NetFlow Services and Applications

Introduction

Rapid growth in Internet and intranet deployment and usage has created a major shift in both corporate and consumer computing paradigms. This shift has resulted in massive increases in demand for network bandwidth, performance, and predictable quality of service as well as multimedia and security oriented network services. Simultaneously, the need has emerged for measurement technology to support this growth by efficiently providing the information required to record network and application resource utilization. Cisco's NetFlow services provide solutions for each of these challenges.

NetFlow also provides the measurement base for Cisco's Internet and Enterprise Quality of Service (QoS) initiatives. NetFlow captures the traffic classification or precedence associated with each flow, enabling differentiated charging based on Quality of Service.

NetFlow Definitions and Benefits

A network flow is defined as a unidirectional sequence of packets between given source and destination endpoints. Network flows are highly granular; flow endpoints are identified both by IP address as well as by transport layer application port numbers. NetFlow also utilizes the IP Protocol type, Type of Service (ToS) and the input interface identifier to uniquely identify flows.

Non-NetFlow enabled switching handles incoming packets independently, with separate serial tasks for switching, security, services and traffic measurements applied to each packet. With NetFlow-enabled switching, security (ACL) processing is applied only to the first packet of a flow. Information from the first packet is used to build an entry in the NetFlow cache. Subsequent packets in the flow are handled via a single streamlined task that handles switching, services and data collection concurrently.

Thus, NetFlow Services capitalizes on the flow nature of traffic in the network to:

- Provide detailed data collection with minimal impact on router performance and
- Efficiently process access lists for packet filtering and security services

NetFlow enables several key customer applications:

- *Accounting/Billing*—NetFlow data provides fine-grained metering (e.g. flow data includes details such as IP addresses, packet and byte counts, timestamps, type-of-service and application ports, etc.) for highly flexible and detailed resource utilization accounting. Service providers may utilize this information to migrate away from single fee, flat rate billing to more flexible charging mechanisms based on time-of-day, bandwidth usage, application usage, quality of service, etc. Enterprise customers may utilize the information for departmental chargeback or cost allocation for resource utilization.
- *Network Planning and Analysis*—NetFlow data provides key information for sophisticated tools such as Netsys to optimize both strategic network planning (e.g. who to peer with, backbone upgrade planning, routing policy planning)

CISCO SYSTEMS

as well as tactical network engineering decisions (e.g. adding additional VIPs to routers, upgrading link capacity) —minimizing the total cost of network operations while maximizing network performance, capacity and reliability.

- *Network Monitoring*—NetFlow data enables extensive near real time network monitoring capabilities. Flow-based analysis techniques may be utilized to visualize traffic patterns associated with individual routers and switches as well as on a network-wide basis (providing aggregate traffic or application based views) to provide proactive problem detection, efficient troubleshooting, and rapid problem resolution.
- *Application Monitoring and Profiling*—NetFlow data enables network managers to gain a detailed, time-based, view of application usage over the network. Content and service providers may utilize this information to plan and allocate network and application resources (e.g. Web server sizing and location) to responsively meet customer demands.
- *User Monitoring and Profiling*—NetFlow data enables network managers to gain detailed understanding of customer/user utilization of network and application resources. This information may then be utilized to efficiently plan and allocate access, backbone and application resources as well as to detect and resolve potential security and policy violations.
- *NetFlow Data Warehousing and Mining*—NetFlow data (or derived information) can be warehoused for later retrieval and analysis in support of proactive marketing and customer service programs (e.g. figure out which applications and services are being utilized by internal and external users and target them for improved service, advertising, etc.). This is especially useful for Internet Service Providers (ISPs), as NetFlow data enables them to create great depth in their service packaging. In addition, NetFlow data gives Internet Market Researchers access to the "who", "what", "where", and "how long" information relevant to Internet consumers.

NetFlow Cache Management and Data Export

The key to NetFlow-enabled switching scalability and performance is highly intelligent flow cache management, especially for densely populated and busy edge routers handling large numbers of concurrent, short duration flows. The NetFlow cache management software contains a highly sophisticated set of algorithms for efficiently determining if a packet is part of an existing flow or should generate a new flow cache entry, dynamically updating per-flow accounting measurements residing in the NetFlow cache, and cache aging/flow expiration determination. Rules for expiring NetFlow cache entries include:

- Flows which have been idle for a specified time are expired and removed from the cache
- Long lived flows are expired and removed from the cache (flows are not allowed to live more than 30 minutes by default, the underlying packet conversation remains undisturbed)
- As the cache becomes full a number of heuristics are applied to aggressively age groups of flows simultaneously
- TCP connections which have reached the end of byte stream (FIN) or which have been reset (RST)

Expired flows are grouped together into "NetFlow Export" UDP datagrams for export from the NetFlow- enabled device. NetFlow Export datagrams may consist of up to 30 flow records for version 5 flow export (25 flow records for version 1 flow export, 28 flow records for version 7 flow export [used only by Catalyst® 5000 NetFlow Feature Card]). Flow datagrams are exported from NetFlow-enabled devices at least once per second, or, as soon as a full UDP datagram of expired flows is available. NetFlow functionality is configured on a per-interface basis. To configure NetFlow Export capabilities, the user simply needs to specify the IP address and application port number of the Cisco NetFlow FlowCollector (a device that provides NetFlow Export data filtering and aggregation capabilities, to be discussed later) configured to receive the exported flow data.

Although configurable, the default main NetFlow cache size for various Cisco platforms is shown in the table below, along with the approximate amount of contiguous DRAM used by the NetFlow cache:

Platform	Default NetFlow Cache Size (entries)	Approximate amount of contiguous DRAM used by NetFlow cache
VIP with 128MB DRAM	128K	8MB
VIP with 64MB DRAM	64K	4MB
VIP with 32MB DRAM	32K	2MB
VIP with 16MB DRAM	2K	128KB
Cisco 7x00, uBR7246, RSP7000	64K	4MB
Cisco AS5800, 4x00, 3600, 2600, 2500, 1600, 1000	4K	256KB

Cisco IOS Router-Based NetFlow Aggregation

Customers can expect a large volume of export data from NetFlow when it is enabled on many interfaces on high-end routers that switch many flows per unit time (such as the Cisco 12000 and Cisco 7500 Series). Designed to significantly reduce NetFlow Export data volume and improve NetFlow scalability, router-based NetFlow aggregation is a Cisco IOS® software feature enhancement that enables limited router-based aggregation of NetFlow Export data. The five provided router-based NetFlow aggregation schemes enable the user to summarize NetFlow Export data on the router before the data is exported to a NetFlow data Collection device such as the Cisco NetFlow FlowCollector v2.0. With this feature enabled, aggregated NetFlow Export data is exported to a Collection device, resulting in lower bandwidth requirements for NetFlow Export data and reduced platform requirements for NetFlow data collection devices. In addition, this feature introduces NetFlow Export Version 8 (v8), a new Export datagram format designed to optimize NetFlow Export performance and bandwidth utilization.

The Router-based NetFlow Aggregation feature enables on-board aggregation by maintaining one or more extra NetFlow caches with different combinations of fields that determine which traditional flows are grouped together. These extra caches are called aggregation caches. As flows expire from the main flow cache, they are added to each enabled aggregation cache. The normal flow ager process runs on each active aggregation cache the same way it runs on the main cache. On demand aging is also supported.

Cisco IOS Router-Based Aggregation with NetFlow v8 is available on all Cisco router platforms that support NetFlow beginning in releases 12.0(3)T and 12.0(3)S. Five aggregation schemes will initially be supported (described later in this document).

The default size for each secondary NetFlow aggregation cache (exported via with v8 NetFlow Export datagrams) is 4096 entries on all platforms that support Cisco IOS NetFlow.

Use of Router-Based NetFlow Aggregation does not preclude the use of traditional NetFlow Services utilizing NetFlow Export v1/v5. Router-Based NetFlow Aggregation (utilizing v8 NetFlow Export datagrams) and traditional NetFlow Services (utilizing v1/v5 NetFlow Export datagrams) may be enabled simultaneously. Although these features can be used together, only the Router-based NetFlow Aggregation feature uses the v8 Export datagram format.

NetFlow Export Version Formats

The NetFlow Export datagram consists of a header and a sequence of flow records.

Figure 1 NetFlow Export Datagram Format

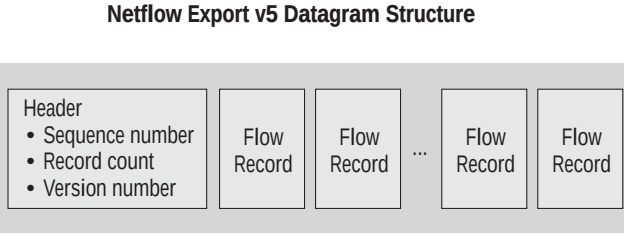

The Version 1 format was the original format supported in the initial Cisco IOS software releases containing NetFlow functionality. The Version 5 format is a later enhancement that adds Border Gateway Protocol (BGP) autonomous system information and flow sequence numbers. The Version 7 format is an enhancement that adds NetFlow support for Cisco Catalyst 5000 series switches equipped with a NetFlow feature card (NFFC). Version 7 is supported only by the Catalyst 5000 NFFC. Versions 2 through 4 and Version 6 were either not released or are not supported by

FlowCollector. Version 8 is the NetFlow Export format used when the Router-Based NetFlow Aggregation feature is enabled on Cisco IOS router platforms. Note that Cisco NetFlow FlowCollector v2.0 does not support collection of NetFlow v8 Export records. This functionality will be introduced in Cisco NetFlow FlowCollector v3.0 in the July 1999 timeframe.

In all four versions, the datagram consists of a header and one or more flow records. The first field of the header contains the version number of the export datagram. Typically, a receiving application that accepts any of the format versions allocates a buffer large enough for the largest possible datagram from any of the format versions and then uses the header to determine how to interpret the datagram. The second field in the header contains the number of records in the datagram (indicating the number of expired flows represented by this datagram) and is used to index through the records. Datagram headers for NetFlow Export versions 5, 7, and 8 also include a "sequence number" field used by NetFlow data consuming applications to check for lost datagrams.

See the following for the contents of the various NetFlow Export header and record formats:

NetFlow Export Version 1 Header Format

```
ushort version;         /* Current version=1*/
ushort count;           /* The number of records in PDU. */
ulong SysUptime;        /* Current time in msecs since router booted */
ulong unix_secs;        /* Current seconds since 0000 UTC 1970 */
ulong unix_nsecs;       /* Residual nanoseconds since 0000 UTC 1970 */
```

NetFlow Export Version 5 Header Format

```
ushort version;         /* Current version=5*/
ushort count;           /* The number of records in PDU. */
ulong SysUptime;        /* Current time in msecs since router booted */
ulong unix_secs;        /* Current seconds since 0000 UTC 1970 */
ulong unix_nsecs;       /* Residual nanoseconds since 0000 UTC 1970 */
ulong flow_sequence;    /* Sequence number of total flows seen */
uchar engine_type;      /* Type of flow switching engine (RP,VIP,etc.)*/
uchar engine_id;        /* Slot number of the flow switching engine */
```

NetFlow Export Version 7 Header Format

```
ushort version;         /* Current version=7*/
ushort count;           /* The number of records in PDU. */
ulong SysUptime;        /* Current time in msecs since router booted */
ulong unix_secs;        /* Current seconds since 0000 UTC 1970 */
ulong unix_nsecs;       /* Residual nanoseconds since 0000 UTC 1970 */
ulong flow_sequence;    /* Sequence number of total flows seen */
ulong reserved
```

NetFlow Export Version 8 Header Format

```
ushort version;         /* Current version */
ushort count;           /* The number of records in PDU. */
ulong SysUptime;        /* Current time in msecs since router booted */
ulong unix_secs;        /* Current seconds since 0000 UTC 1970 */
ulong unix_nsecs;       /* Residual nanoseconds since 0000 UTC 1970 */
ulong flow_sequence;    /* Seq counter of total flows seen */
uchar engine_type;      /* Type of flow switching engine */
uchar engine_id;        /* Slot number of the flow switching engine */
uchar aggregation;      /* Aggregation method being used */
uchar agg_version;      /* Version of the aggregation export=2 */
```

Table 1 NetFlow Flow Record Contents

Contents	V1	V5	V7
source IP address	Y	Y	y, zero in case of destination-only[1] flows
destination IP address	Y	Y	y
source TCP/UDP application port	Y	Y	y, zero in case of source-destination[2] flows or destination-only flows
destination TCP/UDP application port	Y	Y	y, zero in case of source-destination flows or destination-only flows
next hop router IP address	Y	Y	y, always zero
input physical interface index	Y	Y	y, always zero
output physical interface index	Y	Y	y
packet count for this flow	Y	Y	y
byte count for this flow	Y	Y	y
start of flow timestamp	Y	Y	y
end of flow timestamp	Y	Y	y
IP Protocol (for example, TCP=6; UDP=17)	Y	Y	y, zero in case of source-destination flows or destination-only flows
Type of Service (ToS) byte	Y	Y	y, switch sets it to ToS of first packet in flow
TCP Flags (cumulative OR of TCP flags)	Y	Y	y, always zero
source AS number		Y	y, always zero
destination AS number		Y	y, always zero
source subnet mask		Y	y, always zero
destination subnet mask		Y	y, always zero
flags (indicates, among other things, which flows are invalid)			y
shortcut router IP address[3]			y

[1] Catalyst 5000 systems with EARL2 (Encoded Address Resolution Logic) support unicast shortcuts to offload a router from routing between VLANs. EARL2 maintains a shortcut cache for flows being shortcut. In "destination-only" mode, EARL2 maintains only one shortcut cache entry per destination IP address. All flows to a destination use the same shortcut address.
[2] In "source-destination" mode, EARL2 maintains a shortcut cache entry per source-destination IP address pair. All flows between the source and destination use the same shortcut address regardless of the IP protocol ports.
[3] IP address of the router that is shortcutted by the Catalyst 5000 series switch

Please refer to the table below for v8 Flow Record contents:

Table 2 Router-Based Aggregation Schemes—NetFlow v8 Flow Record Contents

	AS	ProtocolPort	SourcePrefix	DestinationPrefix	Prefix
Source Prefix			Y		Y
Source Prefix Mask			Y		Y
Destination Prefix				Y	Y
Destination Prefix Mask				Y	Y
Source App Port		Y			
Destination App Port		Y			
Input Interface	Y		Y		Y
Output Interface	Y			Y	Y
IP Protocol		Y			
Source AS	Y		Y		Y
Destination AS	Y			Y	Y
First Timestamp	Y	Y	Y	Y	Y
Last Timestamp	Y	Y	Y	Y	Y
# of Flows	Y	Y	Y	Y	Y
# of Packets	Y	Y	Y	Y	Y
# of Bytes	Y	Y	Y	Y	Y

Each flow record within a version 1 NetFlow Export datagram has the following format:

```
ipaddrtype srcaddr;    /* Source IP Address */
ipaddrtype dstaddr;    /* Destination IP Address */
ipaddrtype nexthop;    /* Next hop router's IP Address */
ushort input;          /* Input interface index */
ushort output;         /* Output interface index */
ulong dPkts;           /* Packets sent in Duration (milliseconds between 1st & last packet in
                          this flow)*/
ulong dOctets;         /* Octets sent in Duration (milliseconds between 1st & last packet in
                          this flow)*/
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of the flow */
ushort srcport;        /* TCP/UDP source port number (.e.g, FTP, Telnet, etc.,or equivalent) */
ushort dstport;        /* TCP/UDP destination port number (.e.g, FTP, Telnet, etc.,or equivalent) */
ushort pad;            /* pad to word boundary */
uchar prot;            /* IP protocol, e.g., 6=TCP, 17=UDP, etc... */
uchar tos;             /* IP Type-of-Service */
uchar tcp_flags;       /* Cumulative OR of tcp flags */
uchar pad;             /* pad to word boundary */
ushort pad;            /* pad to word boundary */
uchar reserved[8]      /* reserved for future use*/
```

Each flow record within a version 5 NetFlow Export datagram has the following format:

```
ipaddrtype srcaddr;    /* Source IP Address */
ipaddrtype dstaddr;    /* Destination IP Address */
ipaddrtype nexthop;    /* Next hop router's IP Address */
ushort input;          /* Input interface index */
ushort output;         /* Output interface index */
ulong dPkts;           /* Packets sent in Duration (milliseconds between 1st & last packet in
                          this flow)*/
ulong dOctets;         /* Octets sent in Duration (milliseconds between 1st & last packet in
                          this flow)*/
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of the flow */
ushort srcport;        /* TCP/UDP source port number (.e.g, FTP, Telnet, etc.,or equivalent) */
ushort dstport;        /* TCP/UDP destination port number (.e.g, FTP, Telnet, etc.,or equivalent) */
uchar pad;             /* pad to word boundary */
uchar tcp_flags;       /* Cumulative OR of tcp flags */
uchar prot;            /* IP protocol, e.g., 6=TCP, 17=UDP, etc... */
uchar tos;             /* IP Type-of-Service */
ushort dst_as;         /* dst peer/origin Autonomous System */
ushort src_as;         /* source peer/origin Autonomous System */
uchar dst_mask;        /* destination route's mask bits */
uchar src_mask;        /* source route's mask bits */
ushort pad;            /* pad to word boundary */
```

Each flow record within a version 7 NetFlow Export datagram has the following format:

```
ipaddrtype  srcaddr;   /* Source IP Address */
ipaddrtype  dstaddr;   /* Destination IP Address */
ipaddrtype  nexthop;   /* Next hop router */
ushort input;          /* input interface index */
ushort output;         /* output interface index */
ulong dPkts;           /* Packets sent in Duration */
ulong dOctets;         /* Octets sent in Duration */
ulong First;           /* SysUptime at start of flow. */
ulong Last;            /* and of last packet of the flow. */
ushort srcport;        /* TCP/UDP source port number or equivalent */
ushort dstport;        /* TCP/UDP dest port number or equivalent */
uchar flags;           /* Shortcut mode(dest only,src only,full flows*/
uchar tcp_flags;       /* TCP flags */
uchar prot;            /* IP protocol, e.g., 6=TCP, 17=UDP, ... */
uchar tos;             /* IP Type-of-Service */
```

```
ulong src_as;          /* source AS# */
ulong dst_as;          /* destination AS# */
uchar src_mask;        /* source subnet mask */
uchar dst_mask;        /* destination subnet mask */
ushort pad;
ipaddrtype router_sc;  /* Router which is shortcut by switch */
```

The aggregated flow records within version 8 NetFlow Export datagrams have the following format:

For ASMatrix v8 aggregation scheme:

```
ulong flows;           /* Number of flows */
ulong dPkts;           /* Packets sent in Duration */
ulong dOctets;         /* Octets sent in Duration. */
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of flow */
ushort src_as;         /* originating AS of source address */
ushort dst_as;         /* originating AS of destination address */
ushort input;          /* Input interface index */
ushort output;         /* Output interface index */
```

For ProtocolPortMatrix v8 aggregation scheme:

```
ulong flows;           /* Number of flows */
ulong dPkts;           /* Packets sent in Duration */
ulong dOctets;         /* Octets sent in Duration. */
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of flow */
uchar prot;            /* IP protocol, e.g., 6=TCP, 17=UDP, ... */
uchar pad;
ushort reserved;
ushort srcport;        /* TCP/UDP source port number or equivalent */
ushort dstport;        /* TCP/UDP dest port number or equivalent */
```

For SourcePrefixMatrix v8 aggregation scheme:

```
ulong flows;           /* Number of flows */
ulong dPkts;           /* Packets sent in Duration */
ulong dOctets;         /* Octets sent in Duration. */
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of flow */
ipaddrtype src_prefix; /* Source prefix */
uchar src_mask;        /* source address prefix mask bits */
uchar pad;
ushort src_as;         /* originating AS of source address */
ushort input;          /* Input interface index */
```

For DestinationPrefixMatrix v8 aggregation scheme:

```
ulong flows;           /* Number of flows */
ulong dPkts;           /* Packets sent in Duration */
ulong dOctets;         /* Octets sent in Duration. */
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of flow */
ipaddrtype dst_prefix; /* Destination prefix */
uchar dst_mask;        /* destination address prefix mask bits */
uchar pad;
ushort dst_as;         /* originating AS of destination address */
ushort output;         /* Output interface index */
```

For PrefixMatrix v8 aggregation scheme:

```
ulong flows;           /* Number of flows */
ulong dPkts;           /* Packets sent in Duration */
ulong dOctets;         /* Octets sent in Duration. */
ulong First;           /* SysUptime at start of flow */
ulong Last;            /* and of last packet of flow */
ipaddrtype src_prefix; /* Source prefix */
ipaddrtype dst_prefix; /* Destination prefix */
uchar dst_mask; /* destination address prefix mask bits */
uchar src_mask; /* source address prefix mask bits */
ushort reserved;
ushort src_as; /* originating AS of source address */
ushort dst_as; /* originating AS of destination address */
ushort input; /* Input interface index */
ushort output; /* Output interface index */
```

NetFlow data export packets are sent to a user-specified destination, such as the workstation running FlowCollector, either when the number of recently expired flows reaches a predetermined maximum, or every second—whichever occurs first. Recall that, for a Version 1 datagram, up to 24 flows can be sent in a single UDP datagram of approximately 1200 bytes; for a Version 5 datagram, up to 30 flows can be sent in a single UDP datagram of approximately 1500 bytes, and for a Version 7 datagram, up to 28 flows can be sent in a single UDP datagram of approximately 1500 bytes.

For Version 8 datagrams, the following table lists the maximum number of aggregated flow records in and maximum size of each UDP datagram:

V8 Aggregation Scheme	Max number of aggregated export records per datagram	Maximum UDP Packet size (including v8 header) in bytes
ASMatrix	51	1456
ProtocolPortMatrix	51	1456
SourcePrefixMatrix	44	1436
DestinationPrefixMatrix	44	1436
PrefixMatrix	35	1428

Platform and Export Version Support

Table 3 NetFlow Version Matrix

Cisco IOS Software Release Version	Supported NetFlow Export Version(s)	Supported Cisco Hardware Platforms
11.1CA, 11.1CC	v1, v5	Cisco 7200, 7500, RSP7000
11.2, 11.2P	v1	Cisco 7200, 7500, RSP7000
11.2P	v1	Route Switch Module (RSM), 11.2(10)P and later
11.3, 11.3T	v1	Cisco 7200, 7500, RSP7000
12.0	v1, v5	Cisco 2600, 3600, 4500, 4700, AS5800, 7200, uBR7200, 7500, RSP7000, RSM
12.0T	v1, v5	Cisco 1000*,1600*,1720**, 2500*,2600, 3600, 4500, 4700, AS5800, 7200, uBR7200, 7500, RSP7000, RSM, MGX8800 RPM
12.0(3)T and later	v8	Cisco 1000*,1600*,1720**, 2500*,2600, 3600, 4500, 4700, AS5800, 7200, uBR7200, 7500, RSP7000, RSM, MGX8800 RPM
N/A	v7	Catalyst 5000 NetFlow Feature Card (NFFC)

*Support for NetFlow Export v1, v5, and v8 on Cisco 1000, 1600, and 2500 platforms is targeted for Cisco IOS software release 12.0(4)T. NetFlow support for these platforms will not be available in the Cisco IOS 12.0 mainline release.

**NetFlow support is available on Cisco 1720 beginning with Cisco IOS software release 12.0(3)T.

NetFlow Packaging:

- *Cisco 7200/7500/RSM*—Although NetFlow functionality is physically included in all software images for these platforms, customers must purchase a NetFlow Feature License in order to be licensed for its use. NetFlow licenses are sold on a per-node basis.
- *Cisco 1000/1600/2500/2600/3600/4000/AS5800 Series*—NetFlow functionality is supported only in Plus images for these platforms. Customers are required to purchase an appropriate Plus image in order to utilize NetFlow functionality on these platforms.

VIP Distributed NetFlow Export and CEF (Cisco Express Forwarding)

Cisco Express Forwarding (CEF) technology is a scalable, distributed, IP switching solution designed to meet the future performance requirements of the Internet and Enterprise networks. It represents the latest advance in Cisco IOS switching capabilities that includes NetFlow and Distributed Switching.

Existing layer 3 switching paradigms use a route-cache model to maintain a fast lookup table for destination network prefixes. The route-cache entries are traffic-driven in that the first packet to a new destination is routed via routing table information and as part of that forwarding operation, a route-cache entry for that destination is then added. This allows subsequent packets flows to that same destination network to be switched based on an efficient route-cache match. These entries are periodically aged out to keep the route cache current and can be immediately invalidated if the network topology changes. This "demand-caching" scheme—maintaining a very fast access subset of the routing topology information—is optimized for scenarios whereby the majority of traffic flows are associated with a subset of destinations. However, given that traffic profiles at the core of the Internet (and potentially within some large Enterprise networks) are no longer resembling this model, a new switching paradigm was required that would eliminate the increasing cache maintenance resulting from growing numbers of topologically dispersed destinations and dynamic network changes.

CEF avoids the potential overhead of continuous cache churn by instead using a Forwarding Information Base (FIB) for the destination switching decision that mirrors the entire contents of the IP routing table; i.e. there is a one-to-one correspondence between FIB table entries and routing table prefixes; therefore there is no need to maintain a route-cache. This offers significant benefits in terms of performance, scalability, network resilience and functionality, particularly in large complex networks with dynamic traffic patterns.

Cisco's Express Forwarding technology is optimized for information distribution, allowing it to take advantage of the distributed architecture of the high end Cisco IOS routers such as the Cisco 7500. Thus Distributed CEF (dCEF) delivers scalable switching capacity by providing each of the Cisco 7500 Versatile Interface Processors (VIPs) with an identical on-card copy of the FIB database enabling them to autonomously perform Express Forwarding and therefore significantly increase aggregate throughput. CEF also uses Adjacency Tables to hold the layer 2 next hop addresses for all FIB entries so that the associated prepend can be added locally, thus minimizing latency, before switching between linecards. Therefore in the case of dCEF on the Cisco 7500 platform, the Route Switch Processor (RSP) is relieved of any switching operation and so has significantly more CPU power available to perform routing functions, management, network services, etc.

NetFlow leverages the distributed CEF switching infrastructure on Cisco 7500 class routers to provide high performance distributed NetFlow-enabled switching on VIPs. Distributed CEF with NetFlow localizes switching and NetFlow processing to each VIP. Each VIP exports its own set of NetFlow Export datagrams without requiring RSP resources, providing a highly scalable method for gathering traffic accounting information.

dCEF VIP NetFlow Export functionality was introduced in Cisco IOS software release 11.1(19)CC, and is also available in the 12.0 and 12.0T release trains. Note that version 1 (v1.0) of the NetFlow FlowCollector application (to be discussed later in this paper) does not support dCEF VIP NetFlow Export collection. FlowCollector v2.0 is required, as it introduces enhancements to properly process the modified export header identifying each export source.

Supported Interfaces, Encapsulations and Protocols

NetFlow supports IPv4 (and IPv4-encapsulated) routed traffic over a wide range of interface types and encapsulations including Ethernet, Fast Ethernet, FDDI, HSSI, POSIP, CT3, GRE-tunnels, and serial. NetFlow support for fast-switched Token Ring and ISDN interfaces will be introduced in Cisco IOS software release 12.0. NetFlow supports only IP packets. IPX is not supported.

12.0/12.0T–based NetFlow supports accounting for ISL subinterfaces, but does not report the individual sub-interfaces in the flow record. Instead, it reports the main interface index in the flow record. NetFlow Export support for Frame Relay subinterface traffic accounting was introduced in Cisco IOS software 12.0(1) and 12.0(1)T. NetFlow Export support for ATM subinterface traffic accounting was introduced in Cisco IOS software 12.0(1)T. Note that in both cases, although NetFlow Export data will properly account for FR/ATM traffic on a per-subinterface basis, (and properly report the subinterface SNMP indexes in the flow records) NetFlow services are configurable only on a per-physical interface basis (e.g. NetFlow cannot be configured on a per-subinterface basis).

At this time, NetFlow Export data does not capture traffic accounting statistics on IP multicast datagrams. This capability is planned for Cisco IOS software release 12.0(7)T.

NetFlow Activation and Data Collection Strategy

Cisco recommends careful planning of NetFlow deployment with NetFlow services activated on strategically located edge/aggregation routers for service providers and WAN access routers for Enterprise customers which capture the data required for planning, monitoring and accounting applications. Key deployment considerations include the following:

- NetFlow services should be utilized as an edge metering and access list performance acceleration tool and not activated on "hot" core/backbone routers or routers running at very high CPU utilization rates
- Understanding your application-driven data collection requirements: accounting applications may only require originating and terminating router flow information whereas monitoring applications may require a more comprehensive (data intensive) end-to-end view
- Understanding the impact of network topology and routing policy on flow collection strategy: for example, avoid collecting duplicate flows by activating NetFlow on key aggregation routers where traffic originates or terminates and not on backbone routers or intermediate routers which would provide duplicate views of the same flow information
- Service providers in the "transit carrier" business (i.e. carrying traffic neither originating nor terminating on their network) may utilize NetFlow Export data for measuring transit traffic usage of network resources for accounting and billing purposes

NetFlow is an input side measurement technology which should be deployed on appropriate interfaces on edge/aggregation or WAN access routers to gain a comprehensive view of originating and terminating traffic to meet customer needs for accounting, monitoring or network planning data. The key mechanism for enhancing NetFlow data volume manageability is careful planning of NetFlow deployment. NetFlow can be deployed incrementally (i.e. interface by interface) and strategically (i.e. on well chosen routers) —instead of widespread deployment of NetFlow on every router in the network. Cisco will work with customers to determine key routers and key interfaces where NetFlow should be activated based on the customer's traffic flow patterns and network topology and architecture.

Access Control List Acceleration

Access control on Cisco routers is provided via access control lists (ACLs), which enable packet filtering applications to be based on source and destination addresses, protocols, and specific interfaces. With traditional switching mechanisms, each individual packet is matched against a set of access lists to determine if a configured packet filter applies for a particular source and destination address pair.

With NetFlow enabled, only the first packet of a flow follows this process. If the first packet in a flow passes through these filters, an entry is added to the NetFlow flow cache. Subsequent packets in the same flow are then switched based on this cache entry, without needing to be matched against the complete set of access lists. This significant simplification enables NetFlow to maintain high performance when access lists are used for packet filtering. Specific performance will vary based on the number and complexity of the access lists.

In the future, the NetFlow flow cache will be used to accelerate a number of Cisco IOS services, including Cisco IOS Network Address Translation (NAT). CEF-based NetFlow-accelerated Policy-based Routing (NPR) is available on all platforms that support NetFlow beginning with Cisco IOS software release 12.0(3)T. Additional information will be provided as such service acceleration becomes available on a feature-by-feature basis.

NetFlow and Quality of Service (QoS)

NetFlow provides the measurement base for Cisco's Internet and Enterprise QoS initiatives. NetFlow records the Type of Service (ToS) field in the IP header as well as application ports, traffic volumes and timestamps. Thus, service providers can charge premium charges for premium packets as well as charging based on usage, time and application. NetFlow does not accelerate QoS services such as traffic classification or Committed Access Rate (CAR). Note that NetFlow itself is not a QoS feature; it simply captures the Quality of Service level of each flow.

NetFlow Management Applications

Cisco IOS NetFlow software is part of a larger family of products, management utilities and partner applications designed to gather and export flow statistics, collect and perform data volume reduction on the exported statistics, and feed flow detail records to consumer applications such as planning, accounting and monitoring.

Figure 2 NetFlow Infrastructure

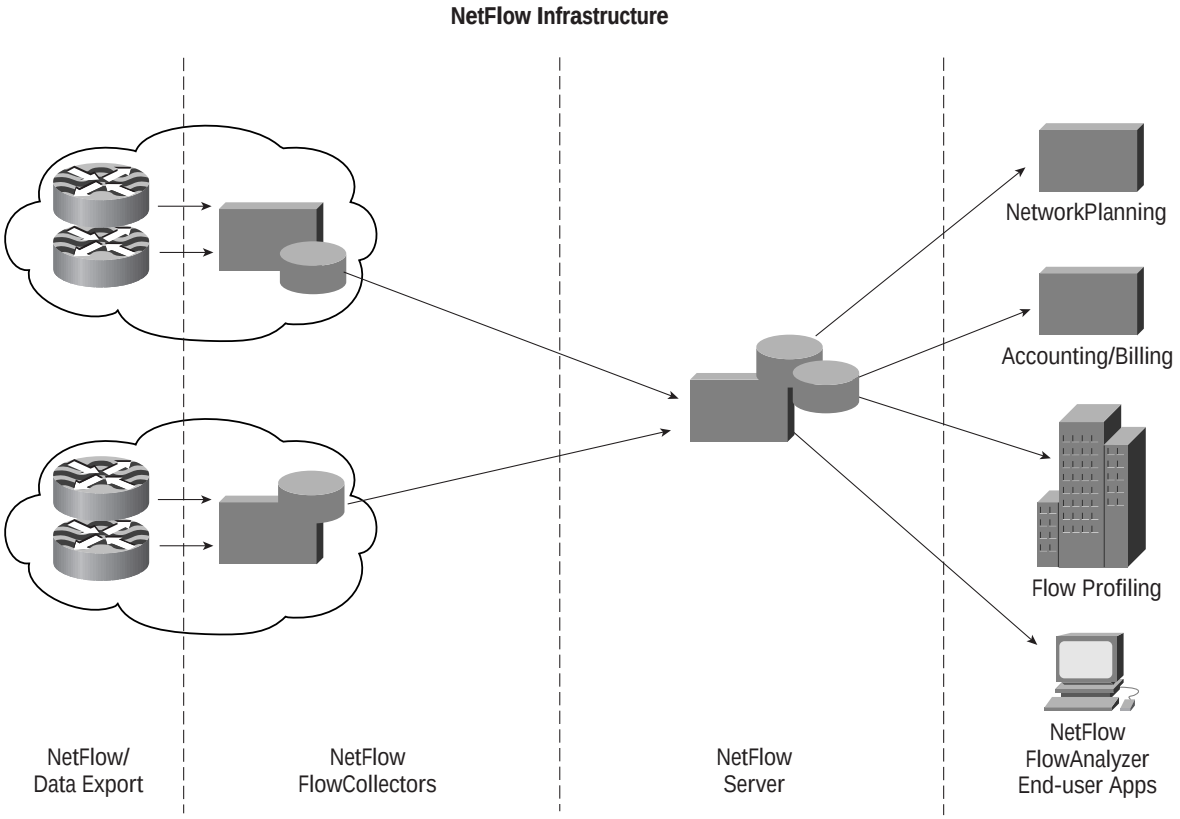

Cisco provides a family of NetFlow management applications which:

- Collect, store and perform data volume reduction on exported NetFlow data
- Provide a scalable and distributed NetFlow data collection and consolidation architecture
- Provide network monitoring, analysis, and troubleshooting tools

NetFlow FlowCollector

The NetFlow FlowCollector provides fast, scalable, and economical data collection from multiple NetFlow Export-enabled devices. A UNIX application supported on Solaris and HP-UX platforms, the FlowCollector:

- Consumes flow datagrams from multiple NetFlow Export-enabled devices
- Performs data volume reduction through selective filtering and aggregation
- Stores flow information in flat files on disc for post-processing by NetFlow data consumers, including third party billing applications, traffic analysis tools, etc.

Today the FlowCollector is a key provider of time-based, granular data measurements to external applications. Service Provider and Enterprise customers can utilize the FlowCollector as an integral component of their distributed data collection processes. Traffic accounting details gathered by FlowCollectors will drive extensive new capabilities in the areas of network planning, accounting/billing and network and application resource monitoring, with minimal impact on router performance and without extensive (and expensive) polling operations. A future proposed NetFlow Server product will provide the capability for external applications to access NetFlow data via SQL from a centralized relational database (discussed later in this paper).

The NetFlow FlowCollector does not provide flow correlation and or flow de-duplication capabilities. For performance and scalability reasons, these functions must be performed downstream in a post-collection application such as the NetFlow Server or third-party billing/accounting application.

Figure 3 NetFlow FlowCollector

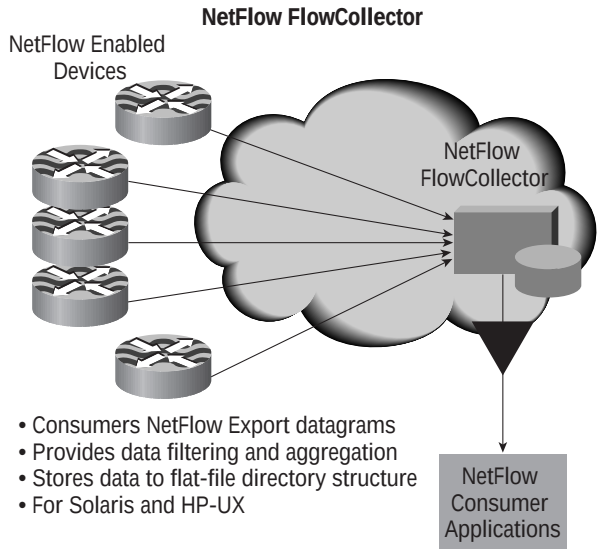

The FlowCollector passively listens to specified UDP ports to receive and process exported NetFlow datagrams. The FlowCollector application provides a high performance, easy-to-use solution that scales to accommodate consumption of NetFlow Export data from multiple devices in order to support key flow-consumer applications including accounting, billing, and network planning/monitoring.

In high traffic service provider and enterprise environments the number of flows and the corresponding flow record data volume may be considerable. The FlowCollector provides highly flexible configuration options for specifying data volume reduction methods to fit the customer's application requirements. This data volume reduction is made possible through the use of *filters* and *aggregation schemes*.

Filters

FlowCollector filters enable the user to specify which flow records should be either accepted for or rejected from further processing and storage. The FlowCollector provides filtering on the flow record data fields specified in the table below:

Table 4 NetFlow FlowCollector Filters

Filter	NetFlow FlowCollector v1.0	NetFlow FlowCollector v2.0
Source IP address	Y	Y
Destination IP address	Y	Y
Next hop router IP address	Y	Y
Source TCP/UDP application port	Y	Y
Destination TCP/UDP application port	Y	Y
Input physical interface index	Y	Y
Output physical interface index	Y	Y
Source AS number		Y
Destination AS number		Y
Type of Service (ToS) byte		Y
Export datagram source IP address		Y

Note that FlowCollector provides several new filtering options, including the capability to filter NetFlow Export datagrams based on the datagram source IP address. This enables the user to configure the FlowCollector to selectively accept datagrams only from specific data sources. This effectively enables the user to disable NetFlow data export on selected device(s), as seen by the FlowCollector, without having to actually reconfigure a NetFlow Export-enabled device(s).

Note: Cisco NetFlow FlowCollector v1.0 reached its "End of Sales" lifecycle milestone when FlowCollector v2.0 was released. Thus, the v1.0 product, though still supported, is no longer available.

Cisco NetFlow FlowCollector v2.0 does not support collection of NetFlow v8 Export records. This functionality will be introduced in Cisco NetFlow FlowCollector v3.0 in the July 1999 timeframe.

Aggregation Schemes

FlowCollector aggregation enables the user to reduce export data volume by storing aggregated summary records instead of raw flow records to flat files. The FlowCollector enables the following aggregations:

- *Source node*—one row is stored for every source IP address. The row consists of source address, total packets sent, total bytes sent and total flow records
- *Destination node*—one row is stored for each destination IP address. The row consists of destination address, total packets received, total bytes received and total flow records
- *Detail destination node*—one row is stored for each destination IP address. The row consists of destination address, source TCP/UDP port, destination TCP/UDP port, protocol, received packet and byte count and flow record total
- *Host Matrix*—one row is stored for each source address-destination address pair. The row consists of source address, destination address, exchanged packet count total and flow record total
- *Detail Host Matrix*—one row is stored for each source-destination address pair. The row consists of source address, destination address, source port, destination port, protocol, exchanged packet and byte count total, flow record total and timestamps of the first and last flows aggregated into this row
- *Source TCP/UDP port*—one row is stored for each transport layer source port. The row consists of source port number, packet and byte transmit total and flow record total
- *Destination TCP/UDP port*—one row is stored for each transport layer destination port. The row consists of destination port number, total received packets and bytes and flow record total
- *Protocol*—one row is stored for each protocol. The row consists of protocol name, byte and packet total sent to or received from this protocol and flow record total
- *Detail Interface*—one row is stored for each input-output physical interface pair. The row consists of source address, destination address, input interface, output interface, next hop address, packet and byte exchanged total and flow record total

- *Autonomous System Matrix*—one row is stored for each autonomous system pair. The row consists of source autonomous system number, destination autonomous system number, packet and byte exchanged total and flow record total
- *Call Record*—provides an aggregation scheme for usage-based billing/accounting applications. Each row consists of IP source address, IP destination address, source port number, destination port number, protocol, type-of-service, packet total, byte total, number of flows summarized into record, earliest timestamp and latest timestamp

New aggregation schemes in NetFlow FlowCollector v2.0:

- *Detail Source Node*—Each row consists of source IP address, source TCP/UDP application port, destination TCP/UDP application port, IP protocol, total packets sent, total bytes sent and total flow records
- *Detail AS Matrix*—Each row consists IP source address, IP destination address, source port number, destination port number, protocol, type-of-service, input interface, output interface, source AS, destination AS, packet total, byte total, and number of flows
- *NetMatrix*—Each row consists of masked source IP address, masked destination IP address, source address mask, destination address mask, input interface, output interface, packet total, byte total, and number of flows

The various aggregation schemes and output are summarized below in Figures 4 and 5:

Figure 4 FlowCollector Aggregation Schemes

Schemes	src addr	dest addr	src port	dest port	prot	ToS	input IF	output IF	next-hop	src AS	dest AS	masked src addr	masked dest addr	src mask	dest mask
v1.0															
Source note	—														
DestNode		—													
HostMatrix	—	—													
SourcePort			—												
DestPort				—											
Protocol					—										
DetailDest-Node		—	—	—	—										
DetailHost-Matrix	—	—	—	—	—										
Detail-Interface	—	—				—	—	—							
CallRecord	—	—	—	—	—	—									
ASMatrix										—	—				
v2.0															
DetailSource-Node	—		—	—	—										
DetailAS-Matrix	—	—	—	—	—		—	—		—	—				
NetMatrix												—	—	—	—

Figure 5 FlowCollector Aggregation Output

	packet count	byte count	flow count	active time	first time-stamp	last time-stamp
v1.0						
Source note	—	—	—			
DestNode	—	—	—			
HostMatrix	—	—	—			
SourcePort	—	—	—			
DestPort	—	—	—			
Protocol	—	—	—			
DetailDest-Node	—	—	—			
DetilHost-Matrix	—	—	—	—	—	—
Detail-Interface	—	—	—			
CallRecord	—	—	—	—	—	—
ASMatrix	—	—	—			
v2.0						
DetailSource-Node	—	—	—			
DetailAS-Matrix	—	—	—			
NetMatrix	—	—	—			

Threads

The FlowCollector also provides a powerful construct for specifying data reduction and storage policies for each collection port called *threads*. Thread policy definitions include applicable filters and aggregators as well as storage directory path, collection time periods and storage directory cleanup specifications. Multiple threads may defined for each collection port to accommodate the data collection requirements of different applications (e.g. network planning vs. accounting/billing).

The *period* parameter is one attribute defined in a thread: it specifies the time period over which data is aggregated for a specified output file (e.g. a period of 30 generates two output files per hour). Thus, a small period parameter results in smaller aggregation files being flushed to hard disk frequently. A large period parameter results in larger aggregation files being flushed to hard disk less frequently. Because export data aggregation is performed in memory prior to being flushed to disk, use of a small period parameter results in lower platform RAM requirements, while increasing disk storage requirements. Conversely, use of a large period parameter results in large platform RAM requirements and comparatively less disk storage capabilities. In any case, FlowCollector memory and disk storage requirements depend on many variables, including:

- Number of NetFlow Export-enabled devices sending Export data to the FlowCollector
- Interface utilization on NetFlow-enabled devices
- Export datagram flow rate to FlowCollector
- Period parameter

Data Storage

The FlowCollector provides a powerful set of storage management capabilities to quickly and reliably:

- Record NetFlow Export data to disk in a well defined flat-file directory structure
- Make the data files available in a format for efficient extraction and post-processing by external applications
- Provide file cleanup and reclamation utilities to minimize overall storage space requirements

The FlowCollector provides a hierarchical directory structure organization and structured file format for thread data output to allow for easy consumption of aggregated data by Cisco or third party management and reporting applications. FlowCollector output files contain a file header specifying parameters such as router name, aggregation scheme used to create the file, and file period parameter (as previously defined) to enable file parsing by consuming applications.

NetFlow FlowCollector User Interface (NFUI)

FlowCollector includes the NetFlow FlowCollector user interface (NFUI), which is an interactive, menu-oriented user interface used to perform the following tasks:

- Display runtime configuration parameters, resource definitions, and statistics
- Modify existing configuration parameters
- Define new configuration parameters

The user interface consists of a main menu that provides access to a series of submenus and information displays. The user interface is self-guiding: it displays some information and then prompts the user to act on that information. The user interface also contains embedded help menus to assist the user in navigating through the user interface and understanding menu operations. The help menus explain all the options available for retrieving, configuring, and reviewing the FlowCollector runtime configuration parameters and statistics.

New Features in FlowCollector v2.0

FlowCollector v2.0 provides the following new functionality:

- *Increased performance*—FlowCollector v2.0 provides approximately five times the performance of its v1.0 predecessor.
- *Three additional Aggregation schemes*—
  - DetailSourceNode
  - DetailASMatrix
  - NetMatrix
- *Additional data filtering options*—FlowCollector v2.0 additionally enables the user to filter NetFlow Export data on the following fields:
  - Source Autonomous System (AS) number (origin or peer)
  - Destination Autonomous System (AS) number (origin or peer)
  - IP Type of Service (ToS) byte
  - Export data source IP address
- *Support for v7 NetFlow Export data* from the Catalyst 5000 NetFlow Feature Card (NFFC).
- *"Router Grouping" feature* enables export data from multiple NetFlow-enabled devices to be aggregated under a single logical name/address.
- *"File Push Hook"*—Enables the FlowCollector to execute a user-supplied script after it has written a new data file. This enables the FlowCollector to inform a client application that new FlowCollector data is available and should be consumed. As a result, the client application need not periodically poll the FlowCollector to determine when new data is available.
- *"FilesReady" File*—A FlowCollector file containing the names and paths of all datafiles recently written to disk. A client application would typically download this file from the FlowCollector periodically to determine the name(s) of new datafiles created since the last fetch, and then issue a fetch for these files. May be used in conjunction with the "File Push Hook."
- *Export Record Version Autodetect*—FlowCollector v2.0 automatically detects the version numbers in received NetFlow Export datagrams. This eliminates the need for additional configuration to specify the version type(s) expected on a given UDP port.
- *"show-tech" utility for debugging*—Provides users an easy to generate the debug information necessary for support and troubleshooting purposes.

NetFlow FlowCollector v2.0 Minimum Platform Requirements

NetFlow FlowCollector v2.0 is available for the following platforms:

- *Solaris—Version 2.5.1 or 2.6*. Recommended to run on a minimum of an ULTRA-1 workstation with at least:
  - 128MB RAM
  - 512MB swap space
  - 4GB disk space
- *HP-UX—Version 11.0*. Recommended to run on a Class C workstation (or higher) with at least:
  - 128MB RAM
  - 512MB swap space
  - 4GB disk space

Note: To prevent NetFlow data export packet loss, the FlowCollector workstation should be installed local to the NetFlow data export device and should not be running other applications.

Year 2000 Compliance

Both NetFlow FlowCollector v2.0 and FlowAnalyzer v2.0 have passed testing for Year 2000 Compliance and are deemed Year 2000 compliant. For more information, please refer to the "Cisco Year 2000 Product Compliance" page at `http://www.cisco.com/warp/customer/752/2000/cptbl_ov.htm`.

NetFlow FlowAnalyzer

The FlowAnalyzer application is a network traffic analysis tool that combines a graphical user interface with other companion modules. Together, these modules enable the user to retrieve, display, and analyze NetFlow data that has been collected from NetFlow FlowCollector flat files.

The FlowAnalyzer provides several major categories of functionality including:

- NetFlow Export data visualization policies (e.g. what to visualize and how to visualize it)
- Graphical data display based on the specified visualization policies
- Data export to external applications (e.g. Excel spreadsheets) for reporting purposes

The individual modules of the FlowAnalyzer application are described in the following sections.

Figure 6 Elements of the NetFlow FlowAnalyzer

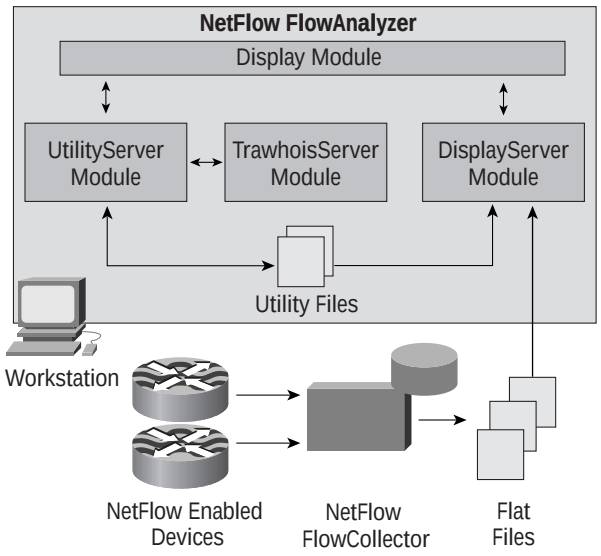

Display Module

The Display module in FlowAnalyzer v2.0 is a stand-alone Java application that can be installed and run on workstations or PCs in your network. This module provides an easy-to-use, graphically oriented user interface to the NetFlow system. The user may install any number of Display modules into your network, as appropriate for the user's particular NetFlow data retrieval and analysis requirements.

The Display module can present NetFlow data in a variety of formats, including the following:

- Tables
- Graphs—Pie charts, bar charts, or histogram charts
- Numeric lists
- Non-numeric lists

Figure 7 Sample FlowAnalyzer Output

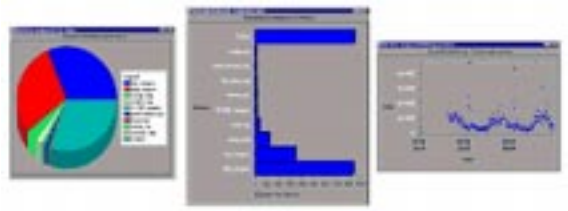

In addition, the Display module incorporates a number of facilities that help the user to use the FlowAnalyzer application conveniently and effectively. Such user aids include the following:

- *Pop-up windows*—Several pop-up windows have been incorporated into the Display module user interface to provide helpful tips and reminders in using the FlowAnalyzer.
- *Movable time slider marks*—Two vertical bars in the upper right area of the Display module window enable the user to establish the applicable time interval for FlowAnalyzer data retrieval and analysis tasks. The effect of shortening the applicable time interval is to limit the breadth and scope of FlowAnalyzer operations, which includes reducing the volume of NetFlow data to be processed in responding to user commands. Conversely, lengthening the applicable time interval for FlowAnalyzer tasks has opposite effects, such as increasing the volume of NetFlow data to be processed, placing heavier demands on system resources, and reducing system performance. The user can position these time slider marks through a wide range to suit the needs of any given FlowAnalyzer task at hand.
- *Get TopN pull-down menu*—Enables the user to select the number of traffic flows to be used in processing NetFlow data for display purposes. For example, the user can select the top "N" packets, bytes, or flows for display purposes, where "N" can be any the user of the following values: 10, 100, 500, 1000, 2000, 5000, or 10000. The default value for "N" is 100.
- *"Sorted by" pull-down menu*—Enables the user to specify a sort key by which NetFlow data is to be sorted during data retrieval operations. This pull-down menu enables the user to "filter" the NetFlow data as octets, packets, or flows for display purposes.
- *Status bar*—Provides an "at a glance" status view of any in-process or completed FlowAnalyzer task.

DisplayServer Module

The DisplayServer module runs on a host workstation in your network and receives and acts on requests for Netflow data that are issued at the console of a Display module. The user can configure any number of DisplayServer modules in the network, as appropriate, to act on user requests for NetFlow data.

The DisplayServer module responds to such user requests by accessing the NetFlow data files stored on a designated FlowCollector workstation in the network and transmitting the requested data to the Display module for presentation on the screen of a host workstation or PC. The Display module formats the data on the screen according to the selected aggregation scheme.

UtilityServer Module

The UtilityServer module provides host, autonomous system (AS), and application port translations, as well as device interface information to the FlowAnalyzer.

TrawhoisServer Module

The TrawhoisServer module receives requests from the UtilityServer module to perform AS translation tasks. In response, the TrawhoisServer module returns AS names (as contained in the router arbiter database) to the UtilityServer module. “Trawhois” is an acronym for trivial routing arbiter whois. The Trawhois Server was developed by Merit Network (http://www.merit.edu).

Note: FlowAnalyzer v1.0 reached its “End of Sales” lifecycle milestone when FlowAnalyzer v2.0 was released. Thus, the v1.0 product, though still supported, is no longer available.

New Features in FlowAnalyzer v2.0

- Search operations:
  - “AS data summaries” for providing output for flows with specified source and destination AS numbers
  - “AS/Protocol summaries” for providing a protocol summary of NetFlow data in addition to selecting only flows with specified source and destination AS numbers
  - “IP data summaries” for providing output for flows with specified source and destination IP addresses to the significance specified in source and destination IP address masks
  - “IP/Protocol summaries” for providing a protocol summary of NetFlow data in addition to selecting only flows with specified source and destination IP addresses to the significance specified in source and destination IP address masks
  - “NetFlow Drilldown” enables the user to view AS, IP address, and protocol data in a hierarchical manner at various levels of detail
- Support for three new aggregation schemes introduced in FlowCollector v2.0:
  - DetailASMatrix
  - DetailSourceNode
  - NetMatrix
- DetailASMatrix aggregation drilldown
- Multiple dataset, multiple router or multiple interface selection (e.g., by AS or POP)
- Simpler installation and start up
- A Web browser & server are no longer needed or used
- Extensive online help
- AS, host and port number to name translation
- Several performance improvements, including load sharing of multiple NFA DisplayServers

Supported and Unsupported Configurations

NetFlow FlowAnalyzer Version 2.0 has been expressly designed to operate with NetFlow FlowCollector Version 2.0. However, for the benefit of users of previous versions of these applications, it is important to note the following:

Supported configurations:

- FlowAnalyzer Version 2.0 processing data collected by FlowCollector Version 2.0
- FlowAnalyzer Version 2.0 processing data collected by FlowCollector Version 1.0.

Unsupported configurations:

- FlowAnalyzer Version 1.0 processing data collected by FlowCollector Version 2.0. FlowAnalyzer Version 1.0 can only process data collected by FlowCollector Version 1.0.

Additional information:

- FlowCollector v2.0 is required if you intend to process data for FlowCollector v2.0’s DetailASMatrix aggregation scheme.

NetFlow FlowAnalyzer v2.0 Minimum Hardware Requirements

The UNIX and PC platform requirements for the FlowAnalyzer application are listed below:

- Solaris—Version 2.5.1 or 2.6 running on an ULTRA-1 workstation, or
- HP-UX—Version 10.2 running on a Class C workstation (or higher)
- Windows NT 4.0—For the FlowAnalyzer Display module only
- The Bourne shell "sh" (/bin/shell) must be available for execution

The workstation on which you run FlowAnalyzer v2.0 must meet the following minimum requirements:

- 256 MB of physical memory (RAM) and 400 MB of free logical memory (this is a requirement only for the host on which the DisplayServer module runs).
- 70 MB of free disk space for the tar and uncompressed installation files (these files can be deleted after installation).
- 50 MB of free disk space for the installed executable (25 MB is required for the NFADisplay executables running on the PC).

The PC on which you run the Display module must meet the following requirements:

- Windows 95 or Windows NT 4.0
- Pentium machine containing at least a 166 MHz processor and 64 MB or more of RAM.

Note: To eliminate potential data loss and to prevent workstation performance degradation during NetFlow data collection and processing, it is recommended (but not mandatory) that you install the NetFlow FlowCollector and the NetFlow FlowAnalyzer applications on different workstations.

NetFlow Server

Customers can expect a large volume of export data from NetFlow due to the size and distributed nature of service provider and large Enterprise customer networks, the granularity of recorded flow data and the rapid traffic growth in networks where NetFlow is deployed. In order to deal with these issues, the NetFlow management utilities will include a distributed data collection paradigm based on the NetFlow Servers collecting and post-processing data from multiple, distributed NetFlow FlowCollectors.

The NetFlow Server will enable the following capabilities:

- The ability to consolidate NetFlow statistics across one or more NetFlow FlowCollectors in order to provide a single repository of network-wide statistics
- The ability to further summarize NetFlow statistics by enabling bi-directional consolidation, as well as creating daily, monthly, quarterly and yearly summaries
- The ability to store NetFlow statistics in a common commercially accepted RDBMS, in order to leverage SQL and other database tools for complex queries and reporting, as well as to provide to end-user applications easy access to NetFlow data
- The ability to encrypt and compress NetFlow statistics when they are transmitted over WAN links

The NetFlow Server collects user-specified datafiles from one or more NetFlow FlowCollectors in the network, optionally stores them in the database, and optionally performs time-based consolidation of the data. The NetFlow makes the data available in a central database, thereby providing a single point of application and end-user access to NetFlow statistics.

Note: In order to better meet customer needs, Cisco has decided to partner with third-party vendors to deliver the NetFlow Server product. We are now targeting a late CY 1999/early CY 2000 delivery date. More information will be disseminated as it becomes available.

Figure 8 NetFlow Server

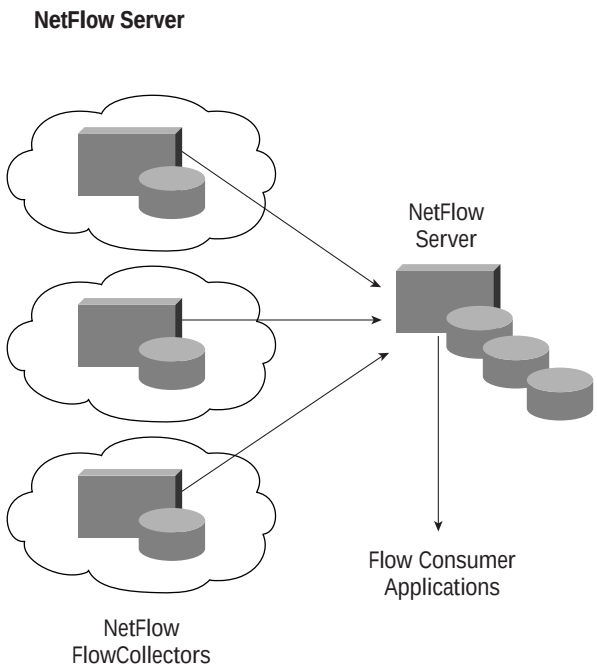

- Central repository for NetFlow data
- Post-collection aggregation and time-based consolidation
- Reliable data collection
- Secure data collection
- Oracle8 RDBMS
- Single point of access for end-user applications

In small to medium size networks, there will likely be a single NetFlow Server that collects the data from all the NetFlow FlowCollectors in the network. In large networks, users can, if required, install two or more NetFlow Servers in their network. When present, multiple NetFlow Servers operate independently of each other, without any coordination among them, and have exclusive "collection rights" to any FlowCollectors from which they are configured to collect from datafiles.

The NetFlow Server application will initially be available for Solaris hosts, specifically Sun Ultra5s or better running Solaris 2.6 and higher. A HP-UX version may be available later. NetFlow Server is targeted to work with the Oracle 8 database product.

NetFlow Server Data Collection

The NetFlow Server collects the datafiles from of one or more NetFlow FlowCollectors in the network. The NetFlow Server can collect datafiles from FlowCollectors on a regularly scheduled basis, on a "continuous" basis, or on an "on-demand" basis. The NetFlow Server allows the user to define a wide variety of data collection parameters, such as which datafiles to be collected, as well as the from which FlowCollectors from which to collect data (based on the concept of a "Collection Profile"), how often, and when to collect datafiles (based on a "Collection Schedule"). The user can also specify whether the collected datafiles are to be stored in the database, or only used to produce consolidated (summarized) data (based on "Consolidation Profile"). Although the NetFlow Server can collect datafiles from many FlowCollectors simultaneously, it will not collect more than one datafile at any one time from the same FlowCollector (to prevent adverse performance impact on the NetFlow FlowCollector). This behavior minimizes disk and CPU usage on each FlowCollector so that its data collection tasks—consuming NetFlow Data Export—are not impacted.

The NetFlow Server uses TCP and application-level protocols to ensure the integrity of the data retrieved from the FlowCollector (e.g. no loss of data and no duplication of data into the database).

In addition, application-level compression may be configured by the user to enable data compression during the data collection process. When transferring datafiles across the network, connection authentication, as well as an optional data encryption/compression mechanism , may also be used.

NetFlow Server Collection Profiles

Collection Profiles and Collection Schedules enable the user to specify the following on the NetFlow Server:

- Which datafiles to collect
- From which FlowCollectors to collect datafiles, and
- When to perform the collection

The Collection Profile specifies the data to be collected, and enables the user to define powerful and flexible criteria for specifying datafiles without explicitly naming them. A Collection Profile allows for the creation of a "filter" that contains a list of Threads, Aggregation Schemes and Source (router) names to match upon.

Collection Profiles enable the administrator to specify which datafiles are to be collected from a given FlowCollector, and whether compression and/or encryption are to be used during file transfers. More than one Collection Profile per FlowCollector is allowed. Collection Profiles also specify whether encryption/compression are to be used for the file transfers. A Collection Profile can specify that all datafiles pertaining to one or more Aggregation Schemes, Threads, or routers, be collected.

The components that make up a Collection Profile are:

- *FlowCollector ID*—The host name or IP address of the FlowCollector to which this Collection Profile pertains
- *Thread ID*—The name(s) of the Threads (as defined on the FlowCollector) for which datafiles are to be collected
- *Aggregation Scheme Name*—The FlowCollector Aggregation Scheme(s) for which datafiles are to be collected. This parameter is used in lieu of the Thread ID parameter when the same Aggregation Scheme may be used across Thread IDs. Further, using Aggregation Scheme names will isolate the NetFlow Server configuration from changes to FlowCollector configurations
- *Router ID*—If specified, it is the IDs of the routers/switches for which datafiles are to be collected. The user can specify a group name as well (as used in the router-groupname feature on the FlowCollector). If not specified, all the routers with data pertaining to the specified Thread IDs or Aggregation Schemes are collected
- *Encryption*—Specifies if the datafiles collected using this profile are to be encrypted en-route from the FlowCollector to the NetFlow Server.
- *Compression*—Specifies if the datafiles collected using this profile are to be compressed before sending them over the network.

The user may opt for either compression (less network traffic), or encryption (better security) or none (less CPU usage on the FlowCollector workstation), or both.

NetFlow Server Scheduled Collections

The "When" of the collection is specified by NetFlow Server "Collection Schedules". They enable the administrator to configure when the NetFlow Server will collect for a given Collection Profile by specifying both time of day and the day (day-of-month, day-of-week). In addition, the user can also optionally specify when a collection "stop-time" must end ("bounded collection"), perform scheduled collections, including a "collection window" on any day-of-the-week(s) or date(s).

This enables the administrator to configure when a Collection Profile is to be used for collecting data. The scheduling scheme enables the administrator great flexibility in determining when and for how long data transfer should be done—thereby allowing for deterministic network usage by NetFlow Server data transfers.

The administrator can realize the following benefits from using Collection Schedules:

- By scheduling data collection at off-peak hours, the impact on the network (amount of data being transferred in-band), the FlowCollector (both in terms of CPU and disk I/O) and the NetFlow Server (many database updates, CPU usage) can be minimized.
- Collection Schedules allow unattended collection activity.
- A Collection Schedule allows the administrator to specify how frequently the data collection for the associated Collection Profile should occur by specifying a day-of-the-week or a day-of-the-month.

The components of a Collection Schedule are:

- Collection Profile name
- *Day*—This can be a date in a month (as in 10th) or a day of the week (as in Sunday)
- *Start Time*—This would be the time when the data collection should start (e.g., 23:00)
- *Stop Time (optional)*—If specified, the task will end at this time, interrupting any collections in progress

A Collection Schedule consists of a collection day (day-of-the-month or day-of-the-week) and a collection window (start time and optionally stop time). An open-ended collection window is a collection window that does not have an associated stop time. Collection activity starts when the specified start time is reached on the specified collection day(s). "Bounded Collections" include a built-in have error recovery mechanism built into them. If, for example, an error is encountered during a bounded collection, the NetFlow Server will perform collection retries until the collection is complete, or until the specified "stop time" is reached.

"On-Demand" NetFlow Server Collections

The On-Demand Collection feature, useful when using Scheduled Collections, enables the user to tell the NetFlow Server to collect datafiles immediately for a specified Collection Profile. The user needs to specify which Collection Profile needs to be collected, and can specify an optional stop time as well. Additionally, the user can also specify a Data Collection Range—this is a begin/end date/time range that further restricts what data will be collected.

"Continuous" NetFlow Server Collections

An "continuous collection" is a collection activity that starts as soon as the NetFlow Server is started, and the FlowCollector in the associated Collection Profile is available. Continuous collections provide a "retry time", which specifies how often continuous collections should be retried if no more datafiles are available from the FlowCollector(s) or a FlowCollector error is encountered. Continuous collections do not have associated stop times.

Benefits of using ongoing collections include:

- Conservation of valuable disk space on FlowCollectors, since datafiles are transferred to the NetFlow server soon after they are generated
- Faster availability of current FlowCollector datafiles at the NetFlow Server
- Reduction of CPU usage on the NetFlow Server, since data arrives continuously in smaller chunks, rather than in a single large batch

NetFlow Server Data Storage

The NetFlow Server allows the administrator to specify various time-based consolidation schemes in order to consolidate the data extracted from NetFlow Server datafiles. Consolidation of Like-Data is carried out across various time periods, such as daily, monthly, quarterly, or yearly, based on user configuration. Like-Data is defined as data that was aggregated using a common NetFlow FlowCollector Aggregation Scheme, such as CallRecord or HostMatrix. Once the data is transferred to the NetFlow Server, it can be configured to perform either, or both, of the following two tasks:

- Store the collected data in the database
  As datafiles are successfully transferred to the NetFlow Server, the datafile contents can be stored in the NetFlow Server's database. All data within the datafiles, as well as datafile attributes needed to identify one file from another, are stored in the database. Such stored data is called As-Is Data.
- Consolidate the collected data, and store the consolidated data in the database (time-based data consolidation)
  Independent of whether or not collected data is stored in the database As-Is, the NetFlow Server can also consolidate the data into periodic summaries, thereby allowing for easy (and fast) access to daily, monthly, etc. consolidations.

"As-Is" Data Storage

As-Is data is available in the database to the end-user with the same level of granularity as was available on the FlowCollector, without any loss of data or detail.

Storing As-Is data enables the NetFlow Server database to be the single point of access to all NetFlow data in the network, thereby eliminating the need for the end-user applications to access the FlowCollector(s) for any data already collected by the NetFlow Server, if desired.

Time-Based Data Consolidation

Time-based consolidation involves summarizing the numeric fields of the collected aggregated flow records (number of flows, packets, bytes, and timestamps) for Like-Data, that is, data produced by the same Aggregation Scheme and sharing the same key. Such data can be consolidated into daily, monthly, quarterly or yearly records.

Daily Consolidation involves summarizing the data for a given day into Daily Consolidated Data by summing the numeric fields for Like-Data for that day. For example, Daily Consolidation might involve summarizing the data from 96 SourceNode datafiles across a particular day from one FlowCollector, where each SourceNode datafile contains 15 minutes of NetFlow data into one file, where each aggregation key appears only once.

Also, as new datafiles are received from the FlowCollectors, "in-progress summaries" are kept updated for time-based consolidations. This enables access to updated and running summaries. When NetFlow Server detects that all the data for a particular time period has been received, it will mark the related consolidations as "closed".

Thus, even at mid-day, the Daily Consolidated Data can reflect up-to-the-minute summaries for all successfully received data pertaining to that day. Similar in-progress summaries for Monthly Consolidated Data, Quarterly Consolidated Data and Yearly Consolidated Data may exist.

Consolidation Schemes

The NetFlow Server provides two general types of consolidation schemes. The first type of scheme, Same-Router Consolidation (SRC), allows the user to aggregate Like-Data retrieved from a single router or router group by one NetFlow Server. The second scheme, Inter-Router Consolidation (IRC), allows the user to consolidate Like-Data across routers, across NetFlow FlowCollectors. Using the IRC scheme, like-Data from multiple source routers can be consolidated into daily, monthly, quarterly and/or yearly summaries, and stored in records under a logical router group name defined by the user.

Bi-directional Consolidation

By definition, NetFlow captures uni-directional flow statistics. For users wanting to consolidate bi-directional flows, that is, to summarize flow statistics which represent the two-way traffic between IP address pairs, the NetFlow Server provides "canned" SQL procedures which the user can leverage without change, or choose to modify/enhance them for different uses. The SQL procedures allow the user to summarize the byte, packet, flow counts and active flow times (when present) for flows between IP address pairs. The user can then use the SQL procedure to make requests such as:

"Summarize all the bi-directional traffic statistics between host-A and every other host that host-A talked with today."

SQL procedures are provided for each FlowCollector aggregation scheme, which contains IP address pairs in the aggregation scheme key. The SQL procedures can be applied to As-Is Data, or data produced by any time-based consolidation using either SRC or IRC schemes.

NetFlow Server Data Retention

It is reasonable to expect the volume of data collected and consolidated from the datafiles across multiple FlowCollectors to be large in some scenarios. As a result, some kind of method to effectively age the data and remove it from the database over time is required.

The NetFlow Server implements very flexible data retention functionality that will allow users to fine-tune retention to their specific needs. When used properly, it will ensure that NetFlow Server will always remove any data that it adds to the database, over a period of time.

The retention scheme will allow a maximum age for each FlowCollector Aggregation Scheme, and will allow the user to specify the following policies:

- *As-Is Data retention*—number of days for which As-Is Data is held. As new As-Is Data is created, any As-Is Data older than the retention specified is removed.
- *Daily Consolidated Data retention*—number of days for which Daily Consolidated Data is to be held. As new Daily Consolidated Data is created, any Daily Consolidated Data beyond the retention specified is removed.
- *Monthly Consolidated Data retention*—number of months for which Monthly Consolidated Data is to be held. Any Monthly Consolidated Data older than the specified retention value is removed.
- *Quarterly Consolidated Data retention*—number of quarters for which Quarterly Consolidated Data is to be held. Any Quarterly Consolidated Data older than the specified retention is removed.
- *Yearly Data retention*—number of years for which Yearly Consolidated Data is to be held. Any Yearly Consolidated Data beyond this specified retention is removed. It is expected that the administrator has an effective archival/backup policy in place that stores the Yearly Consolidated Data off-line, if necessary, before it is aged out.

Flow De-duplication

"De-duplicating" flows is defined as recognizing when two NetFlow data export records are redundant, i.e., where a single flow was seen and reported by two different routers in the network. In a billing/accounting application, it would be important to not double-count the statistics of such a flow.

Given the NetFlow Export data currently available, a simple flow correlation technique involves matching the aggregation keys of two or more records output by FlowCollector(s), and then compare the starting and ending flow timestamps to determine if an overlap exists. If an overlap exists, we can conclude that at least some portion of the flows is redundant. The NetFlow Server will provide flow de-duplication functionality, with some basic limitations:

- For very short flows (a few packets), or for flows which are measured by routers with non-trivial time differences between them, the timestamps may not overlap. In this case, it is impossible to recognize the flow records as being redundant. A single flow with this profile would be reported by each router and then double-counted by the NetFlow Server.
- Only a few of the FlowCollector aggregation schemes, such as DetailHostMatrix and CallRecord, contain sufficient detail (starting and ending flow timestamps) to be able to correlate duplicate flows in the network. The other aggregation schemes are lacking this detail and, therefore, cannot be used to recognize duplicate flows unless their records were extended to include timestamp fields (with possible performance and disk space impacts).
- Consider a single flow sent through two parallel routers using per-packet load balancing available with Cisco Express Forwarding (CEF). With NetFlow enabled, the two routers would each report a flow with similar timestamps. Given the timestamp approach to recognizing duplicate flows, the NetFlow Server would identify the NetFlow Export data from the two routers as being redundant even though they are not.

NetFlow Server Availability

Not yet available, the NetFlow Server is targeted for availability in mid-CY 1999.

NetFlow's Open Interfaces

Cisco provides a set of NetFlow-related open interfaces on routers/switches and NetFlow management utility products for use by our partners and customers in developing powerful, value-added network management and accounting applications including the following:

- *Routers*—The NetFlow Export interface specification is documented in the NetFlow Developer's Interface Specification document available on the Cisco Web site. The current Cisco IOS NetFlow code supports both version 1 and version 5 of NetFlow Export
- *FlowCollector*—The FlowCollector file formats including the Flow Detail Record format as well as the filters, aggregation schemes and thread language are documented in the FlowCollector User's Manual and Release notes on the Cisco Web site
- *FlowAnalyzer*—The FlowAnalyzer can export data to Excel Spreadsheet format on demand. The FlowAnalyzer user's manual and release notes are available on the Cisco Web site

Several customers (e.g. ANS and BBN with the cflowd suite) and partners (including HP, NetScout, Solect, Portal Software, X-Cel Communications, Belle Systems, Sequel Technologies, XaCCT Technologies, and Concord Communications) are utilizing the NetFlow interfaces either to build tools for internal use or to enhance existing products or build new applications.

Netsys Service Management Suite and NetFlow

NetFlow's ability to provide lightweight, granular data collection makes it a key feature enhancement source for several categories of software products including network planning, billing/accounting, network monitoring and data mining for outbound marketing purposes. In particular, the Netsys suite utilizes NetFlow data to provide additional customer value.

The combination of Cisco's Netsys Service-Level Management suite (Connectivity Service Manager and Performance Service Manager) and NetFlow data gives the most accurate and complete visibility into usage of the entire network of any available management tools. Network

planners can verify that the current network supports all required connectivity and services, and set service level agreements based on calculated loads. In addition they can test the network for reliability by performing failure analysis tests and examining both reachability and loading levels in off-line what-if tests.

To plan for growing user and traffic loads, planners can use the Netsys Service-Level Management suite to scale existing flows, then modify the topology, router configurations, or router models to best support the new levels. After performing the off-line what-if tests, planners can export actual Cisco IOS command files for loading on the live network.

Netsys Service Management Suite and NetFlow for Network Troubleshooting

Cisco's Netsys Service Management Suite adds significant network-wide context to the network monitoring process while operating almost exclusively in a non-intrusive mode.

Using the HostMatrix aggregation scheme of NetFlow Collector data, the Connectivity Service Manager shows the end-to-end path taken by a flow. By loading data from many flows simultaneously, users can see flows converge and diverge as they cross the network. This unique route path analysis enables network managers to improve path selection and correct security breaches.

The Performance Service Manager uses packet and byte counts at both protocol and application level to diagnose the source and content of flows contributing to over-threshold conditions. Combined with the route path analysis of the Connectivity Service Manager, this exposes to unprecedented level exactly who and what is causing load at any given point in the network. The Performance Service Manager even extends NetFlow capabilities to the Frame Relay environment, by showing to the circuit level the load and application mix contributed by each flow.

The typical use of the Netsys/NetFlow combination is to diagnose the causes of traps captured by a network management system (NMS). For instance, having been alerted by the NMS to an over-threshold condition on a given link, the network manager uses Netsys to pull the last "n" minutes of NetFlow data from the NetFlow FlowCollector. The manager is graphically shown all end systems, round trip paths, and applications causing the high utilization. He can then perform off-line what-if analyses to test a fix for the situation, and download the exact Cisco IOS commands required via CiscoWorks or Telnet to the router.

NetFlow Customer Applications

Many customers have leveraged NetFlow's powerful capabilities to build the following value-added applications and services:

Usage-Based Billing

In order to promote integration of NetFlow data into advanced usage-based Internet billing applications, Cisco is working closely with several third-party billing/accounting vendors, including Portal Software, Solect, Belle Systems, Saville Systems, X-Cel Communications, and Kenan Systems. These partners understand the value of NetFlow data for Internet billing applications, and are currently expending considerable development resources to integrate NetFlow as yet another billing data source in their offerings. In the future, NetFlow data will be a key enabler for Internet data usage-based billing, quality-of-service based Internet usage billing, time-of-day Internet usage billing, and Internet-based telephony billing (e.g. Voice over IP [VoIP]).

Network Architecture Planning

NetFlow provides the information to be the key source of data for powerful network architecture planning including peering planning and backbone and access transport network planning. A powerful example is the cflowd toolset created and utilized by ANS and BBN (refer to the following URL for details: `http://engr.ans.net/cflowd/index.html`). Quoting from Daniel McRobb of ANS:

> *"In years past, ANS collected traffic data (packet-sampled) from the NSS routers. The aggregate data included such things as port and protocol data (packets and bytes per port and per protocol) and net matrix data (how many packets and bytes were sent from network X to network Y). This data was used at ANS for a few things, but most importantly for network architecture planning. In the old days, ANS was able to collect this data at every ingress to ANSnet because an NSS sat at each ingress (or close enough to the ingress to be useful). Later, as the NSFNET service model changed and ANSnet expanded into more commercial areas, ANS enabled data collection on some interfaces on core NSS routers to allow us to continue collecting data from which to make decisions about our core architecture. When the NSS routers were decommissioned, ANS lost the ability to collect this data, and hence lost a source of*

very useful information. cflowd is intended to replace the system used on the NSS routers, by collecting data from Cisco routers via flow-export and flow-switching."

The NetFlow data is utilized to create an extensive set of graphs and visuals, which can be viewed at the above URL.

Dial User Profiling

Internet Service providers also plan to utilize NetFlow data in conjunction with data captured from other sources such as Radius servers to build detailed profiles of dial users acquired from dial servers both for network planning purposes as well as marketing purposes. The NetFlow data allows service providers to understand application usage by each customer by time, data volume and destination and to construct high impact marketing programs, which leverage this information.

Distance-Based Charging

Internet Service Providers are utilizing NetFlow's source and destination autonomous system (AS#) number recording capabilities both for network planning purposes (who to peer with and bandwidth requirements) as well as accounting purposes. In particular, the AS# can be used to determine if packets are bound for a local, domestic, regional or distant international destination and both distance and volume based usage charged accordingly to offset scarce network resource utilization.

NetFlow and RMON

NetFlow constitutes only one component of a comprehensive management data collection strategy. NetFlow is not intended as a replacement for SNMP and/or RMON/ RMON2 capabilities. Instead, Cisco suggests that SNMP, RMON and NetFlow data collection capabilities be combined to maximize network monitoring, management, and planning capabilities. Several recommendations follow:

- Utilize carefully designed SNMP polling policies to gather key statistics on backbone/core routers and on MIB objects not related to flow-by-flow measurements including interface errors and memory and CPU utilization statistics required for real time monitoring
- Utilize RMON capabilities for detailed drilldown including application tracking, interface error analysis and packet capture for problem diagnosis and resolution and for threshold driven network management via the events and alarms capability

NetFlow statistics gathered in Cisco routers and switches can also be exported to Cisco SwitchProbe devices. Ethernet SwitchProbes with software version 4.2 support NetFlow Export versions 1 and 5. SwitchProbe software version 4.5 will introduce support for NetFlow Export version 7, enabling data collection from the Catalyst 5000 series NetFlow Feature Card (NFFC). SwitchProbe software version 4.5 will also enable NetFlow capabilities on the Fast Ethernet SwitchProbe.

When NetFlow Monitor is installed in a SwitchProbe device, the device creates a special internal interface used for NetFlow data consumption. CWSI's traffic management application can be configured to monitor this interface. This data can be proxied (mapped) to RMON and RMON2 for further traffic analysis, thereby providing RMON and RMON2 support on the router backbone. One probe is required for each router. Multiple router support will be implemented in the SwitchProbe software version 4.5, targeted for availability in Q4 CY 98.

For additional information on Cisco's SwitchProbe product line, please see: `http://www.cisco.com/warp/customer/cc/cisco/mkt/enm/sprobe/index.shtml`.

Summary

NetFlow technology efficiently provides the metering base for a key set of applications including accounting/billing, network planning, network monitoring and outbound marketing for both service provider and enterprise customers. Cisco also provides a set of NetFlow management utilities to collect flow export data, perform data volume reduction, post-processing and storage and make flow detail records available to consumer applications in a convenient format. Cisco is also working with a growing set of application partners including HP, NetScout Systems, Concord Communications, Solect, Portal Software, X-Cel Communications, and Belle Systems to integrate NetFlow Export data and usage records with key applications including network monitoring and billing, rating and provisioning.

NetFlow does not require adoption of new or proprietary protocols or new generations of networking equipment. NetFlow is available today on many Cisco platforms. NetFlow may be deployed incrementally, on an interface-by-interface basis on strategically located edge, aggregation or WAN access routers. NetFlow data collection and export will also serve as a key enabler for flexible, differentiated IP services based on Cisco IOS QoS capabilities. NetFlow provides an incremental path to high performance, services rich networking environments while providing maximal investment protection for the installed base of network equipment.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
http://www-europe.cisco.com
Tel: 33 1 69 18 61 00
Fax: 33 1 69 28 83 26

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at http://www.cisco.com/offices.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela

3/99 LW