CHAPTER 2

Traffic Processing Overview

This module describes how the Cisco Network Module Enhanced Application Performance Assurance (NME-APA) processes traffic.

The module also defines the main elements (configuration entities) and explains how they relate to each other.

Routing Environment

The Cisco Application Performance Assurance solution operates in a Symmetric routing environment. Inbound and outbound traffic is routed through one NME-APA. For a marginal number of flows only one direction goes through the NME-APA.

Traffic Processing

There are two stages of traffic processing:

- Traffic classification— analyses traffic flows and determines their type (for example, browsing, e-mail, file sharing, or voice).
- Traffic accounting and reporting— performs bookkeeping and generates Raw Data Records (RDRs) that let you analyze and monitor the network.

These two stages are described in the following sections.

You control how classification and reporting are performed by editing configurations and applying them to the NME-APA.

Traffic Classification

Traffic processing starts with traffic classification , which categorizes network sessions into classes.

For each network application traversing an enterprise's infrastructure, a corresponding class is defined in the Cisco Application Performance Assurance solution. You can use this class to classify and identify the traffic and report on its usage.

Classes

In the traffic classification process, categorizes network sessions into *classes* .

Classes are the building blocks for:

- Class configurations
- Aggregated usage reporting

From a enterprise’s point of view, a class is is usually a network application—such as browsing, e-mail, file sharing, or voice—that the user uses. From a technical point of view, a class consists of one or more class elements, each of which enables a decision about the class associated with a network traffic flow type.

A number of classes are predefined in the default configuration. You can modify these classes and add additional classes to a configuration.

A configuration can contain up to 500 classes.

The classification process occurs when a flow is established. The process examines the first few packets of the session and decides to which class the session belongs. The session is then assigned a class ID that remains the same during the session’s life cycle.

Traffic is classified and mapped to classes on the basis of some or all of the following class elements:

- Protocol—The protocol used. This allows, for example, the mapping of browsing flows and e-mail flows to separate classes.
- Zone—Lists of IP addresses of the network-side host of the flow. This allows, for example, the mapping of all voice flows going to a specified server to a specific class.
- Flavor—Specific Layer 7 properties such as host names of the network-side host of the flow. This allows, for example, the mapping of all HTTP flows where the URL matches a certain pattern to a specific class.

The NME-APA uses these flow mappings to map each network connection passing through it to a class.

Class Elements

A class consists of one or more class elements; different network traffic flow types are mapped to different class elements.

A class element maps a specific protocol, initiating side, zone, and flavor to the selected class. Some or all of these parameters can take wild-card values.

Note When asymmetric routing classification mode is enabled, the flavor of a class element is always the wild-card value.

A traffic flow is mapped to a specific class if it meets all four of the following criteria:

- The flow uses the specified *protocol* of the class element.
- The flow matches the *initiating side* specified for the class element.

- The destination of the flow is an address that belongs to the specified *zone* of the class element.
- The flow matches the specified *flavor* of the class element.
- If a flow matches two class elements and one is more specific than the other, the flow will be mapped to the more specific of the two. For example: Class A is defined for browsing and Class B is defined for browsing to a specific list of URLs. A browsing flow to a URL on Class B's list matches both classes, but will be mapped to Class B.
- If a flow matches one parameter of one class element and a different parameter of another class element, precedence will be given first to matching flavors, then to protocols, then to zones, and finally to the initiating side. For example: Class A is defined for e-mail and Class B is defined for all traffic to a specific network zone. An e-mail flow to the specific network zone matches both classes, but will be mappedto Class A.

Examples of Classes

The following table contains examples of classes and their network parameters.

Table 2-1 Examples of Classes and Class Parameters

Class Name	Protocol	Initiating Side	Zone	Flavor
Web Browsing	HTTP HTTPS	User-initiated		
Web Hosting (network-initiated browsing)	HTTP HTTPS	Network-initiated		
Local SMTP	SMTP		Local-mail servers (215.53.64.0/24)	

Signatures

The NME-APA examines traffic flows using its deep-packet-inspection capabilities, and compares each flow with an installed set of protocol signatures to identify the network application that generated the flow.

The NME-APA comes with a set of predefined signatures for common network applications and protocols, such as browsing, e-mail, file sharing, and VoIP.

Protocols

One of the main classifications of a flow is the protocol of a session (that is, of the network application that generated the session).

A protocol, as defined in the system, is a combination of one or more signatures, one or more port numbers, and a transport type. The protocol of the network flow is identified according to these parameters. For example, if the port number is 80, the transport type is TCP, and content matches the HTTP signature, the NME-APA maps the flow to the HTTP protocol.

The default configuration contains a long list of predefined protocols. You can add additional protocols.

When a TCP or UDP flow does not match a specific protocol definition, the NME-APA maps the flow to the Generic TCP or Generic UDP protocol.

When a non-TCP/UDP flow does not match a specific protocol definition, the NME-APA maps the flow to the Generic IP protocol.

Protocol Elements

A protocol is a collection of protocol elements.

A protocol element maps a specific signature, IP protocol, and port range to the selected protocol. Some or all of these parameters can take wild-card values; port numbers can take range values.

A traffic flow is mapped to a specific protocol if it meets all three of the following criteria:

- The flow matches the specified signature of the protocol element.
- The flow protocol matches the IP Protocol of the protocol element.
- The flow matches the specified port range of the protocol element.

  If a flow matches two protocol elements and one is more specific than the other, the flow will be mapped to the more specific of the two. For example: Protocol A is defined for flows that match the FTP signature and Protocol B is defined for flows that match the FTP signature on TCP port 21. An FTP flow on port 21 matches both protocols, but will be mapped to Protocol B.

  If a flow matches the signature of one protocol element and the port of another protocol element, it will be mapped to the matching signature. For example: Protocol A is defined for flows that match the FTP signature and Protocol B is defined for flows on TCP port 21. An FTP flow on port 21 matches both protocols, but will be mapped to Protocol A.

Initiating Side

The Application Performance Assurance solution is located between the enterprise's users and the network. User-initiated flows are initiated by the user toward the network; network-initiated flows are initiated from the network toward the user.

You can monitor some flow-types to one initiating side. For example, with HTTP you can monitor the user-initiated traffic separately from the network-initiated traffic. HTTP is always user-initiated when the user ventures outward to surf the Internet. If the direction of the HTTP flow is network-initiated, this probably means that a web server is open on the user's local machine for receiving incoming HTTP traffic. The enterprise can monitor the network-initiated HTTP and use other criteria to evaluate if the traffic is legitimate.

Zones

A zone is a collection of network-side IP addresses.

You configure zones by arranging IP addresses in groups connected by a common purpose. A user's network flow mapped to a class may be applied to a zone. In practice, zones often define geographical areas.

Zones are used to classify network sessions; each network session can be assigned to a class element based on its destination IP address.

Zone Elements

A zone is a collection of related zone elements.

A zone element is an IP address or a range of IP addresses.

Table 2-2 Examples of Zone Items

Network Address	Example
IP address	123.123.3.2
IP address range (and mask)	123.3.123.0/24 This means that the first 24 bits of the IP address must be included as specified and the final 8 bits can take any value. (That is, all IP addresses in the range 123.3.123.0 to 123.3.123.255.)

EXAMPLES OF ZONES:

- A “walled garden”—A range of IP addresses of a server farm with premium video content, for which the enterprise would like to limit access to specific users and to assure traffic priority.
- A zone to differentiate between off-net and on-net flows.

EXAMPLE OF ASSIGNING A ZONE TO A SESSION:

- Zone A and Zone B are two user-defined zones. Zone A includes the IP address range 10.1.0.0/16, and Zone B includes the IP address range 10.2.0.0/16. Analysis of a new session shows that its network IP address is 10.1.1.1—the session belongs to zone A.

Flavors

Flavors are advanced classification elements that classify network sessions according to signature-specific Layer 7 properties.

Flavors provide an additional level of granularity in defining classes in the Cisco Application Performance Assurance solution. A protocol flavor uses an additional protocol attribute in classifying a class, making this class a flavor of the class based on the protocol only. For example, the user-agent attribute of the HTTP protocol could be added as a protocol flavor, enabling the definition of all HTTP traffic generated by the same browser type (indicated in the user-agent field) as one class.

Examples of flavor types are HTTP User Agent and SIP Source Domain.

Flavor Elements

A flavor is a collection of flavor elements .

The type of a flavor element depends on the flavor type.

The default configuration includes some predefined flavors, such as HTTP Streaming Agents (a flavor of HTTP) and Vonage (a flavor of SIP).

Mapping Flow Attributes to Classes

The following figure illustrates the mappings of flow elements of a session to class elements of a class.

Figure 2-1

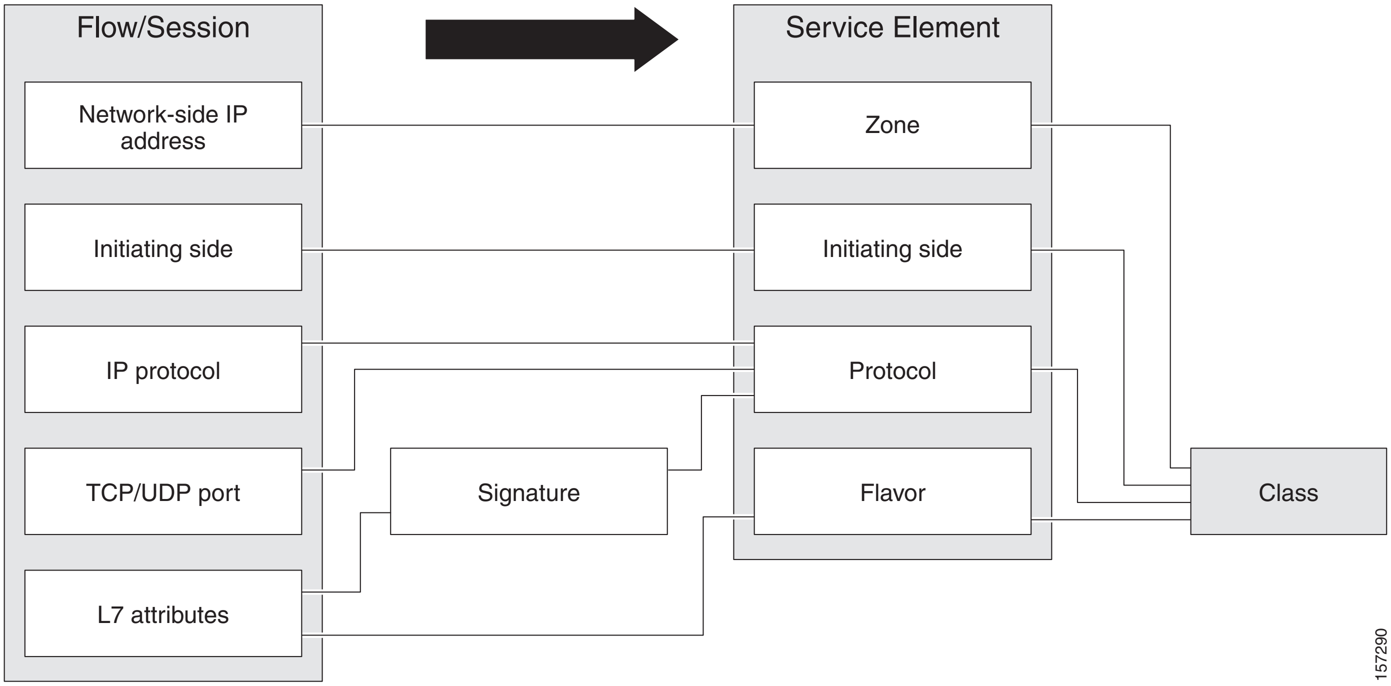

Traffic Monitoring

You can use data gathered by the Application Performance Assurance solutions for reporting.

Various metrics are collected in different scopes—global (per entire link), per class (or group of classes), per policy (or policy profile), and per user—based on user-defined usage counters.

The values from the usage counters can be either pushed or pulled:

- The Application Performance Assurance solution generates and transmits Raw Data Records (RDRs) that contain flow, usage, and other data.
- The Application Performance Assurance solution maintains an SNMP MIB that can be queried by external systems.

Usage Monitoring

The NME-APA collects and maintains various network metrics, per class, in different scopes.

The network metrics are:

- Upstream volume (L3 kilobytes)
- Downstream volume (L3 kilobytes)
- Sessions
- Active users
- Concurrent sessions
- Session duration

Note For VoIP classes, such as SIP and MGCP, the concurrent sessions usage counter counts concurrent voice calls, and the session duration usage counter measures voice call duration.

Per class accounting takes place in the following scopes:

- Per user
- Per group of users (package)
- Per link (global)

Several classes may share the same class usage counter. For example, in the default configuration, the SMTP service and the POP3 service share the E-Mail Counter. The assignment of classes to usage counters is determined by the class hierarchy, as explained in the following section.

The Class Hierarchy

Classes are arranged in a hierarchal tree. A single default class is at the root, and you can place each new class anywhere in the tree.

Classes inherit the matching rules of their parents.

Class Usage Counters

The class hierarchy provides a way to share usage counters and to organize classes according to their semantics. Classes are accounted in groups, as defined in the class hierarchy. Each class is assigned usage counters.

There are two categories of usage counters for classes:

- Global—Used for Link Usage RDRs and reports
- User—Used for Real-Time User Usage RDRs and reports

A global usage counter and a user usage counter are assigned to each class. The use of a class can be accounted either exclusively for traffic classified to it or in conjunction with the traffic of its parent class. For example, if a class called Premium Video Content is defined as a child of Streaming, the operator can either define a special usage counter for Premium Video Content or configure it to use the same usage counter as Streaming. The global usage counter and the user usage counter are independent; for the same class, one usage counter may be the same for parent and child, whereas the other is exclusive to the child.

Reporting

Application Performance Assurance solutions running generate and accumulate Raw Data Records (RDRs) that contain information relevant to the enterprise.

The following are the main categories of RDRs:

- Usage RDRs—Generated periodically. These RDRs contain the state of the usage counters, per class and per accounting scope. There are four types of usage RDRs:
  - Link Usage RDRs—Global usage per class, for the entire link.
  - Usage RDRs—Usage per group of users, per class.

  - User Usage RDRs—Usage per user, per class. These RDRs are generated for all users. The Cisco Application Performance Assurance solution uses these RDRs to generate top-user reports and aggregated usage records.
  - Real-Time User Usage RDRs—Generated for selected users only. The Cisco Application Performance Assurance solution uses these RDRs by to generate detailed user activity reports.
- Transaction RDRs—Generated for a sample of the flows. These RDRs are used to create statistical histograms such as Top TCP Ports.
- Transaction Usage RDRs—Generated for every flow according to user-defined filters. These RDRs contain detailed Layer 7 information for browsing, streaming, and voice flows. They are used for flow-based reporting.

Configurations

A *configuration* implements and enforces the enterprise's business strategy and vision.

A configuration can take effect only after it is propagated to the appropriate NME-APA. The NME-APA enforces the configuration by analyzing the network traffic passing through it.

A configuration consists of:

- Traffic classification settings—Classes, such as web browsing, file sharing, and VoIP. Each class consists of elements that define how network traffic is mapped to the class. The configuration building blocks of classes are protocols, zones, flavors, and signatures.
- Traffic monitoring and reporting settings—Settings that determine how traffic flows and network usage are reported.

In practice, defining configurations is an iterative process.

It is recommended that you use the following sequence of steps:

1. Set up the system.
2. Apply the default configuration.
3. Gather data.
4. Analyze.
5. Continue traffic discovery by partitioning the traffic into (additional) classes.