



Introduction to Traffic Processing

This chapter describes how the Cisco SCA BB installed on a Cisco Service Control Engine (Cisco SCE) platform processes traffic.

The chapter also describes the main elements (service configuration entities) of the Cisco SCA BB system and explains how they relate to each other.

This chapter consists of these sections:

- [Routing Environment](#) , page 1
- [Traffic Processing](#) , page 2
- [Service Security](#) , page 18
- [Traffic Filters](#) , page 20
- [Traffic Forwarding to Value-Added Services Servers](#) , page 20
- [Service Configurations](#) , page 21

Routing Environment

Traffic processing depends on the routing environment. The Cisco Service Control solution can operate in two typical routing schemes:

- **Symmetric (Normal)**—For most flows the inbound and outbound traffic is routed through one Cisco SCE platform. For a marginal number of flows, only one direction goes through this Cisco SCE platform.
- **Asymmetric**—For a significant number of flows, only one direction (inbound or outbound) is routed through the Cisco SCE platform. For other flows, both directions go through this Cisco SCE platform.

A flow is bidirectional when the inbound and outbound traffic of the flow passes through the same Cisco SCE platform. A unidirectional flow is one where only one of the inbound traffic and the outbound traffic go through the Cisco SCE platform.

The Cisco Service Control solution can handle both unidirectional and bidirectional flows. The Cisco SCE platform can be configured to operate in either a symmetric or an asymmetric routing environment. The traffic processing capabilities of the Cisco SCE platform in the asymmetric environment are a subset of its capabilities in the symmetric environment.

When the Cisco Service Control solution is deployed in an asymmetric routing environment, and unidirectional classification is enabled, the Cisco SCE platform classification is better tuned to identify traffic based on a single direction. The Cisco SCE platform handles unidirectional flows independently, with no synchronization with other Cisco SCE platforms that might handle the opposite direction of the flow.

Traffic Processing

There are three stages of traffic processing:

- Traffic classification—Cisco SCA BB analyses traffic flows and determines their type (for example, browsing, e-mail, file sharing, or voice).
- Traffic accounting and reporting—Cisco SCA BB performs bookkeeping and generates Raw Data Records (RDRs) that let you analyze and monitor the network.
- Traffic control—Cisco SCA BB limits and prioritizes traffic flows according to their service, subscriber-package, subscriber quota state, and so on.

You can control how classification, reporting, and control perform by editing the service configurations and by applying these configurations to the Cisco SCE platform.

The three stages are described in these sections:

Traffic Classification

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

For each commercial service that a provider offers to its subscribers, a corresponding service is defined in the Cisco Service Control solution. You can use this service to classify and identify the traffic, report on its usage, and control it.

Cisco SCE internal architecture has two concepts that aid traffic classification:

- Hardware flow—created entirely in hardware, with a maximum limit of 32 million flows on Cisco SCE 10000.
- Software flow—created in software, with a maximum limit of 16 million flows on Cisco SCE 10000.

Each flow context is unidirectional. Flows are opened based on the following logic:

- If the flow is on *filter list* or *traffic rule* with *ignore*, it is ignored and bypassed
- If the packet is Non-IP, it is ignored and bypassed
- If the packet is larger than 1600 bytes, it is ignored and bypassed
- If the packet is a TCP-retransmit packet or has a wrong checksum, it is ignored and bypassed
- If the packet matches any of the active attack filters, it is ignored and bypassed
- If the packet is TCP and the flow is in half-open state (3 way handshake), hardware flow is created for each direction
- If the packet is TCP and is in established state, software flows (2 unidirectional) are created for the first payload packet

- If the packet is UDP, hardware flows are created for first packet in each direction.
- If the packet is UDP, software flow is created for the 5th packet.

Creating flow on the fifth packet helps to avoid creation of software flows for port-scans, and thus, protect Cisco SCE from DoS conditions. Port-scans are still detected because their flows are opened in hardware temporarily. Also, some flows are still opened on the first packet, based on SCA-BB GUI options (Advanced settings).

- If the flow is non-TCP, non-UDP but still IP (for example, ICMP), hardware flow is opened for each direction on first packet
- If the flow is non-TCP, non-UDP but still IP (for example, ICMP), software flow is opened for each direction on second packet

User counters, Service Counters, and Protocol counters are updated, and RDRs are generated only for software flows.

Services

In the traffic classification process, Cisco SCA BB categorizes network sessions into services.

Services are the building blocks for:

- Service configurations (because Cisco SCA BB can enforce different rules on different services)
- Aggregated usage reporting

From the point of view of a provider, a service is a network product sold to a subscriber. The service is usually a network application—such as browsing, e-mail, file sharing, or voice—that the subscriber uses. From a technical point of view, a service consists of one or more service elements, each of which enables a decision about the service associated with a network traffic flow type.

A number of services are predefined in the default service configuration. You can modify these services and add additional services to a service configuration. A service configuration can contain up to 500 services. See the Default Service Configuration Reference Tables chapter of the *Cisco Service Control Application for Broadband Reference Guide* for a list of services.

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

Traffic is classified and mapped to services based on some or all of the following service elements:

- Protocol—The protocol used. This classification allows, for example, the mapping of browsing flows and e-mail flows to separate services.
- Initiating side—Whether the subscriber side or the network side generated the flow. This classification allows, for example, the mapping of subscriber-initiated and network-initiated peer-to-peer traffic to separate services.
- Zone—Lists of IP addresses of the network-side host of the flow. This classification allows, for example, the mapping of all voice flows going to a specified server to a specific service.
- Flavor—Specific Layer 7 properties such as host names of the network-side host of the flow. This classification allows, for example, the mapping of all HTTP flows where the URL matches a certain pattern to a specific service.



Note Flavors are not used for classification when unidirectional classification is enabled.

Cisco SCA BB uses these flow mappings to map each network connection passing through it to a service. You define rules for the different services to implement control policies. The classification rules can contain Layer 3 and Layer 4 parameters (such as port numbers and IP addresses), and also Layer 7 parameters (such as host name and user agent for HTTP connections).



Note Cisco SCA BB cannot achieve 100% classification of all P2P services, because some P2P applications are persistent in trying to connect. They use many alternate protocols and connection schemes. Their native protocol is encrypted and this encryption tends to change whenever a new version is released. This means that if you try to block the P2P traffic, the client may eventually connect in some cases. A better approach may be to limit bandwidth for this traffic to make it ineffective instead of trying for a complete block.

Service Elements

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

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



Note When unidirectional classification is enabled, the flavor of a service element is always the wild-card value.

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

- The flow uses the specified protocol of the service element.
- The flow matches the initiating side specified for the service element.
- The destination of the flow is an address that belongs to the specified zone of the service element.
- The flow matches the specified flavor of the service element.

If a flow matches two service elements and one is more specific than the other, the flow is mapped to the more specific of the two. For example, Service A is defined for browsing and Service B is defined for browsing to a specific list of URLs. A browsing flow to a URL on the list of Service B matches both services, but is mapped to Service B.

If a flow matches one parameter of one service element and a different parameter of another service element, precedence is given first to matching flavors, then to protocols, then to zones, and finally to the initiating side. For example, Service A is defined for e-mail and Service B is defined for all traffic to a specific network zone. An e-mail flow to the specific network zone matches both services, but is mapped to Service A.

Examples of Services

Table 1: Examples of Services and Service Parameters

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

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 Cisco SCA BB 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, Cisco SCA BB maps the flow to the HTTP protocol.

The default service 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, Cisco SCA BB maps the flow to the Generic TCP or Generic UDP protocol.

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

When unidirectional classification is enabled protocol classification is performed in the normal way, with one exception: unidirectional UDP flows. In this case, Cisco SCA BB tries to classify the protocol using the destination port of the first packet. If no exact match is found, Cisco SCA BB tries to classify the protocol using the source port.

Easy Definition of Port-Based Protocols

All generic (unclassified) traffic on a specific port can be assigned to a protocol, by adding the protocol-element in the form <“Generic” signature, specific port> to that protocol. When the “Generic” signature on a specific port is assigned to a protocol, the “Behavioral” signatures are automatically assigned to that protocol as well. For example, in the default configuration, the “Generic” signature on port 555 is assigned to the H20 protocol, and therefore the “Behavioral Upload/Download” signature on port 555 is also automatically assigned to the H20 protocol.

This assignment is done automatically, so you do not need to do the assignment manually. These protocol-elements that are added automatically are not displayed in the GUI. If, on the other hand, you want

to assign the “Behavioral Upload/Download” signature on a specific port to a different protocol, you can do it by creating an appropriate protocol-element and assigning it to the other protocol.


Note

In the default configuration, the HTTP protocol definition accepts not just the HTTP signature, but also all other generic (unclassified) traffic on port 80, by including the protocol-element <“Generic” signature, port 80>. As described previously, when a protocol-element in this form, <“Generic” signature, specific port>, is used in a certain protocol definition, the Cisco SCE maps both the generic and the behavioral signatures, on the specified port, to that protocol. For HTTP traffic, this means that traffic on port 80, which is classified as “Behavioral Upload/Download” signature, would also be assigned to the HTTP protocol. As described earlier, the purpose of this behavior is to allow easy definition of port-based protocols. Nevertheless, this behavior can be avoided, by adding the protocol-element <“Behavioral” signature, specific port> to a different 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.

If a traffic flow meets all the following criteria, it is mapped to a specific protocol:

- 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 is 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 is mapped to Protocol B.

If a flow matches the signature of one protocol element and the port of another protocol element; it is 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 is mapped to Protocol A.

Signatures

Cisco SCA BB examines traffic flows using the deep-packet-inspection capabilities of the Cisco SCE platform, and compares each flow with an installed set of protocol signatures to identify the network application that generated the flow.

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

When unidirectional classification is enabled and a unidirectional flow (inbound or outbound) passes through the Cisco SCE platform, the flow is matched against a special set of unidirectional protocol signatures. When a bidirectional flow passes through the Cisco SCE platform, the protocol library tries to match it to one of its standard (bidirectional) protocol signatures.

Cisco periodically publishes protocol packs containing new signatures and updates to existing signatures. You can use these protocol packs to update the set of signatures installed on Cisco SCA BB, enhancing its classification capabilities.

Dynamic Signatures

Most signatures used by Cisco SCA BB are predefined and hard-coded. Cisco SCA BB also allows you to add dynamic signatures, which can be user-defined.

You can create and edit dynamic signatures in the Signature Editor tool. The Dynamic Signature Script (DSS) engine in Cisco SCA BB carries out the classification using these user-defined signatures in addition to the predefined signatures.

Initiating Side

The Cisco SCE platform is usually located between the subscribers of the provider and the network. Based on the initiating side, flows are called Subscriber-initiated flows and network-initiated flows. Flows initiated by the subscriber towards the network are called subscriber-initiated flow, while the flows initiated from the network towards the subscriber are called network-initiated flows.

You can limit some flow-types to one initiating side. For example, with HTTP you can restrict the direction of the flow to subscriber-initiated, because HTTP is always subscriber-initiated when the subscriber ventures outward to surf the Internet. A network-initiated HTTP-flow means, that probably a web server is open on the local machine of the subscriber for receiving incoming HTTP traffic. The provider can block network-initiated HTTP.

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 network flow of the subscriber mapped to a service 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 service element based on its destination IP address.

Examples of Zones:

- A “walled garden”—A range of IP addresses of a server farm with premium video content, for which the provider would like to limit access to specific subscribers 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.

Zone Items

A zone is a collection of related zone items.

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

Table 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.)

For details on managing zones and zone items, see the “Managing Zones” section.

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 services in the Cisco Service Control solution. A protocol flavor uses an additional protocol attribute in classifying a service, making this service a flavor of the service 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 service.

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



Note

Flavors are not used for traffic classification when unidirectional classification is enabled.

Flavor Items

A flavor is a collection of flavor items.

The type of a flavor item depends on the flavor type. For a list of available flavor types, see [Flavor Types and Parameters](#) section.

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

DSCP ToS

One flavor type is TOS. This allows DSCP ToS to be used as a classification criterion so that a packet carrying a specific marking can be assigned to a predefined service with, for example, unlimited bandwidth or reported. The DSCP ToS classification process takes precedence over other classification mechanisms to allow external devices, such as a voice gateway, to dictate how the flow is treated. DSCP ToS-based classification is an excellent way of marking proprietary managed services where Cisco SCA BB does not recognize the applications but identifies them via the DSCP ToS field.

Content Filtering

Content filtering involves classification and control of HTTP flows according to the requested URL. The classification of the URL is performed by accessing an external database.

Service providers require effective Web filtering for their subscribers, for various purposes such as avoiding litigation and providing parental control. The problem is that the Web is huge and constantly growing, and Cisco SCA BB and the Cisco SCE platform are not designed to track and maintain the huge database of URLs required for effective filtering.

Cisco SCA BB provides content filtering by integrating with SurfControl Content Portal Authority (CPA). SurfControl's technology enhances Cisco SCA BB URL classification capabilities by eliminating the need for a network administrator to manage a URL database or interact with the server, while creating a powerful filtering solution. It provides complete coverage of the web's most trafficked sites and access to the most accurate and relevant database of URLs classified by risk category, such as sexually explicit, racist, hacker, and so on.

The integration of SurfControl's CPA into Cisco SCA BB provides the required web-filtering solution. Cisco SCA BB, running on the Cisco SCE platform, contacts a CPA server to categorize the website that a subscriber requests. The returned category is then used to classify the HTTP flow. This classification is then used for the normal Cisco SCA BB traffic control and reporting.




Note

Cisco SCA BB includes an internal database of URLs used by the HTTP URL flavor classification. When a URL is found in both the internal database and the external content filtering database, the URL is classified according to the internal database.

Flow Attributes to Services Mapping

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

Figure 1: Mapping Flow Elements of a Session to Service Elements of a Service



Traffic Accounting and Reporting

You can use data gathered by the Cisco SCE platforms for real-time signaling, billing, and reporting.

Various metrics are collected in different scopes—global (per entire link), per service (or group of services), per package (or group of packages), and per subscriber—based on user-defined usage counters.

- Global control bandwidth is based on Layer 1 volume.
- Subscriber bandwidth control (and accounting and reporting) is based on Layer 3 volume.

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

- The Cisco SCE platform generates and transmits Raw Data Records (RDRs) that contain flow, usage, and other data.
- The Cisco SCE platform maintains an SNMP MIB that external systems can query.

Usage Accounting

Cisco SCA BB collects and maintains various network metrics, per service, in different scopes.

The network metrics are:

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



Note

For VoIP services, 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 service accounting takes place in the following scopes:

- Per subscriber
- Per group of subscribers (package)
- Per link (global)

Several services may share the same service usage counter. For example, in the default service configuration, the SMTP service and the POP3 service share the E-Mail Counter. The service hierarchy determines how to assign services to usage counters, as explained in the following section. Similarly, several packages may share the same package usage counter, and the package hierarchy determines how to assign packages to usage counters. For details, see [The Package Hierarchy](#) section.

The Service Hierarchy

Services are arranged in a hierarchal tree. A single default service is at the root, and you can place each new service anywhere in the tree. For more information see, [Services](#) section.

Services inherit the rule of their parents. When a rule is defined for a particular service (in a specific package), unless explicitly specified, the same rule of the parent package controls all the child services.

Service Usage Counters

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

There are two categories of usage counters for services:

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

A global usage counter and a subscriber usage counter are assigned to each service. The use of a service can be accounted either exclusively for traffic classified to it or with the traffic of its parent service. For example, if a service 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 subscriber usage counter are independent. For the same service, one usage counter may be the same for parent and child, whereas the other is exclusive to the child.

The Package Hierarchy

Packages are arranged in a hierarchal tree. A single default package is the root of the tree, and you can place new packages anywhere in the tree. For more information see [Packages](#) section.

Package Usage Counters

The package hierarchy allows you to organize packages according to their semantics and provides for sharing package usage counters. You can define a maximum of 1024 different exclusive package usage counters per service configuration, one of which is used for the Unknown Subscriber Traffic package.

Usage reporting at a package level is grouped as follows:

- Package assigned an exclusive package usage counter—All traffic associated with this package is accounted separately in the assigned counter, along with any children that are not assigned exclusive counters.
- Package *not* assigned an exclusive package usage counter—All traffic associated with this package is accounted together with its parent package.

In the figure Example Package Tree, if the Mail & Web Baseline package is allocated an exclusive counter, but neither child package is assigned an exclusive counter, then all Package Usage RDRs and derived reports (such as “Package Bandwidth per Service”) would group usage of subscribers assigned to all three packages.

However, if the Mail & Web Boost package also had an exclusive counter, the traffic for Main & Web Baseline and Mail & Web Captive HTTP would be accounted together, but traffic for Mail & Web Boost would be

accounted separately. (In general, this is not an efficient configuration. You should use the hierarchical structure to group packages that can share the same counter.)

Figure 2: Example Package Tree



Reporting

Cisco SCE platforms running Cisco SCA BB generate and transmit Raw Data Records (RDRs) that contain information relevant to the service provider.

RDRs contain a wide variety of information and statistics, depending on the configuration of the system.

RDRs are transmitted using a Cisco proprietary protocol. To use RDRs, you require the Cisco Service Control Collection Manager (CM) or to develop software to process the RDRs.

The data in some RDRs can also be exported using the NetFlow reporting protocol, which has become an industry standard. NetFlow reporting allows the Cisco SCA BB solution to be more easily integrated with your existing data collectors.

This section contains these topics:

Raw Data Records (RDRs)

The following are the main categories of RDRs:

- Usage RDRs—Generated periodically. These RDRs contain the state of the usage counters, per service and per accounting scope. There are four types of usage RDRs:
 - Link Usage RDRs—Global usage per service, for the entire link.
 - Package Usage RDRs—Usage per group of subscribers, per service.
 - Subscriber Usage RDRs—Usage per subscriber, per service. These RDRs are generated for all subscribers. The Cisco Service Control Collection Manager (CM) and Cisco Service Control Application (SCA) Reporter use these RDRs to generate top-subscriber reports and aggregated usage billing records.
 - Real-Time Subscriber Usage RDRs—Generated for selected subscribers only. The Cisco Service Control Collection Manager and SCA Reporter use these RDRs by to generate detailed subscriber 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 billing.
- Real-Time Signaling RDRs—Generated to indicate specific network events such as flow start or end. These RDRs are used to signal external systems to allow real-time actions across the network.
- Malicious Traffic RDRs—Generated to indicate that the Cisco SCE platform has detected a traffic anomaly, such as a DDoS attack. These RDRs are used to detect attacks and attackers to mitigate them.

NetFlow

The following information can be exported using the NetFlow protocol

- Usage—Generated periodically. These RDRs contain the state of the usage counters, per service and per accounting scope.
- Malicious Traffic—Generated to indicate that the Cisco SCE platform has detected a traffic anomaly, such as a DDoS attack.

Traffic Control

Traffic Control provides means to block, limit, or prioritize traffic flows according to service, subscriber package, subscriber quota state, and so on.

Packages

A package is a collection of rules describing subscriber policy. The package defines the group of services delivered to a specific group of subscribers and the behavior of the system for each service. It may contain restrictions on network flows, guidelines for prioritization of the flows, and instructions about how to report flows.

Each subscriber in the network is provided with a reference to a package to which that subscriber belongs. The following list describes how the system references each subscriber in the network:

- 1 Maps each network flow to a service by matching the flow with a service element
- 2 Identifies the subscriber to whom the flow pertains, according to the network ID of the subscriber (usually the IP address of the subscriber)
- 3 Identifies the package to which the subscriber belongs
- 4 Applies the correct rule to the service of the network flow of the subscriber

Another scheme is described in the following section:

Virtual Links Mode

In normal mode, you define bandwidth controllers for each package (see [Bandwidth Management](#)). In Virtual Links mode, you define template bandwidth controllers. The actual bandwidth parameters are assigned to a

subscriber when the subscriber enters the system. These parameters depend on the package of the subscriber and the direction of the virtual link.

For more information, see [Quota Management section](#).

Unknown Subscriber Traffic

The Cisco SCE platform tries to identify the subscriber responsible for every traffic flow that it processes. The platform looks at the IP address or VLAN tag of the traffic flow, and checks its internal database for a subscriber identified by this IP Address or VLAN tag. If such a subscriber is not found in the database, the traffic flow is mapped to the Unknown Subscriber Traffic category.

Rules

A rule is a set of instructions that tell the Cisco SCE platform how to treat network flows of a specific service. A rule can:

- Specify that a flow should:
 - be blocked
 - be granted a certain amount of bandwidth
 - have the DSCP ToS of its packets marked with a given value (see [DSCP ToS Marking](#))
 - Define an aggregate volume or session limit, after which a set of different restrictions are enforced on the flow
 - Specify how a flow is reported for billing or analysis purposes

Calendars

You can use calendars to divide the hours of the week into four time frames.

After you have configured a calendar, you can add Time-Based Rules to a package that uses the calendar.

Time-Based Rules

A time-based rule is a rule that applies to only one time frame. Time-based rules allow you to set rule parameters that are only applied at specific times. You might, for example, want to define different rules for peak, off-peak, nighttime, and weekend usage.

You can add time-based rules to any rule. If a time-based rule is not defined for a time frame, the parent rule is enforced.

Often, you need rules for different time frames to be similar. When you add a time-based rule, the settings of the parent rule are copied to the new time-based rule; you can make any needed changes. Subsequent changes to the parent rule do not affect the time-based rule.

Related Topics

[Global Bandwidth Control](#) , on page 15

Bandwidth Management

The physical link bandwidth is an absolute limit on the bandwidth that can pass through the system. You can limit the total bandwidth passing through the Cisco SCE platform to a value lower than the physical link bandwidth. For example, if another device connected to the Cisco SCE platform on the IP stream has limited BW capacity, you can limit the bandwidth passing through the Cisco SCE platform to match the capacity of the other device.

Bandwidth control in Cisco SCA BB is accomplished in two stages:

- Global control— based on Layer 1 volume.
- Subscriber bandwidth control—and accounting and reporting is based on Layer 3 volume.

Global Bandwidth Control

Global controllers control the total bandwidth use. Global controllers are virtual queues in Cisco SCE platforms. You configure them for the entire system, rather than for individual subscribers.

Global controllers provide constraints for large, global volumes of traffic, such as “Total Gold Subscriber Traffic”, or “Total P2P Traffic”. Each global controller defines the maximum percentage of total available bandwidth allocated to all traffic of a particular type. Using a global controller, you can limit total traffic of services such as P2P in the system to any bandwidth between 16 kb/s and 1000 Mb/s. In this way, you keep the total bandwidth consumed by this traffic under control.

The upstream and downstream interfaces are each assigned one default global controller that, by default, controls 100 percent of the link traffic. You can add up to 1023 more global controllers for each interface on Cisco SCE Gigabit Ethernet hardware and up to 4095 more global controllers on Cisco SCE 10 Gigabit Ethernet hardware and you can assign a maximum percentage of the total link limit to each global controller separately.

For each global controller, you can define separate values for the maximum percentage of total available bandwidth separately for each time frame.

In dual-link systems, you can define different bandwidth values for each link. You can also set a limit on the aggregated bandwidth passing on the two links.

Virtual Links mode uses template global controllers. Template global controllers are templates of virtual queues; they are applied to as many separate physical links as exist in the system. For each physical link, actual bandwidth parameters depend on the link. (For more information, see [“Quota Management” section](#).)

Related Topics

[Calendars](#) , on page 14

[Quota Management](#)

Subscriber Bandwidth Control

Subscriber BW Controllers (BWCs) controls the bandwidth used by individual subscribers.

Each BWC controls available bandwidth for selected services. Services controlled by a particular BWC are defined per package, but bandwidth control is per service.

The following parameters specify a BWC:


- Committed Information Rate (CIR)—The minimum bandwidth that must be granted to the services that a BWC control.
- Peak Information Rate (PIR)—The maximum bandwidth that can be allocated to the services that a BWC control.
- Global Controller—The global controller to which this BWC links
- Assurance Level (AL)—The rate of change of available bandwidth under conditions of traffic congestion

The Bandwidth Control Levels figure illustrates the maximum available bandwidth (Admitted Information Rate [AIR]) ranges between the CIR and the PIR. The actual consumed bandwidth is always less than the AIR.

The BWC has a third parameter that controls how the AIR is determined at different congestion conditions. When the network is not congested the system allows the PIR and when the network is highly congested the system provides the CIR. In between these two extremes, a third parameter—Assurance Level (AL)—determines the AIR. The AL controls how fast the AIR would decrease from the PIR to the CIR as congestion builds, or increase from the CIR to the PIR as congestion decreases. A higher AL ensures a higher AIR compared to a similar BWC with a lower AL.

The BWC ensures that even when the network is congested (PIR-congestion) at least the CIR is granted. Similarly, the BWC ensures that even when there is little traffic associated with a BWC the PIR is not exceeded.

Figure 3: Bandwidth Control Levels




Bandwidth may be thought of in terms of a virtual pipe of adjustable width. The PIR is the maximum allowed width of the virtual pipe. The CIR is the minimum width to which the pipe can contract. The actual pipe width is the AIR. During Network congestion, the system contracts each pipe differently to differentiate between subscribers and between their services.

Primary and Internal Bandwidth Control

In Cisco SCA BB each subscriber has an independent set of BWCs, consisting of a single Primary (Total) BWC (tBWC) that controls the total bandwidth available to the subscriber and several Internal BWCs (iBWCs)

that control the available bandwidth of some services of that subscriber, as illustrated in Figure. For example, one BWC may control the Streaming Service; another may control the Download and E-mail Services together. The PIR defines the maximum bandwidth for the associated services; the CIR defines the minimum bandwidth for them.

Figure 4: Bandwidth Control on Two Levels



You can link iBWCs to traffic in the following way:

- 1 In the package general definitions, add a subscriber BWC, defined by its CIR, PIR, AL, and CoS.
- 2 When defining a rule, assign each service to one subscriber BWC.

Quota Management

You can assign subscribers a quota limit on selected services.

Each subscriber has 16 quota buckets, each of which you can define for volume or sessions. When a subscriber uses a certain service, the amount of consumed volume or number of sessions is subtracted from one of the buckets.

The service configuration determines which bucket to use for each service. Consumption of volume buckets is measured in units of L3 kilobytes. Consumption of session buckets is measured by the number of sessions. For example, you can define that the Browsing and E-Mail services consume quota from Bucket #1, that the P2P service consumes quota from Bucket #2, and that all other services are not bound to any particular bucket.

External quota provisioning systems can use the Quota Provisioning API to modify the quota in each bucket dynamically. For example, you can increase the quota of a certain bucket when a subscriber purchases additional quota. These external systems can also query the amount of remaining quota in each bucket. This can be used, for example, to show subscribers in a personal web page how much of their quota remains. For details on Quota Provisioning API, see the *Cisco Service Control SCE Subscriber API Programmer's Guide*.

External quota provisioning can also be acquired using the Quota Manager (QM), an off-the-shelf solution provided by Cisco. For more information about the installation and operation of the QM, see the *Cisco Service Control Management Suite Quota Manager User Guide*.

External quota provisioning can also be acquired using the Gy quota model and Gx quota model. For more information, see the *Cisco Service Control Mobile Solution Guide*.



Note External quota provisioning is not supported when unidirectional classification is enabled.

The internal Cisco SCA BB quota provisioning system replenishes each quota bucket by a fixed amount at fixed intervals.

Subscribers can be notified when they breach the quota in any bucket.

Subscriber Notification

The subscriber notification feature lets you push web-based messages (such as notifications of quota depletion) to a subscriber by redirecting the subscriber HTTP traffic to relevant web pages. HTTP redirection starts when the subscriber notification is activated and ceases when the notification is dismissed.



Note Subscriber notification is not supported when unidirectional classification is enabled.

Service Security

Cisco SCA BB includes service security functionality to help protect network operators and their subscribers from attacks and malicious traffic:

- DoS attacks
- DDoS attacks
- VoIP threats
- Worms
- Hacker activity
- Malicious takeover of subscriber computers:
 - Spam zombies
 - E-mail based viruses

Although it is never possible to provide complete protection from network threats, the Cisco Service Control solution provides insight into malicious activity in a network, and can mitigate large-scale eruptions of malicious activity that compromise overall network performance.

Networks operators can use Cisco SCA BB to:

- Monitor network traffic for suspicious activity
- Block malicious traffic

- Notify subscribers that are creating or have been affected by malicious traffic

Detecting Malicious Traffic

Cisco SCA BB uses four threat detection mechanisms:

- Anomaly Detection—This set of mechanisms monitors the rate of connections (both successful and unsuccessful) to and from each host IP address. High connection rates or a low ratio between successful and unsuccessful connections indicate malicious activity.

Anomaly detection characteristics can indicate the following categories of malicious activity:

- IP sweep—Scanning multiple IP addresses, all on the same port (a behavior typical of worms)
- Port scan—Scanning all ports at one IP address (a behavior typical of hackers)
- DoS attack—An attack (on a single IP address) from a single IP address
- DDoS attack—An attack (on a single IP address) from multiple IP addresses



Note

Cisco SCA BB identifies a DoS attack with spoofing (using many fake IP addresses instead of one real address) as a DDoS attack.

The anomaly detection mechanism is effective in addressing new threats as they appear. It does not need knowledge about their exact nature and Layer 7 signatures, but is based on the characteristics of their network activity.

- Mass mailing activity detection—This mechanism monitors SMTP session rates for individual subscribers (using Cisco SCE platform subscriber-awareness; it can work in subscriber-aware or anonymous subscriber mode). A high rate of SMTP sessions from an individual subscriber is usually an indicator of malicious activity that involves sending e-mail (either mail-based viruses or spam-zombie activity).
- Signature-based detection—The stateful Layer 7 capabilities of the Cisco SCE platform are used to detect malicious activity that is not easily detectable by the other mechanisms. Operators can add signatures for such threats, achieving a quick response time in addressing new threats.
- RFC compliance detection—This mechanism monitors the SMTP traffic for RFC compliance. Non-compliant traffic is marked as spam.

Responding to Malicious Traffic

You can define the following actions when configuring the detection mechanisms described in the preceding section:

- Monitor the network for malicious activity detected by each of these mechanisms.
You can display graphs in the Console based on data collected for malicious activity analysis.
- Automatically block malicious activity detected by the Cisco SCE platform to avoid threat propagation and adverse effects to the network.
- Notify subscribers that are involved in malicious activity by redirecting their web sessions to a captive portal.

Cisco SCA BB provides a high level of flexibility in tuning the detection methods to define malicious activity and in configuring the actions to be taken when malicious activity is detected.

Traffic Filters

Filter rules are part of service configurations. Filter rules allow you to instruct the Cisco SCE platform to ignore some types of flow (based on the Layer 3 and Layer 4 properties of the flow) and to transmit the flows unchanged.

When a traffic flow enters the Cisco SCE platform, the platform checks whether a filter rule applies to the flow. If a filter rule applies to this traffic flow, the Cisco SCE platform performs one of the following actions:

- **Bypass**—The Cisco SCE platform passes the traffic flow to its transmit queues without generating any RDRs (the flow does not appear in records generated for analysis purposes) and without enforcing any service configuration rules.
- **Quick forward**—A flow filter rule action whose aim is to ensure low latency for delay sensitive flows. The packets of quick-forwarded flows are duplicated and sent through different paths: one copy goes directly to the transmit queue and thus suffers only a minimal delay, the other copy goes through the normal packet path.

A filter rule can also set the DSCP ToS value of the filtered traffic.

It is recommended that you add filter rules for OSS protocols (such as DHCP) and routing protocols (such as BGP) that might traverse the Cisco SCE platform. These protocols usually should not be affected by policy enforcement, and their low volume makes them insignificant for reporting.

A number of filter rules are included in the default service configuration.

There are 64 flow filter rules for both IPv4 and IPv6 addresses on the Cisco SCE 10000 devices. You can add from 0 to 60 flow filter rules using the Cisco SCA BB console and the other three flow filter rules are reserved.

Flows of certain protocols can also be filtered according to the Layer 7 characteristics of the flow.

DSCP ToS Marking

DSCP ToS marking is used in IP networks to signal the type and priority of a flow between network elements. Typically, those elements that have an insight on how to treat the traffic throughout the network performs the DSCP ToS marking. Such an element can be the element generating traffic—a voice gateway, for example. Cisco SCA BB, being application aware, can, for example, allocate bandwidth resources based on the business model and the specific needs of latency sensitive applications.

Traffic Forwarding to Value-Added Services Servers

Traffic forwarding to Value Added Services (VAS) servers allows the Cisco Service Control solution to use an external expert system (VAS server) for additional traffic processing. The Cisco SCE reroutes traffic to the preconfigured location of the VAS server. After processing, the traffic is sent back to the Cisco SCE, which then sends it to its original destination.



Note VAS traffic forwarding is not supported when unidirectional classification is enabled.

Service Configurations

A service configuration implements and enforces the business strategy and vision of the provider.

A service configuration can take effect only after it is propagated to the appropriate Cisco SCE platform. Cisco SCA BB enforces the service configuration by analyzing the network traffic passing through them.

A service configuration consists of:

- Traffic classification settings—Services, such as web browsing, file sharing, and VoIP. Each service consists of elements that define how network traffic is mapped to the service. The configuration building blocks of services are protocols, zones, flavors, and signatures.
- Traffic accounting and reporting settings—Settings that determine how traffic flows and network usage accounting are reported.
- Traffic control settings—Packages, which consist of a set of rules (such as bandwidth rate limit and quota limits) defined for different services. The main configuration building blocks of packages are rules, quota buckets, subscriber BWCs, and global controllers.

Defining Service Configurations in Practice

In practice, defining service configurations is an iterative process. It is recommended that you use the following sequence of steps:

Procedure

- Step 1** Set up the system.
 - Step 2** Apply the default service configuration.
 - Step 3** Gather data.
 - Step 4** Analyze.
 - Step 5** Do one or both of the following:
 - Continue traffic discovery by partitioning the traffic into (additional) services.
 - Create rules to limit and prioritize traffic according to services and subscriber packages.
-

