow

Product Sheets

- NetUsage Product Suite Overview www.netusage.net/products.htm
- NetUsage Traffic Reporter www.netusage.net/traffic_reporter.htm
- NetUsage Business Reporter www.netusage.net/business_reporter.htm
- NetUsage Cost Reporter www.netusage.net/cost_reporter.htm

White Papers

- Cisco NetFlow Briefing Paper www.netusage.net/case_studies.htm
- WAN Deployment Paper www.netusage.net/case_studies.htm

Contact Address

Apoapsis Limited
The Bridge
12-16 Clerkenwell Road
London EC1M 5PQ
Email: info@netusage.net

