Emerald

Network Collector

Version 5.0.52

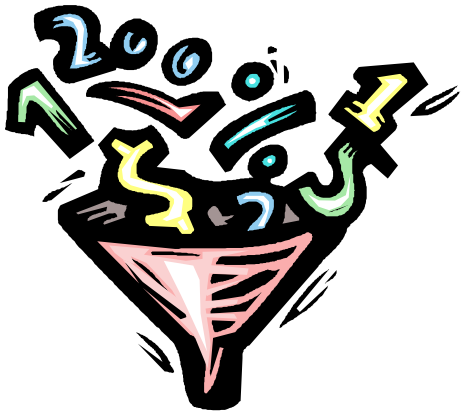

Emerald Management Suite

IEA Software, Inc.

Software License Agreement

By purchasing or installing RadiusNT or RadiusX, you indicate your acceptance of the following License Agreement.

Ownership of Software You acknowledge and agree that the computer program(s) and associated documentation contained with RadiusNT or RadiusX (collectively, the Software) are owned exclusively by IEA Software, Inc. and/or its licensors. The Software contained in the package is protected under copyright laws and all copyright and other intellectual property rights relating to the Software are and remain the exclusive property of IEA Software, Inc. and/or its licensors. You may not rent or lease the Software, but you may transfer the Software and accompanying materials on a permanent basis provided you retain no copies and the recipient agrees to the terms of this Agreement.

License IEA Software, Inc. grants to you, and you accept, a limited, non-exclusive and revocable license to use the Software. You agree to use the Software in machine-readable object code form only as authorized in this License Agreement. This License Agreement does not convey any title or interest in the Software to you.

Scope of License You may not make any changes or modifications to the Software, and you may not de-compile, disassemble, or otherwise reverse engineer the Software. You may not lend, rent, lease or sublicense the Software or any copy to others for any purpose. RadiusNT or RadiusX may only be installed on a single WindowsNT, Solaris, Linux or FreeBSD workstation or server. Additional servers may be purchased separately. You agree to use reasonable efforts to protect the Software from unauthorized use, modifications, reproduction, distribution and publication. You are not permitted to make any uses or copies of the Software that are not specifically authorized by the terms of this License Agreement. Your adherence to this License Agreement will allow IEA Software, Inc. to continue developing innovative and useful products and providing a high level of customer service and support. If you do not comply with the terms of this License Agreement, your license will be revoked.

Updates and Support All software updates and fixes are available via the IEA Software, Inc. Web site. Major version upgrades are not included or covered as part of the basic purchase agreement. Technical support is currently available via methods listed on our Web site Support section at http://www.iea-software.com/support.

Restricted Rights The Software is provided with U.S. Governmental Restricted Rights. Use, duplication, or disclosure by the Government is subject to restrictions set forth in subparagraph ©(1)(ii) of The Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or subparagraphs ©(1) and (2) of the Commercial Computer Software - Restricted Rights at 48 CFR 52.227-19 as applicable. The Software is also protected by International Treaty Provisions. Manufacturer is IEA Software, Inc. PO BOX 1170 Veradale WA, 99037.

Miscellaneous This License Agreement shall be construed, interpreted and governed by the laws of the State of Washington. Should any term of this License Agreement be declared void or unenforceable by any court of competent jurisdiction, enforcement of the remaining terms shall not be affected. Failure of either party to enforce any rights or to take action against the other party in the event of any breach of this Licensing Agreement shall not be deemed a waiver of any subsequent enforcement of rights.

Limitations of Liability and Remedies In no event shall IEA Software, Inc. or its licensors be liable for any loss of profit or any other commercial damage, including but not limited to special, incidental, consequential or other damage, even if IEA Software, Inc. or its licensors are advised, in advance, of the possibility of such damages. IEA Software, Inc. and its licensor's entire liability and your exclusive remedy shall be, at IEA Software's option, either (a) return of the price paid, or (b) repair or replacement of the Software. To the maximum extent permitted by applicable law, IEA Software, Inc. and its licensors disclaim all other warranties, either express or implied, including but not limited to, implied warranties with regard to the Software and the accompanying material. This Limited Warranty is void if failure of the Software has resulted from accident, abuse or misapplication. You may have other specific legal rights, which vary from state/jurisdiction to state/jurisdiction.

Return Policy It is our goal to provide customers with the highest level of satisfaction possible. In order to ensure that our products work well in your environment, IEA Software offers a 30-day FULL functioning software trial that includes documentation and support. If you require more than 30 days to evaluate the software, we are happy to work with you to extend the trial to a length that fits your timetable. This gives you, the user, an opportunity to ensure that the product fully meets your needs. (Please test the software in a non-production environment.) In light of the trial period and opportunity to fully test our software, IEA Software maintains the policy that no refunds will be offered. We will, however, address any problems with the software.

Should a software anomaly occur, our Development and Support Teams will work to correct the problem. Please note that you must be using the application normally, as defined, and you must ensure that the bug is not due to anomalies in other programs, the operating system, your hardware, or data.

In order to address any problems, please note that the bug must be able to be reproduced. Our Development and Support Teams will require full documentation of the steps taken by the user that caused the error in the software as well as necessary data and scenario files to reproduce the error.

Contact Should you have any questions concerning this license agreement, please contact IEA Software, Inc. PO BOX 1170 Veradale, WA 99037 U.S.A. (509) 444-BILL (2455).

Trademarks

Emerald Management Suite, *RadiusNT* and *RadiusX* are trademarks of IEA Software, Inc. All images, photographs, animations, audio, video and text incorporated into the Software are owned by IEA Software, Inc., unless otherwise noted by Trademark. *Intel* is a registered trademark of Intel Corporation. *Transact-SQL* is a registered trademark and *DB-Library* is a trademark of Sybase, Inc. *Sun Solaris* are trademarks of Sun Microsystems. *Cisco* is a trademark of Cisco Systems. All other trademarks are the property of their respective owners.

Table Of Contents

Software License Agreement **1**
Software License Agreement **2**
Trademarks 3
Overview **5**
Installation **5**
Configuration **5**
General Options 5
Rating Options 7
Flow Filter Options 7
Database Configuration 7
Emerald Configuration 7
Debug and Logging 8
Collector Monitoring **8**
Rating upload queue 8
Rating Statistics 8
Flow filter statistics 10
Reload rating rules 10
Configuring NetFlow & IPFIX Clients **10**
Cisco Routers 10
fprobe 11
ntop 11
Network Monitoring 11

Overview

EmerNet is an integrated aggregation and rating module to the Emerald Management Suite enabling Emerald to summarize and bill for network traffic. For network traffic collection Cisco Netflow versions 1, 5, 7, 9 and IPFIX (RFC 5101) are supported as well as passive monitoring of Ethernet traffic over the wire.

Installation

EmerNet is included as part of the Emerald distribution however it is an optional component and must be explicitly selected for install during the installation of Emerald. If you will be installing EmerNet on a separate server dedicated for network data collection you may choose to install only the EmerNet component. Please see the Emerald documentation for more information on installing Emerald. **Note: On the windows platform before starting EmerNet for the first time the network capture driver located in the Emerald folder (winpcap_*) must be installed by executing this file.**

Configuration

Configuring EmerNet is accomplished via the EmerNet Configuration server. To start the EmerNet Configuration server in initial configuration mode follow these steps:

Windows platform:

1. From your Start Menu, Program Files, Emerald, Server menu, select EmerNet Server.
2. From your Start Menu, Program Files, Emerald menu, select EmerNet Config.

UNIX platform or manual installation:

1. From the command line switch to the folder EmerNet has been installed (/usr/local/emerald). Within the directory execute the command ‘./emernet –config’
2. From a web browser connect to the server on port 8018. http://127.0.0.1:8018

General Options

General options menu configures the main data collection method and miscellaneous management related options. The two data collection methods available are Cisco Netflow (versions 1,5, 7, 9 and IPFIX) and passive collection by placing the collectors Ethernet interface in promiscuous mode.

Cisco Netflow requires a router capable of exporting traffic flow summaries. The Cisco Netflow formats are supported in hardware by many Non-Cisco vendors and additionally software solutions such as fprobe (http://fprobe.sourceforge.net) that translate local traffic into Cisco Netflow exports. EmerNet does not support aggregated export versions such as Cisco flow version 8 or sampled IPFIX as they do not provide enough information for billing to occur or be properly accounted for by the Emerald-rating engine. Cisco Netflow and IPFIX are ‘one sided’ protocols. They offers no security against spoofing flow records and no retransmission options should the collector not be running or there is too much traffic to account for. When using EmerNet in production using Netflow and IPFIX it is recommended it be connected directly to a dedicated network interface on a dedicated network segment to ensure security and reliability of traffic collection.

Promiscuous mode data collection is an alternative when Cisco Netflow is not available. In this mode all traffic going over the same Ethernet segment as the collector is summarized and rated. Promiscuous mode collection in EmerNet is currently limited to Ethernet interfaces and works only with IPv4. In most Ethernet environments you will need to configure a ‘Monitor’ or ‘Mirror’ port in the Ethernet switch to send all traffic to the EmerNet collector so that it can see the networks data traffic in order to properly summarize and rate it.

Field	Description
Configuration server	When enabled the configuration web server is started with Emernet. When disabled the configuration server is only available when started manually by running ‘./emernet –config’
Configuration server port	TCP port the configuration server listens for incoming HTTP requests.
Data collection	Selects the data collection method. Only one data collection method can be enabled for a collector. When set to Netflow & IPFIX EmerNet listens for flow data sent to it. When set to ‘Local Packet capture’ EmerNet puts the local Ethernet interface in into promiscuous mode and captures data directly from the Ethernet interface.
Concurrent flow processors	The number of threads dedicated to rating flow records. This should reflect the number of concurrent execution threads supported by the collectors CPU. For example with a dual core processor you would set this to 2 or 4 for a dual core hyper-threaded processor. Setting the number of flow processors higher than the sum of all CPU’s available concurrent execution threads will decrease performance.
Process queue	Maximum amount of flow records having been evicted from the flow cache that can be queued for processing by the rating engine. A value of between 1000 and no higher than 10000 is recommended. If the process queue is full the flow record is discarded.
Flow cache size	Approx amount of unique traffic flows that can be stored and aggregated in main memory. Increasing the flow cache size can significantly improve processing performance by aggregating more traffic flow data before being rated at the expense of higher memory utilization. Each flow cache entry requires about 270 bytes of contiguous main memory.
Flow cache max TTL	Number of seconds a cached flow can remain cached before being evicted to the process queue. Increasing this value improves performance by providing more opportunity for flows to be combined before being rated at the expense of minor delay in the rating of collected flows.
Netflow/IPFIX listen port	UDP port to listen for flow exports. The default port is 4739.
Allowed NetFlow/IPFIX export hosts	List of IP Addresses allowed to send netflow data to EmerNet. Note: flow data can be trivially spoofed - the allowed export list must not be relied on and should be used in conjunction with access control lists or other methods to verify source address.
Packet Capture interface	Name of the local ethernet interface to capture network traffic from when the Data collection mode is set to ‘Local packet capture’

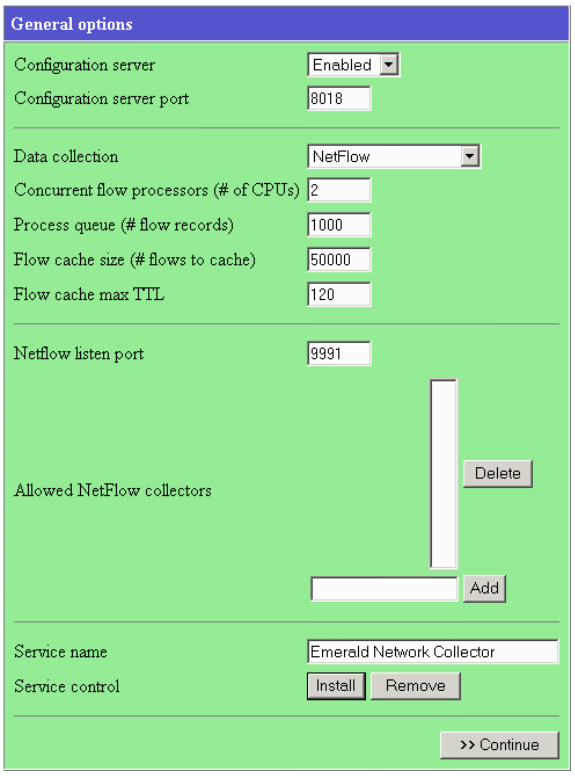

Rating Options

This menu provides configuration for Emerald rating engine.

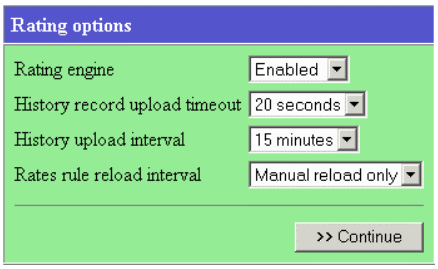

Field	Description
Rating Engine	This enables and disables the rating engine. It should always be enabled.
History record upload timeout	Query execution timeout when updating the RateHistory table. This is intended only as a backup to break deadlocks should one occur and the RDBMS lock manager is not able to clear it. It should not be set lower than 20 seconds.
History upload interval	Sets the interval at which rated flow records are uploaded to the database. Higher values place less load on the database server at the expense of less frequent updates.
Rates rule reload interval	This option determines how often rating rules and rating classifiers are refreshed from the database. Reloading this information often will ensure the working rating configuration is kept up to date.

Flow Filter Options

This menu is provided for compatibility with flow logging features used for rating in earlier versions of Emerald and is no longer supported.

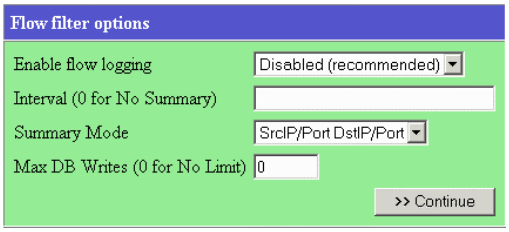

Database Configuration

Select the Database Configuration menu option. After changing the settings, click continue. Please note, these settings are commonly shared between all Emerald services. Therefore, if you already have Emerald or other Emerald services installed and configured on this machine, these settings should already be set and you should ignore the database configuration menu. Changing these settings will affect all other Emerald suite applications as well.

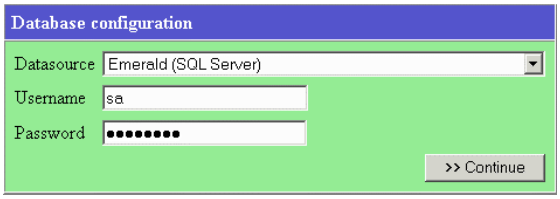

- Select a data source item. This is the ODBC DSN that you will use to connect to the Emerald database. If you do not have a DSN defined, you can select (new) to create a new DSN.
- Fill in the username and password to connect to the database as.

Emerald Configuration

Rating of network flow traffic depends on the assignment of the “Netflow IP Address” and “Netflow Collector IP” custom data fields to all service types that will be billed for netflow traffic.

“Netflow IP Address” represents the IP address or address mask of the IP or network that traffic destined to or originating from will be billed. This field must be specified in order for the end user to be billed for their network traffic. It is recommended this data field be set to required for netflow based service types to prevent operators from

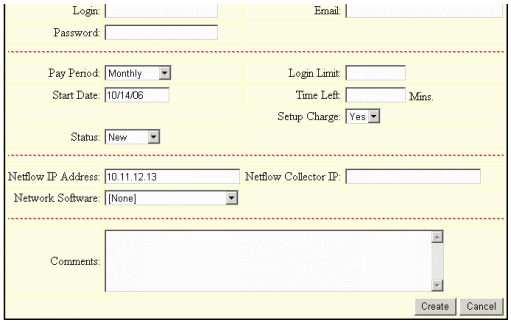

forgetting to enter the end users IP. “Netflow Collector IP” is an optional field seting the collector address responsible for collection of traffic for this account. This is useful only when there are multiple flow probes throughout the border and cores of the network and you need to prevent transit between routers from causing the customer to be double billed. Please see the Rating section of the Emerald Administrators guide for details on configuring rating for network traffic.

Debug and Logging

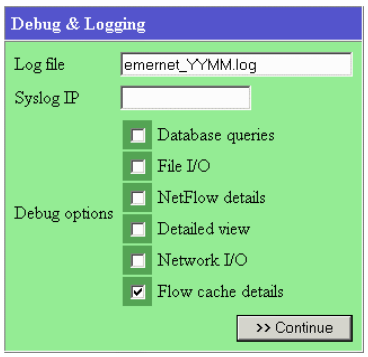

This section is used to assist in troubleshooting and monitoring the EmerNet server. We recommend all debug options be disabled unless there is a specific reason to enable them. This way only salient warning and error messages are sent to the log file. The log file name accepts special characters ‘DD’, ‘YY’ and ‘MM’ for Day, Year and Month respectively. When these characters are entered the resulting log file name is replaced with the values for the current day. In addition to the log file you may send logging messages to a central Syslog server including the Syslog server included with the Emerald suite.
If you need to trouble shoot the EmerNet collector an easy way to do so monitoring its actions in real-time is by running it in debug mode from the command line. To do this first stop the Emerald netflow collector if it is running as a windows service or UNIX background task. Next run ‘./emernet –debug 15’ from the Emerald folder to start the collector in debug mode. When in this mode all messages are sent immediately to the screen and not routed to either the log file and or syslog server.

Collector Monitoring

EmerNet includes online monitoring options ‘Rating upload queue’ used to view rated network flows pending upload to the database and ‘Rating statistics’ used to monitor the performance of the network collector.

Rating upload queue

Shows the results of all rated flows that have not yet been uploaded to the database. This view is normally reset as the database is updated. Each row reflects the rated usage of a single customer. If multiple rates are defined a customer may have more than one row associated with them.

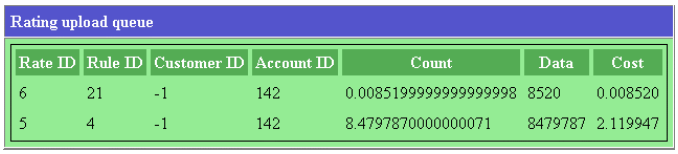

Field	Description
Rate ID	The rate identifier of a configured Emerald Rate available from the ID column of the rate listing within Emeralds Admin / Rating / Rates menu.
Rule ID	Rating rule invoked to rate the applicable flows available from the ID column of the Rule Set listing within Emerald Admin / Rating / Rule Sets.
Customer ID	The Emerald MBR being billed for usage. In most cases Customer ID is displayed as ‘-1’ meaning that the MBR directly associated with the service responsible for generating the actual network usage (See Account ID below). Generally whenever Customer ID is not ‘-1’ the rate is being used to charge a reseller or other third party not directly associated with the service.
Account ID	The Emerald service responsible for the generating network usage.
Count	Reflects the number of ‘Intervals’ as configured in the rating rule set (See RuleID above) that have been rated.
Data	Data always reflects the number of bytes recorded.
Cost	The configured cost based on Count, Data and possibly specifics of individual flows and classifiers. Costs are defined from the Emerald Admin / Rating / Rule Sets menu.

Rating Statistics

Displays the status and current performance of the rating engine.

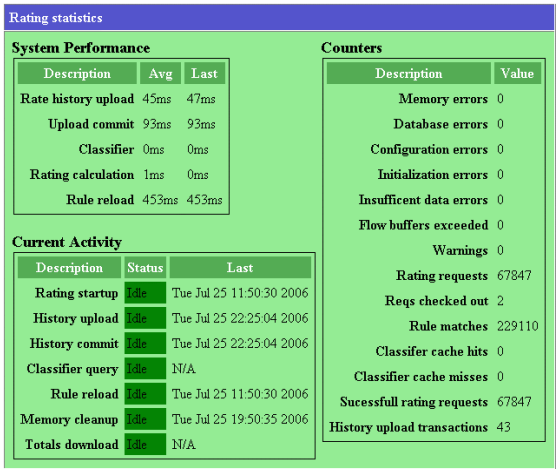

System Performance / Current Activity

Field	Description
Rate history upload	Amount of time to upload a single usage record to the database.
Upload commit	Time needed to commit a history upload batch transaction
Classifier	Time required to match a network flow with an Emerald service.
Rating calculation	Time required to rate a single network flow record
Rule reload	Time required to refresh all applicable rating rules and classifiers.

Counters

Field	Description
Memory errors	Count of all memory allocation errors encountered. If this counter ever increases consider adding more physical or virtual memory to the collector. If you have configured an excessive flow cache size 1>million records lowering the value may correct memory allocation failures.
Database errors	Total count of database errors encountered, normally this may increment slightly if connections need to be reestablished to the database server or during times when the database server is unavailable. You should consult the EmerNet log file to view detailed information about any database errors.
Configuration errors	Configuration errors are caused by invalid or inconsistent rating configurations for example the referencing of rating rules that don't exist or selecting an unknown rule match type. If EmerNet is running with a clean configuration with no inconsistent data it will not replace its configuration during a rule refresh if the new configuration is inconsistent. The Emerald user interface effectively prevents the possibility of configuration errors however user customizations or direct configuration may lead to the problem. If this counter is incremented view the EmerNet log file for detailed information about the configuration error and how to correct it.
Initialization errors	Initialization errors point to an internal problem within EmerNet itself and should be reported to your support representative.
Insufficient data errors	These occur when there is not enough or incorrect information presented to be able to properly rate a given flow record. See the EmerNet log file for details.
Flow buffers exceeded	When this field is incremented it means EmerNet was not able to rate all incoming flows in time and was forced to discard a flow record. This condition can often be corrected by increasing 'Flow Cache Size' and 'Flow Cache Max TTL' values in the General options menu. If you run into this problem and are using Netflow export from a Cisco router or a netflow probe you may be able to increase the local 'aggregation' times so that less flow records are ultimately exported to EmerNet. If the buffer exceeds seem to happen in spurts rather than evenly over time increasing the 'Process queue' value in the General options menu help but do this only if the above fails and do not exceed 10,000 flow records. The EmerNet log file will periodically provide summaries of the number of flow buffers exceeded over time.
Missed flow packets	When Netflow V9 or IPFIX is used this field may be incremented to reflect flow packets not successfully delivered to EmerNet. This represents an estimate of missing flows based on gaps in the sequence numbers of successive received flows.
Sequence violations	When Netflow V9 or IPFIX is used a small number of sequence violations can occur normally due to common events such as a router reboot or as a result of periods of network connectivity problems between EmerNet and a device emitting flow records. Large numbers of sequence violations may indicate attempts to exploit the collection of traffic flows, or incorrect IPFIX implementation.
Incomplete flow records	When Netflow V9 or IPFIX is used EmerNet requires at the very least Ipv4 source, destination and byte count fields be present in exported flow data. If the incomplete flow counter is incrementing devices exporting netflow may need to be reconfigured to specifically provide this information or disable incompatible aggregation formats. IP information is required by the Emerald rating engine to properly account for customer usage.
Template lookup failures	When Netflow V9 or IPFIX are used templates are sent to EmerNet enabling it to decode flow data records. Whenever template information is not available for a particular flow this field is incremented and the associated network flows are not recorded. This field may normally increment slightly but large counts may indicate a problem with the system exporting netflow. You may need to reconfigure the template export intervals to occur more frequently.
Warnings	Count of warning level messages recorded to the EmerNet log file.

Rating Requests	Total number of rating requests processed.
Reqs checked out	This always translates to the number of flow processors configured (See 'Concurrent flow processors' in the General options menu)
Rule matches	Amount of rating rule matches, usually many times higher than the number of rating requests.
Classifier cache hits	This should never increment, using rating classifiers based on database queries with EmerNet must be avoided. All netflow classifiers included with Emerald load all their data into memory using upload attributes.
Classifier cache misses	This should never increment, using rating classifiers based on database queries with EmerNet must be avoided. All netflow classifiers included with Emerald load all their data into memory using upload attributes.
Successful rating requests	The number of rating requests successfully processed, ideally this number should match the 'Rating Requests' number above.
History upload transactions	Count of historical uploads transactions done since EmerNet startup. This does not reflect the number of Emerald (RateHistory table) updates rather the number of history update transactions that usually contain several individual updates.

Flow filter statistics

Flow filters statistics are only available when flow filtering is enabled. This is provided only for backwards compatibility for those migrating from previous versions of Emerald. Use of flow filtering is not supported and not recommended.

Reload rating rules

Refreshes all applicable rating rules and rating classifiers from the Emerald database. EmerNet can be configured to refresh this information periodically via the 'Rates rule reload interval' setting found in the 'Rating options' menu.

Configuring NetFlow & IPFIX Clients

There are many choices for NetFlow & IPFIX clients. Most router vendors are capable of sending flow packets based on a global or per interface settings. Alternately EmerNet can directly monitor flows on Ethernet segments by configuring a 'Mirror' or 'Monitor' port at the Ethernet switch.

There are several flow formats routers may send to EmerNet. The five currently supported flow formats are Cisco Netflow V1, V5, V7, V9 and IPFIX . Although each version contains different data, all versions contain the base information EmerNet needs to perform filter and aggregation functions.

Aggregated flow formats including Cisco Netflow version 8 or IPFIX configured to emit aggregated summaries are not supported and must not be used with EmerNet. These formats enable efficient collection of statistics for network management and planning purposes however they lack necessary specificity to guarantee network flows are accurately associated and billed to end users.

Cisco Routers

To configure a Cisco router for NetFlow requires two steps:

1. Enable NetFlow at the Global level. You will define the version of the flow to send and what address to send the flow to.

2. Enable NetFlow per interface you want to collect statistics for.

For more information on configuring NetFlow switching in a Cisco router, please see the follow URL on the Cisco website:

http://www.cisco.com/en/US/docs/ios/12_1/switch/configuration/guide/xcdnfc.html

fprobe

fprobe is network monitoring server for UNIX scanning local traffic on its network segment and sending Netflow exports to a remote server such as EmerNet for processing.

http://fprobe.sourceforge.net/

Command line example of using fprobe: fprobe 1.2.3.4:4739

ntop

ntop is a network monitoring utility that can send NetFlow packets. You can find out additional information about ntop from the ntop website: http://www.ntop.org

For more information on ntop, please see the ntop website listed above.

Network Monitoring

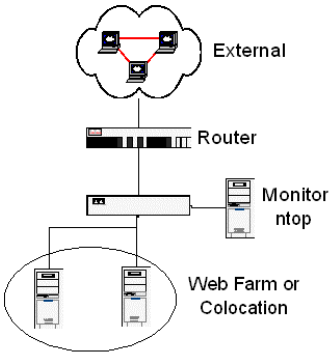

You need to plan where to enable Netflow on your network. In a simple scenario, you can enable NetFlow at your Internet or external connection. This will send all traffic passing out through the Internet to your NetFlow collector. However, if you want to collect only for a limited number of computers (a web farm or co-location facility) this may cause a significant amount of unwanted traffic between the NetFlow client and EmerNet.

Another scenario possible uses an Ethernet switch and EmerNet in local packet capture mode. In this configuration, you would enable the port the EmerNet collector plugs into to see all traffic across the switch. This is often called the “monitor” or “mirror” port. For 3Com switches, this is often called the analysis port. Please consult your switch documentation for further details on how to configure this setup. If you switch does not have a monitor or analysis function, you can use a hub to accomplish the same function. You would plug into the hub a connection to the switch for your web farm or colocation, the monitor computer, and a connection to the Internet or external network. This limits the traffic the monitor sees only to that passing through the hub.