Filter Policies

In This Chapter

The SROS supports filter policies for services and network interfaces (described in this chapter), subscriber management (integrated with service filter policies with the subscriber management specifics defined in the SROS Triple Play Guide), and CPM security and Management Interface (described in SROS Router Configuration Guide).

Topics in this chapter include:

ACL Filter Policy Overview

ACL Filter policies, also referred to as Access Control Lists (ACLs) or filter for short, are sets of ordered rules specifying packet match criteria and actions to be performed upon a match. Filters are applied to services or network ports to control network traffic into (ingress) or out of (egress) a service access port (SAP) or network. There are three main types of filter policies: IPv4, IPv6, and MAC filter policies. The same filter can be applied to ingress traffic, egress traffic, or both. Ingress filters affect only inbound traffic destined for the routing complex, and egress filters affect only outbound traffic sent from the routing complex.

Configuring an entity with a filter policy is optional. By default, there are no filters associated with services or interfaces, and therefore, all traffic is allowed on the ingress and egress interfaces. The filter must be explicitly created and associated. There are different types of filter policies as defined by the scope argument of the filter policy. An exclusive filter is intended to be used by a single SAP/interface, a template filter is intended to be shared by multiple SAP/interfaces in the system, and an embedded filter is intended to define common filter rules that can then be used (embedded) by other filters in the system. Filter policies are created with a unique filter ID, but each filter has also a unique filter name argument that can be defined once the filter policy has been created. Either filter ID or filter name can then be used throughout the system to manage filter policies and their associations.

On a Layer 2 SAP, either a single IP (v4 or v6) or a single MAC filter policy can be applied in a given direction. On a Layer 3 SAP and network interfaces, a single IP (v4 or v6) can be applied in a given direction. The ingress and egress direction policies can be same or different. For dual stack IPv4/IPv6 SAPs/interfaces, if both IPv4 and IPv6 filter policies are defined, the policy applied will be based on the outer IP header of the packet. Note that non-IP packets are not hitting an IP filter policy, so the default action in the IP filter policy will not apply to these packets.

Filter Policy Entities

A filter policy is applied to packets coming through the system, in the ascending order the entries are numbered in the policy. When a packet matches all the parameters specified in a filter entry's match criteria, the system takes the specified action defined in that entry. If a packet does not match the entry parameters, the packet is compared to the next higher numerical filter entry, and so on. If the packet does not match any of the entries, the system executes the default action specified in the filter policy. Each filter policy is assigned a unique filter ID. Each filter policy is defined with:

- Scope
- Default action
- Description
- Filter Name that can be optionally used in CLI to reference this filter policy instead of Filter ID (some exceptions apply)
- At least one filter entry

Each filter entry contains:

- Match criteria
- An action
- In addition, in a filter policy entry, an operator can also:
  - ∅ configure log ID to enable filter logging for this entry.
  - ∅ control how cflowd sampling is done for an IP interface based on IP interface cflowd configuration and the filter entry cflowd configuration.

Applying Filter Policies

Filter policies can be associated with the following entities:

Table 8: Applying Filter Policies

IPv4 Filter	MAC Filter	IPv6 Filter
Security CPM filter	Security CPM filter	Security CPM filter
CRON TOD-suite	CRON TOD-suite	CRON TOD-suite
Router interface	N/A	Router interface
Egress multicast group	Egress multicast group	Egress multicast group
IES interface SAP, spoke SDP	N/A	IES interface SAP, spoke SPD subscriber-interface
VPRN interface SAP, spoke SDP	N/A	VPRN interface SAP, spoke SDP
VPLS mesh/spoke SDP, SAP	VPLS mesh/spoke SDP, SAP	VPLS mesh/spoke SDP, SAP
Epipe SAP, spoke SDP	Epipe SAP, spoke SDP	Epipe SAP, spoke SDP
Fpipe SAP, spoke SDP	Fpipe SAP, spoke SDP	Fpipe SAP, spoke SDP
Ipipe SAP, spoke SDP	Ipipe SAP, spoke SDP	Ipipe SAP, spoke SDP
Pseudowire template	Pseudowire template	Pseudowire template

-

ACL Filter Policy Scale

Release 11R4 introduces an enhanced flexibility in defining per service or per customer filter policies across services and interfaces that the router supports.

Prior to release 11.0, the number of filter policies supported in the system was equal to the number of filter policies supported by a single FlexPath on a line card. All policies would be downloaded to all line cards, regardless whether a policy was needed by a line card or not. Starting with Release 11.0R1, the number of filter policies that can be configured in a single system is now greater than the number of filter policies that can be used on a single FlexPath forwarding complex. The operator can manage the standard filter policies at a system-level (with system-wide policy identifiers) and SR-OS automatically maps and downloads policies to each FlexPath only as needed by services and interfaces configured on that FlexPath.

Statistics for filters aggregate all statistics across all FlexPaths that have a given filter entry active and will show zero (0) if a filter entry is not downloaded to any line card. The statistics are also reset to zero (0) when a given filter is removed from one of the line cards. When a filter is downloaded to a new line card as result of another service using that filter, the statistics continue incrementing. Figure 15 shows the new model for IPv4 ACL filter policies.

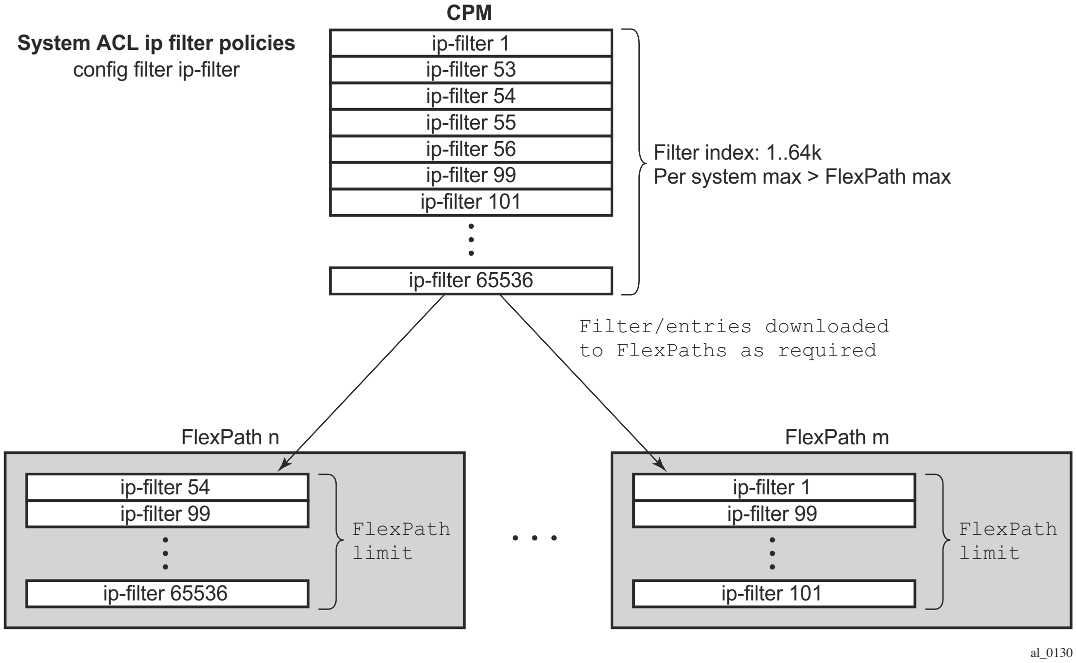

Figure 15: FlexPath ACL IPv4 Filter Policy Management Example

Assignment of filter policies to Interfaces, SAPs and SDPs is allowed up to the maximum number of filter policies supported per FlexPath (unchanged). If a maximum supported on a given FlexPath is breached, the configuration change to a filter policy is blocked. Due to a co-existence of dynamic filter policy entries in the system, an operator-configured filter policy may still fail to be installed in hardware later on. If that is the case, a trap will be raised for the impacted filter policies. It is recommended that the operator remove extra filter entries as operational conditions, such as an IOM reset for example may cause different filter entries to be activated when FlexPath limits are exceeded.

Note that a filter policy applied to a spoke, or to a SAP/Interface that resides on multiple line cards (for example, when a LAG or G.8032 rings with links on multiple cards are used), will be downloaded to all FlexPaths on all cards on which the SAP/interface resides.

Since only the active filter policies are downloaded to a given line card, counters for filter entries are available only for those filters that are downloaded to one or more line cards.

Match-list for Filter Policies

Figure 16 depicts an approach to implement logical OR on a list of matching criterion (IPv4 address prefixes in this example) in one or more filter policies prior to introduction of match list.

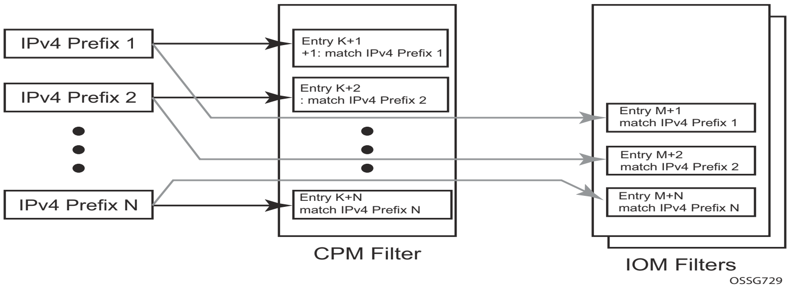

Figure 16: IOM/CPM Filter Policy using Individual Address Prefixes

An operator has to create one entry for each address prefix to execute a common action. Each entry defines a match on a unique address prefix from the list plus any other additional match criteria and the common action. If the same set of address prefixes needs to be used in another IOM or CPM filter policy, an operator again needs to create one entry for each address prefix of the list in those filter policies. Same procedure applies (not shown above) if another action needs to be performed on the list of the addresses within the same filter policy (when for example specifying different additional match criteria). This process can introduce large operational overhead, especially when a list contains many elements or/and needs to be reused multiple times across one or more filter policies.

Match list for CPM and IOM filter policies are introduced to eliminate above operational complexity by simplifying the IOM and CPM filter policy management on a list of a match criterion. Instead of defining multiple filter entries in any given filter, an operator can now group same type of the matching criteria into a single filter match list, and then use that list as a match criterion value, thus requiring only single filter policy entry per each unique action. The same match list can be used in one or more IOM filter policies as well as CPM filter policies.

The match lists further simplify management and deployment of the policy changes. A change in a match-list content is automatically propagated across all policies employing that list in their match criteria, thus only a single configuration change is required to trigger policy changes when a list is used by multiple entries in one or more filter policies.

Figure 17 depicts how the IOM/CPM filter policy illustrated at the top of this section changes with a filter match list usage (using IPv4 address prefix list in this example).

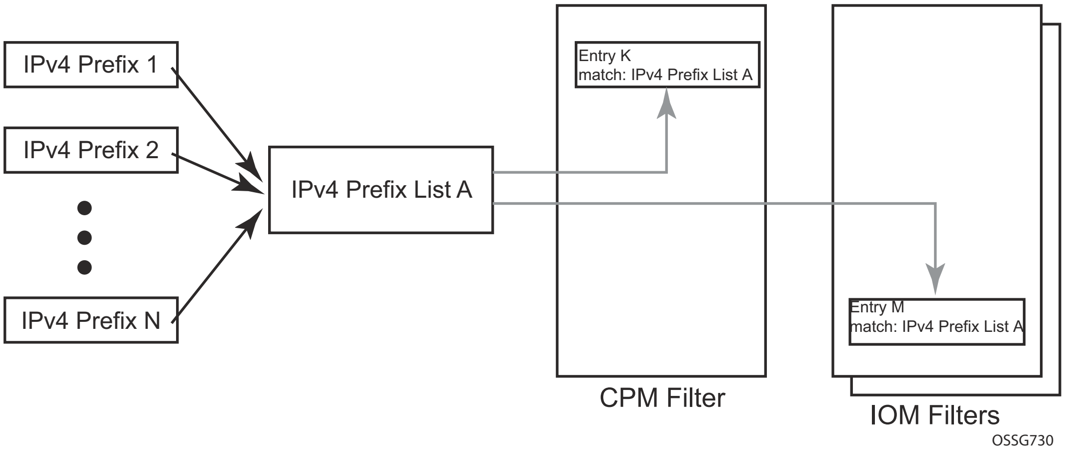

Figure 17: IOM/CPM Filter Policy Using an Address Prefix Match List

Note: The hardware resource usage does not change whether filter match lists are used or whether operator creates multiple entries (each per one element of the list): however, a careful consideration must be given to how the lists are used to ensure only desired match permutations are created in a filter policy entry (especially when other matching criteria that are also lists or ranges are specified in the same entry). The system verifies that a new list element, for example, an IP address prefix, cannot be added to a given list or a list cannot be used by a new filter policy unless resources exist in hardware to implement the required filter policy (ies) that reference that list. If that is not the case, addition of a new element to the list or use of the list by another policy will fail.

Some use cases like those driven by dynamic policy changes, may result in acceptance of filter policy configuration changes that cannot be programmed in hardware because of the resource exhaustion. If that is the case, when attempting to program a filter entry that uses a match list(s), the operation will fail, the entry will be not programmed, and a notification of that failure will be provided to an operator.

Please refer to SROS Release Notes for what objects can be grouped into a filter match list for IOM and CPM filter policies.

Auto-generation of Filter-policy Address Prefix Match Lists

It is often desired to automatically update a filter policy when the configuration on a router changes. To allow such a touch-less filter policy management, SROS allows auto-generation of address prefixes for IPv4 or IPv6 address prefix match lists based on operator-configured criteria. When the configuration on a router changes, the match lists address prefixes are automatically updated and, in-turn, all filter policies (CPM or IOM) that use these match lists are automatically updated.

When using auto-generation of address prefixes inside an address prefix match list operators can:

- Specify one or more *regex* expression matches against SROS router configuration per list.
- Specify wildcard matches by specifying *regex* wildcard match expression (“.*”).
- Mix auto-generated entries with statically configured entries within a match list.

The following additional rules apply to auto-generated entries:

- Operational and administrative states of a given router configuration are ignored when auto-generating address prefixes.
- Duplicates are not removed when populated by different auto-generation matches and static configuration.
- A configuration will fail if auto-generation of address prefix would result in filer policy resource exhaustion on a filter entry, system, or line-card level.

NOTE: See Release notes and CLI section for details on what configuration supports address prefix list auto-generation.

The following may apply to this feature:

If filter policy resources are not available for newly auto-generated address prefixes when a BGP configuration changes, new address-prefixes will not be added to impacted match lists or filter policies as applicable. An operator must free resources and change filter policy configuration or must change BGP configuration to recover from this failure.

Embedded Filter Support for ACL Filter Policies

When a large number of standard filter policies are configured in a system, a set of policies will often contain one or more common blocks of entries that define, for example, system-wide and/or service-wide security rules. Prior to introduction of the embedded filters, such common rules would have to be configured separately in each exclusive/template policy.

To simplify management of such common rules across multiple filter policies, operator can now use embedded filter policies. An embedded filter policy is a special type of a filter policy that cannot be deployed directly but instead is used to define a common filter policy rules that are then included in (embedded by) other filter policies in the system. Thanks to embedding, a common set of rules can now be defined and changed in a single place but deployed across multiple filter policies. The following main rules apply when embedding an embedded filter policy:

1. An operator can explicitly define an offset at which to embed a given embedded filter into a given embedding filter—the embedded filter entry number X becomes an entry (X + offset) in the embedding filter.
2. An exclusive/template filter policy may embed multiple embedded filter policies as long as the embedded entries do not overlap.
3. A single embedded filter policy may be embedded in many exclusive/template filter policies.
4. When embedding an embedded filter, an operator may wish to change or deactivate an embedded filter policy entry in one of the embedding filter, thus allowing for customizing of the common embedded filter policy rules by the embedding filter. This can be achieved by either defining an entry in the embedding filter that will match ahead of the embedded filter entry or by overwriting the embedded filter entry in the embedding filter.

   For example: If embedded filter 99 has entry 20 that drops packets that match IP source address **src_address**, and filter 200 embeds filter 99 at offset 100, then to *deactivate* the embedded entry 20, an operator could define an entry 120 (embedded entry number 20 + offset 100) in filter policy 200, that has the same match criteria and has either no action defined (this will deactivate the embedded entry and allow continued evaluation of filter policy 200), or has action forward defined (packets will match the new entry and will be forwarded instead of dropped, evaluation of filter policy 200 will stop).
5. Any embedded policy rule edits are automatically applied to all filter policies that embed that embedded filter policy.
6. The system verifies whether system and h/w resources exist when a new embedded filter policy is created, changed or embedded. If resources are not available, the configuration is rejected. In rare cases, filter policy resource check may pass but filter policy can still fail to load due to a resource exhaustion on a line card (for example when other filter policy entries are dynamically configured by applications like RADIUS in parallel). If that is the case, the embedded filter policy configured will be de-activated (configuration will be changed from **activate** to **inactivate**).

7. An embedded filter is never embedded partially into an exclusive/template filter; that is, resources must exist to embed all embedded filter entries in a given exclusive/template filter. Although a partial embedding into a single filter will not take place, an embedded filter may be embedded only in a subset of embedding filters (only those where there are sufficient resources available).

Figure 18 shows implementation of embedded filter policy using IPv4 ACL filter policy example with an embedded filter 10 being used to define common filter rules that are then embedded into filter 1 and 20 (with filter 20 overwriting rule at offset 50):

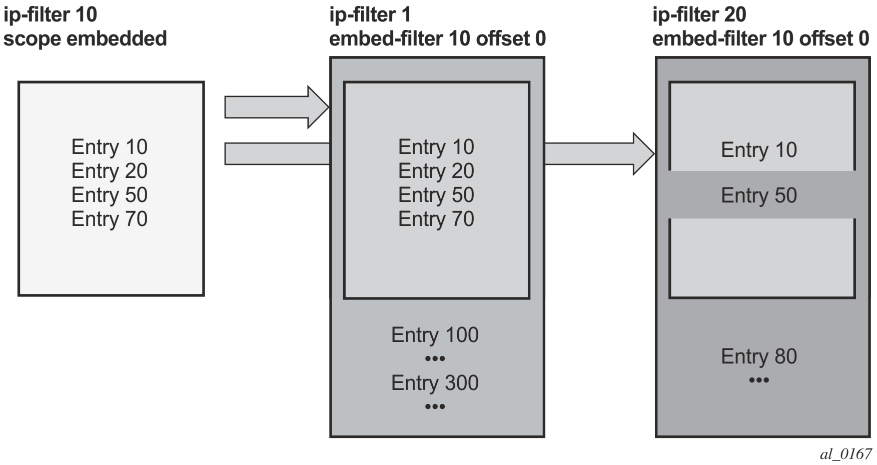

Figure 18: Embedded Filter Policy

NOTE: Embedded filter policies are supported for line card IP(v4) and IPv6 filter policies only.

Redirect Policies

SROS-based routers support redirect policies. Redirection policies are used to identify cache servers (or other redirection target destinations) and define health check test methods used to validate the ability for the destination to receive redirected traffic. This destination monitoring greatly diminishes the likelihood of a destination receiving packets it cannot process.

Redirection identifies packets to be redirected and specifies the method to reach the web cache server. Packets are identified by IPv4 filter entries. The redirection action is accomplished and supported with Policy Based Routing. Only IPv4 routed frames can be redirected. Bridged IP packets that match the entry criteria will not be redirected.

Redirection policies can contain multiple destinations. Each destination is assigned an initial or base priority describing its relative importance within the policy. The destination with the highest priority value is selected.

There are no default redirect policies. Each redirect policy must be explicitly configured and specified in an IPv4 filter entry.

To facilitate redirection based on a redirection policy, an IPv4 filter must be created and applied to the appropriate ingress IP interfaces where redirection is required. The entry criteria for the filter entry must specify a redirect policy to enable the appropriate IPv4 packets to be redirected from the normal IPv4 routing next hop. If packets do not meet any of the defined match criteria, then those packets are routed normally through the destination-based routing process.

The redirection policy is referenced within the action context for an IPv4 filter entry, binding the filter entry to the policy and the IPv4 destinations managed by the policy. The policy specifies the destination IPv4 address where the packets matching the filter entry will be redirected. When the policy determines the destination for packets matching the filter, the action on the filter entry is similar to provisioning that destination IPv4 address as an indirect next hop Policy Based Route (PBR) action.

Web Redirection (Captive Portal)

Web redirection policies can be configured on 7750 SR devices. Redirection policies were designed for testing purposes. The new redirection policy can now block a customer's request from an intended recipient and force the customer to connect to the service's portal server. 255 unique entries with http-redirect are allowed.

Traffic Flow

The following example provides a brief scenario of a customer connection with web redirection.

1. The customer gets an IP address using DHCP (if the customer is trying to set a static IP he will be blocked by the anti-spoofing filter).
2. The customer tries to connect to a website.
3. The router intercepts the HTTP GET request and blocks it from the network
4. The router then sends the customer an HTTP 302 (service temporarily unavailable/moved). The target URL should then include the customer's IP and MAC addresses as part of the portal's URL.
5. The customer's web browser will then close the original connection and open a new connection to the web portal.
6. The web portal updates the ACL (directly or through SSC) to remove the redirection policy.
7. The customer connects to the original site.

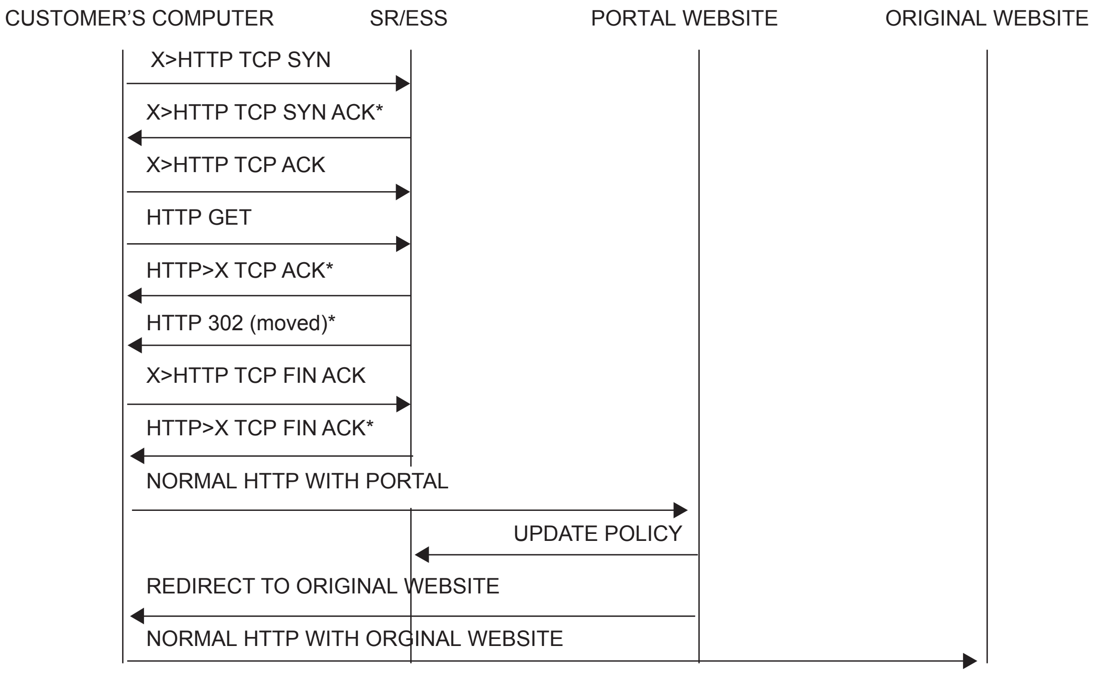

Figure 19: Web Redirect Traffic Flow

Starred entries (*) are items the router performs masquerading as the destination, regardless of the destination IP address or type of service.

The following displays information that can optionally be added as variables in the portal URL (http-redirect url):

- $IP – Customer's IP address
- $MAC – Customer's MAC address
- $URL – Original requested URL
- $SAP – Customer's SAP
- $SUB – Customer's subscriber identification string"

Note that the subscriber identification string is available only when used with subscriber management. Refer to the subscriber management section of the SROS Triple Play Guide and the SR OS Router Configuration Guide.

Since most web sites are accessed using the domain name the router allows either DNS queries or responds to DNS with the portal's IP address.

ISID Filters

ISID filters are a type of MAC filters that allows filtering based on the ISID values rather than L2 criteria used by MAC filters of type "normal" or "vid". ISID filters can be deployed on iVPLS PBB SAPs and ePipe PBB SAPs in the following scenarios:

The MMRP usage of the mrp-policy ensures automatically that traffic using Group BMAC is not flooded between domains. However; there could be a small transitory periods when traffic originated from PBB BEB with unicast BMAC destination may be flooded in the BVPLS context as unknown unicast in the BVPLS context for both IVPLS and PBB Epipe. To restrict distribution of this traffic for local PBB services ISID filters can be deployed. The mac-filter configured with ISID match criterion can be applied to the same interconnect endpoint(s), BVPLS SAP or PW, as the mrp-policy to restrict the egress transmission any type of frames that contain a local ISID. The ISID filters will be applied as required on a per B-SAP or B-PW basis just in the egress direction.

The ISID match criteria are exclusive with any other criteria under mac-filter. A new mac-filter type attribute is defined to control the use of ISID match criteria and must be set to ISID to allow the use of ISID match criteria.

VID Filters

VID Filters are a type of MAC filters that extend the capability of current Ethernet Ports with null or default SAP tag configuration to match and take action on VID tags. Service delimiting tags (for example QinQ 1/1/1:10.20 or dot1q 1/1/1:10, where outer tag 10 and inner tags 20 are service delimiting) allow fine grain control of frame operations based on the VID tag. Service delimiting tags are exact match and are stripped from the frame as illustrated in Figure 20. Exact match or service delimiting Tags do not require VID filters. VID filters can only be used to match on frame tags that are after the service delimiting tags.

With VID Filters operators can choose to match VID tags for up to two tags on ingress or egress or both.

- The outer-tag is the first tag in the packet that is carried transparently through the service.
- The inner-tag is the second tag in the packet that is carried transparently through the service.

VID Filters add the capability to perform VID value filter policies on default tags (1/1/1:* or 1/1/1:x.*, or 1/1/1/:*.0), or null tags (1/1/1, 1/1/1:0 or 1/1/1:x.0). The matching is based on the port configuration and the SAP configuration.

In the industry the QinQ tags are often referred to as the C-VID (Customer VID) and S-VID (service VID). The terms outer tag and inner tag allow flexibility without having to refer to C-TAG and an S-TAG explicitly. The position of inner and outer tags is relative to the port configuration and SAP configuration. Matching of tags is allowed for up to the first two tags on a frame. Since service delimiting tags may be 0, 1 or 2 tags.

The meaning of inner and outer has been designed to be consistent for egress and ingress when the number of non service delimiting tags is consistent. Service 1 in Figure 20 shows a conversion from qinq to a single dot1q example where there is one non-service delimiting tag on ingress and egress. Service 2 shows a symmetric example with two non-service delimiting tags (plus and additional tag for illustration) to two non-service delimiting tags on egress. Service 3 illustrates single non-service delimiting tags on ingress and to two tags with one non-service delimiting tag on ingress and egress.

SAP-ingress QoS setting allows for MAC-criteria type VID which uses the VID filter matching capabilities QoS and VID Filters (moved to QoS guide) on page 313.

A VID filter entry can also be used as a debug or lawful intercept mirror source entry.

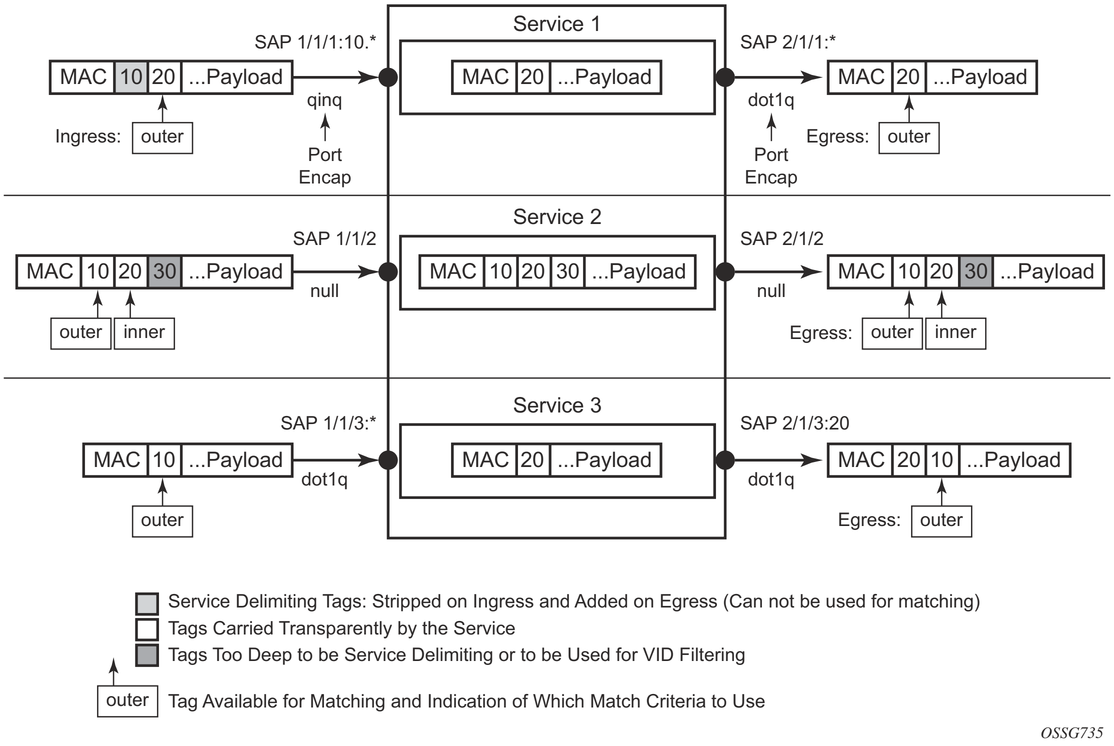

Figure 20: VID Filtering Examples

VID filters are available on Ethernet SAPs for Epipe, VPLS or I-VPLS including eth-tunnel and eth-ring services.

Arbitrary Bit Matching of VID Filters

In addition to matching an exact value, a VID filter mask allows masking any set of bits. The masking operation is ((value & vid-mask) = = (tag and vid-mask)). For example: A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6. VID filters allow explicit matching of VIDs and matching of any bit pattern within the VID tag.

When using VID filters on SAPs only VID filters are allowed on this SAP. Filters of type normal and ISID are not allowed.

An additional check for the “0” VID tag may be required when using certain wild card operations. For example frames with no tags on null encapsulated ports will match a value of 0 in outer tag and inner tag because there are no tags in the frame for matching. If a zero tag is possible but not desired it can be explicitly filtered using exact match on “0” prior to testing other bits for “0”.

Note that **configure>system>ethernet>new-qinq-untagged-sap** is a special QinQ function for single tagged QinQ frames with a null second tag. Using this in combination with VID filters is not recommended. Note that the outer-tag is the only tag available for filtering on egress for frames arriving from MPLS SDPs or from PBB services even though additional tags may be carried transparently.

Port Group Configuration Example

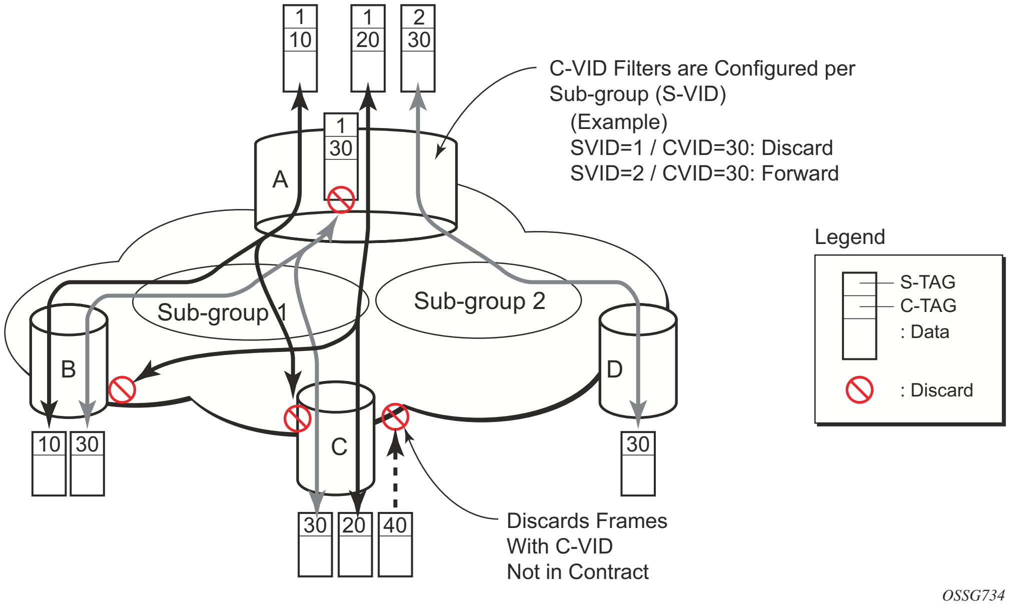

Figure 21: Port Groups

Figure 21 shows a customer use example where some VLANs are prevented from ingressing or egressing certain ports. In the example, port A sap 1/1/1:1.* would have a filter as shown below while port A sap 1/1/1:2.* would not.:

```
mac-filter 4 create
     default-action forward
             type vid
             entry 1 create
                 match frame-type ethernet_II
                     outer-tag 30 4095
                 exit
                 action drop
             exit
         exit
```

Creating and Applying ACL Policies

Figure 22 displays the process to create a redirect policy and to apply that policy to a service SAP or router interface.

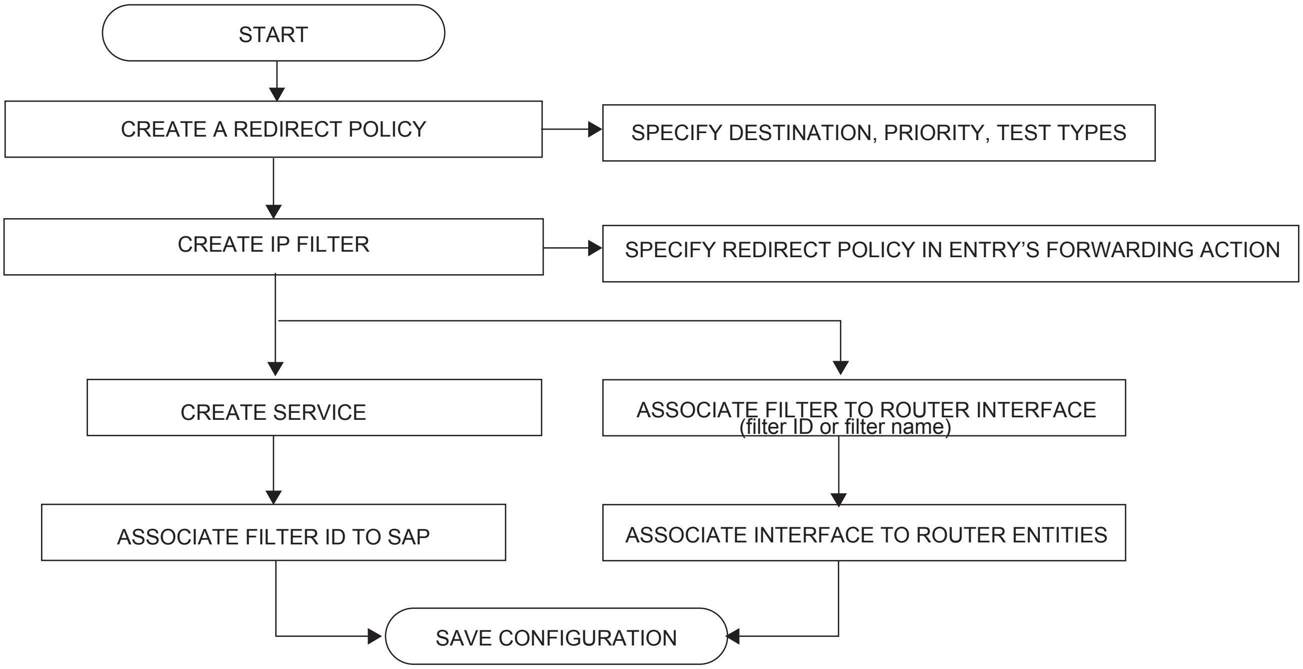
START
CREATE A REDIRECT POLICY
SPECIFY DESTINATION, PRIORITY, TEST TYPES
CREATE IP FILTER
SPECIFY REDIRECT POLICY IN ENTRY'S FORWARDING ACTION
CREATE SERVICE
ASSOCIATE FILTER TO ROUTER INTERFACE
(filter ID or filter name)
ASSOCIATE FILTER ID TO SAP
ASSOCIATE INTERFACE TO ROUTER ENTITIES
SAVE CONFIGURATION

Figure 22: Filter Creation and Implementation Flow

Figure 23 displays the process to create a filter policy and apply that policy to a service or network port.

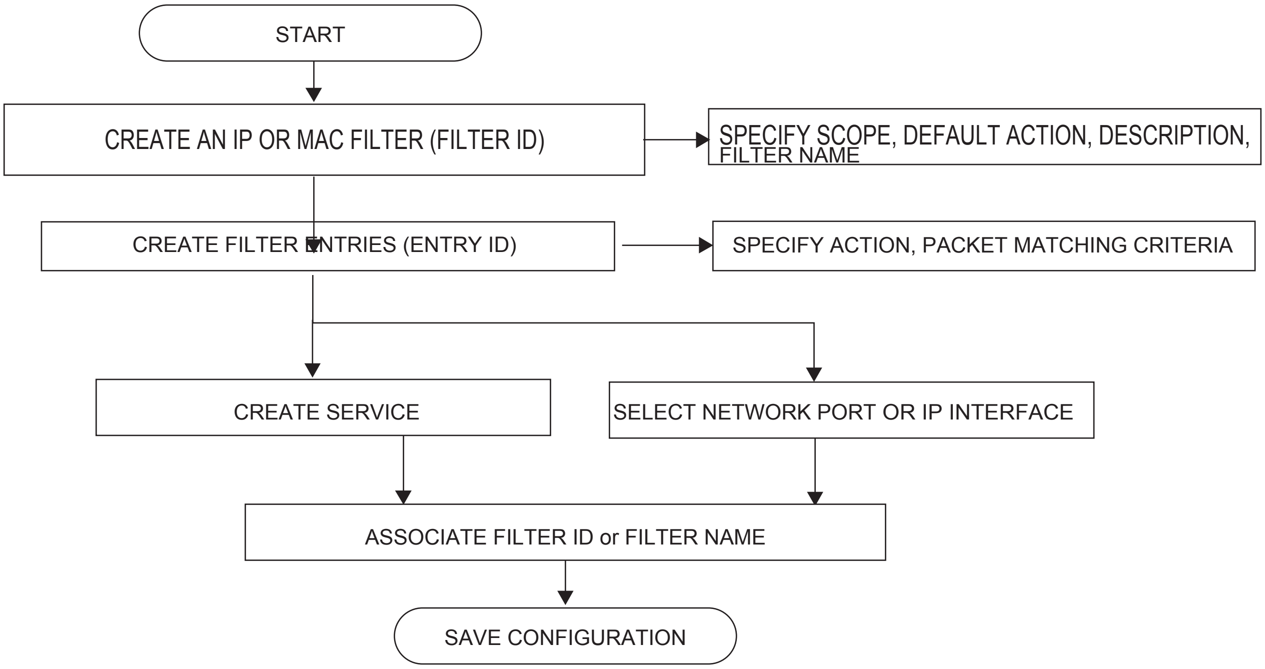
START
CREATE AN IP OR MAC FILTER (FILTER ID)
SPECIFY SCOPE, DEFAULT ACTION, DESCRIPTION, FILTER NAME
CREATE FILTER ENTRIES (ENTRY ID)
SPECIFY ACTION, PACKET MATCHING CRITERIA
CREATE SERVICE
SELECT NETWORK PORT OR IP INTERFACE
ASSOCIATE FILTER ID or FILTER NAME
SAVE CONFIGURATION

Figure 23: Creating and Applying Filter Policies

Applying Filters

After filters are created, they can be applied to the following entities:

- Applying a Filter to a SAP on page 409
- Applying a Filter to a Network Port a Network IP on page 409

Applying a Filter to a SAP

During the SAP creation process, ingress and egress filters are selected from a list of qualifying IP and MAC filters. When ingress filters are applied to a SAP, packets received at the SAP are checked against the matching criteria in the filter entries. If the packet completely matches all criteria in an entry, the checking stops and the entry's action is preformed. If the packet does not match any filter entries, the default filter action is applied.

When egress filters are applied to a SAP, packets received at the egress SAP are checked against the matching criteria in the filter entries. If the packet completely matches all criteria in an entry, the checking stops. If the packet does not match filter entries, the default filter action is applied.

Filters can be added or changed to an existing SAP configuration by modifying the SAP parameters. Filter policies are not operational until they are applied to a SAP and the service is enabled.

Applying a Filter to a Network Port a Network IP

An IP (v4 and/or IPv6) filter can be applied to a network port IP interface. Packets received on the interface are checked against the matching criteria in the filter entries. If the packet completely matches all criteria in an entry, the checking stops and the entry's action is performed. If the packets do not match any filter entries, they are discarded or forwarded based on the default action specified in the policy.

Packet Match Criteria

SROS-based routers/switches support L2, L3 and L4 and above match criteria in IPv4, IPv6 and MAC filters. Type and scale of each criteria supported depends on the platform, please see your Alcatel-Lucent representative for further details. As few or as many match parameters can be specified as required, but all conditions within a single filter policy entry must be met in order for the packet to be considered a match and the specified action performed. Any match criteria will be ignored unless explicitly defined. The process stops when the first complete match is found with triggers execution of the action defined in the entry.

IP filter policy entry match criteria includes the following:

- src-ip/dst-ip

  Match for the specified DSCP value against the Differentiated Services Code Point/Traffic Class field of the outer IPv4/IPv6 header of the packet.
- Destination IP address and mask — Destination IP address and mask values can be entered as search criteria.
- protocol — Match for the specified protocol against the Protocol field (for example, TCP, UDP, IGMP) of the outer IPv4 header of the packet.
- next-header — Match for the specified upper layer protocol (for example, TCP, UDP, IGMPv6) against the Next Header field of the outer IPv6 header of the packet. Note: next-header matching allows also to match on presence of some of the IPv6 extension headers. See CLI section for details on which extension header match is supported. An option to match either source or destination (Logical OR) using a single filter policy entry is supported for some filter policies by using a single **port** command.
- src-port/dst-port — When protocol (IPv4) or next-header (IPv6) specifies TCP, UDP, or both for this entry, it matches against the Source Port Number/Destination Port Number of the outer IPv4/IPv6 header of the packet.
- Destination port/range — Entering the destination port number or port range allows the filter to search for matching TCP or UDP values .
- dscp — Match for the specified DSCP value against the Differentiated Services Code Point/Traffic Class field of the outer IPv4/IPv6 header of the packet. See Table 9, DSCP Name to DSCP Value Table, on page 412.
- icmp-code — Match for the specified value against the Code field of the ICMP/ICMPv6 header of the packet.
- icmp-type — Match for the specified value against the Type field of the ICMP/ICMPv6 header of the packet.
- fragment — Enable fragmentation support in filter policy match. For IPv4, match against MF bit or Fragment Offset field to determine whether the packet is a fragment or not. For IPv6, match against Next Header Field for Fragment Extension Header value to determine

whether the packet is a fragment or not. For IPv6, match on initial fragment is also supported.

- ip-option — Match for the specified option in the first option of the IPv4 packet.
- option-present — Match for the presence or absence of the IP options in the IPv4 packet. Padding and EOOL are also considered as IP options.
- multiple-options — Match when an IPv4 packet contains multiple IP options or not.
- src-route-option — Match when a packet contains IP Option 3 or 9 (Loose or Strict Source Route) in the first 3 IP Options or if a packet has more than 3 IP Options.
- tcp-ack/tcp-syn — When protocol (IPv4) or next-header (IPv6) specify TCP, match for the TCP ACK/TCP SYNC flag presence/absence in the TCP header of the packet.

MAC filter policies match criteria includes the following:

- frame-type — Entering the frame type allows the filter to match for a specific type of frame format; for example, Ethernet-II will match for only ethernet-II frames.
- src-mac— Entering the source MAC address range allows the filter to search for matching a source MAC address and/or range.
- dst-mac— Entering the destination MAC address range allows the filter to search for matching a destination MAC address and/or range.
- dot1p — Entering an IEEE 802.1p value or range allows the filter to search for matching 802.1p frame.
- etype— Entering an Ethernet type II Ethertype value to be used as a filter match criterion. The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame.
- ssap— Specifying an Ethernet 802.2 LLC SSAP value allows the filter to match a source access point on the network node designated in the source field of a packet.
- dsap— Specifying an Ethernet 802.2 LLC DSAP value allows the filter to match a destination access point on the network node designated in the destination field of a packet.
- snap-pid— Specifying an Ethernet IEEE 802.3 LLC SNAP PID allows the filter to match the two-byte protocol ID that follows the three-byte OUI field. The DSAP and mask accepts decimal and hex in the range of 0 to 65535.
- isid — Specifying an Ethernet IEEE 802.1ag ISID from the I-TAG value allows the filter to match the 24 bits ISID value from the PBB I-TAG. This match criteria is mutually exclusive with all the other match criteria under a particular mac-filter policy. The resulting mac-filter can be applied as required on a BVPLS SAP or PW basis just in the egress direction. A new **mac-filter type** attribute is defined to control the use of ISID match criteria and must be set to **isid** to allow the use of isid match criteria.
- inner-tag/outer-tag — Specifying inner-tag/outer-tag VLAN ID values allows the filter to match on the non-service delimiting tags as described earlier in this document. This match

criteria is mutually exclusive with all other match criteria under a particular mac-filter policy. A new mac-filter type attribute is defined to control the use of inner-tag/outer-tag match criteria and must be set to vid to allow the use of inner-tag/outer0-tag match criteria.

DSCP Values

Table 9: DSCP Name to DSCP Value Table

DSCP Name	Decimal DSCP Value	Hexadecimal DSCP Value	Binary DSCP Value
default	0	*	
cp1	1		
cp2	2		
cp3	3		
cp4	4		
cp5	5		
cp6	6		
cp7	7	*	
cs1	8		
cp9	9		
af10	10	*	
af11	11	*	
af12	12	*	
cp13	13		
cp14	14		
cp15	15		
cs2	16	*	
cp17	17		
af21	18	*	
cp19	19		
af22	20	*	
cp21	21		
af23	22	*	
cp23	23		
cs3	24	*	
cp25	25		

Table 9: DSCP Name to DSCP Value Table (Continued)

DSCP Name	Decimal DSCP Value	Hexadecimal DSCP Value	Binary DSCP Value
af31	26	*	
cp27	27		
af32	28	*	
cp29	29		
af33	30	*	
cp21	31		
cs4	32	*	
cp33	33		
af41	34	*	
cp35	35		
af42	36	*	
cp37	37		
af43	38	*	
cp39	39		
cs5	40	*	
cp41	41		
cp42	42		
cp43	43		
cp44	44		
cp45	45		
ef	46	*	
cp47	47		
nc1	48	*	(cs6)
cp49	49		
cp50	50		
cp51	51		
cp52	52		
cp53	53		
cp54	54		
cp55	55		
cp56	56		
cp57	57		

Table 9: DSCP Name to DSCP Value Table (Continued)

DSCP Name	Decimal DSCP Value	Hexadecimal DSCP Value	Binary DSCP Value
nc2	58	*	(cs7)
cp60	60		
cp61	61		
cp62	62		

IP Option Values

Table 10: IP Option Values

Copy	Class	Number	Value	Name	Description
0	0	0	0	EOOL	End of options list
0	0	1	1	NOP	No operation
0	0	7	7	RR	Record route
0	0	10	10	ZSU	Experimental measurement
0	0	11	11	MTUP	MTU probe
0	0	12	12	MTUR	MTU reply
0	0	15	15	ENCODE	
0	2	4	68	TS	Time stamp
0	2	18	82	TR	Traceroute
1	0	2	130	SEC	Security
1	0	3	131	LSR	Loose source router
1	0	5	133	E-SEC	Extended security
1	0	6	134	CIPSO	Commercial security
1	0	8	136	SID	Stream id
1	0	9	137	SSR	Strict source route
1	0	14	142	VISA	Experimental Access Control [Estrin]
1	0	16	144	IMITD	IMI Traffic Descriptor
1	0	17	145	EIP	Extended Internet Protocol
1	0	19	147	ADDEXT	Address Extension
1	0	20	148	RTRALT	Router alert
1	0	21	149	SDB	Selective directed broadcast
1	0	22	150	NSAPA	NSAP addresses
1	0	23	151	DPS	Dynamic packet state
1	0	24	152	UMP	Upstream multicast packet
1	2	13	205	FINN	Experimental flow control

Ordering Filter Entries

When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit entry. Filter matching ceases when a packet matches an entry. The entry action is performed on the packet. To be considered a match, the packet must meet all the match criteria defined in the entry.

Packets are compared to entries in a filter policy in an ascending entry ID order. To reorder entries in a filter policy, edit the entry ID value; for example, to reposition entry ID 6 to a more explicit location, change the entry ID *6* value to entry ID *2* using the **renum** filter policy command.

When a filter consists of a single entry, the filter executes actions as follows:

- If a packet matches all the entry criteria, the entry's specified action is performed (drop or forward).
- If a packet does not match all of the entry criteria, the policy's default action is performed.

If a filter policy contains two or more entries, packets are compared in ascending entry ID order (1, 2, 3 or 10, 20, 30, etc.):

- Packets are compared with the criteria in the first entry ID (the lowest numberical entry ID value).
- If a packet matches all the match criteria defined in the entry, the entry's specified action is executed.
- If a packet does not match, the packet continues to the next entry, and so on until a match is found or until all entries are compared.
- If a packet does not completely match any entries, then the default action is performed.

Figure 24 displays an example of several packets forwarded upon matching the filter criteria and several packets traversing through the filter entries and then dropped.

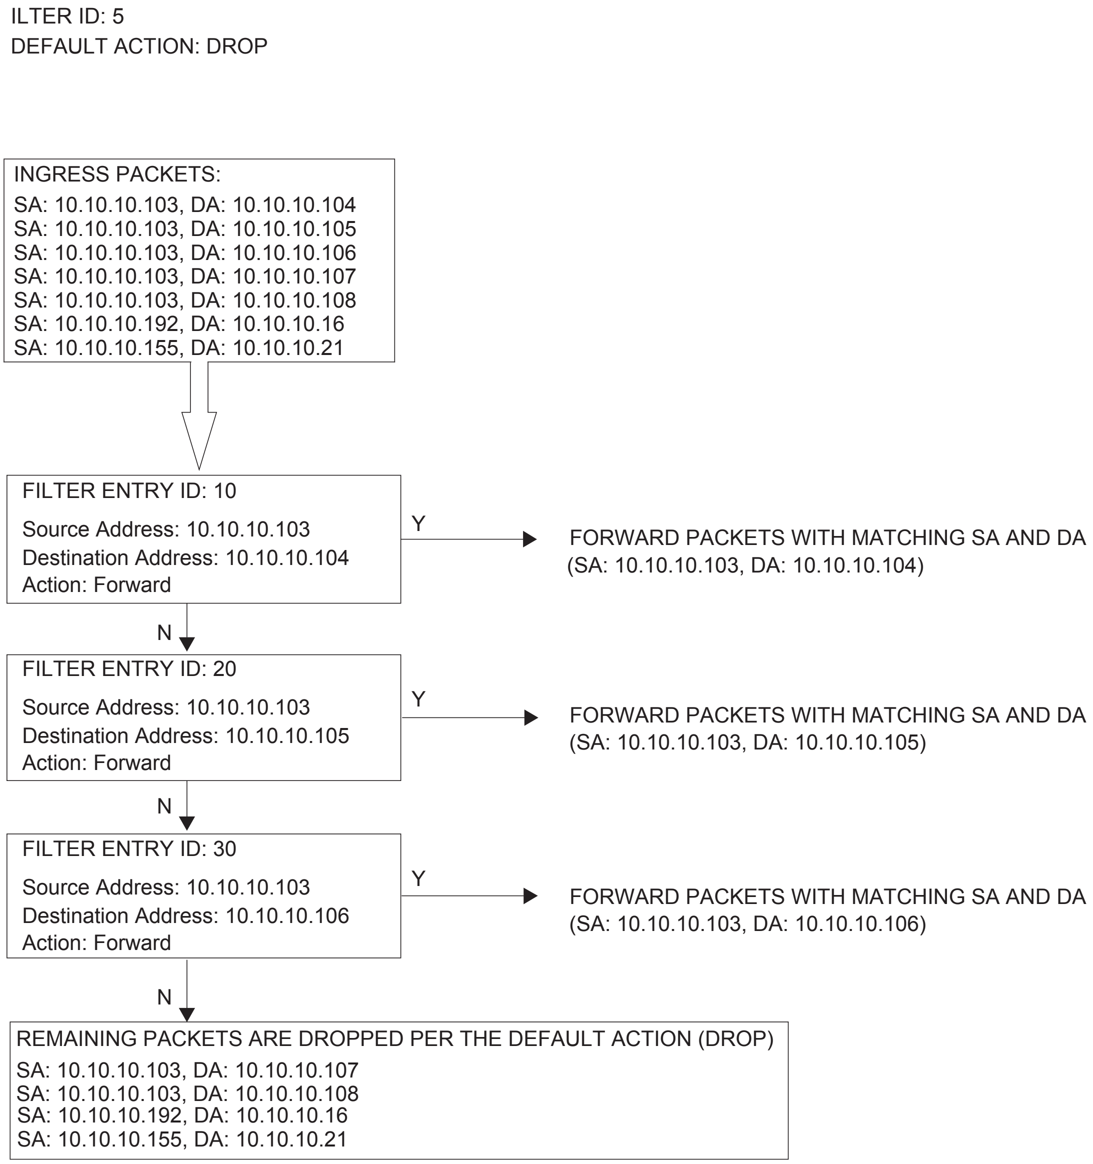
ILTER ID: 5
DEFAULT ACTION: DROP
INGRESS PACKETS:
SA: 10.10.10.103, DA: 10.10.10.104
SA: 10.10.10.103, DA: 10.10.10.105
SA: 10.10.10.103, DA: 10.10.10.106
SA: 10.10.10.103, DA: 10.10.10.107
SA: 10.10.10.103, DA: 10.10.10.108
SA: 10.10.10.192, DA: 10.10.10.16
SA: 10.10.10.155, DA: 10.10.10.21
FILTER ENTRY ID: 10
Source Address: 10.10.10.103
Destination Address: 10.10.10.104
Action: Forward
Y
FORWARD PACKETS WITH MATCHING SA AND DA
(SA: 10.10.10.103, DA: 10.10.10.104)
N
FILTER ENTRY ID: 20
Source Address: 10.10.10.103
Destination Address: 10.10.10.105
Action: Forward
Y
FORWARD PACKETS WITH MATCHING SA AND DA
(SA: 10.10.10.103, DA: 10.10.10.105)
N
FILTER ENTRY ID: 30
Source Address: 10.10.10.103
Destination Address: 10.10.10.106
Action: Forward
Y
FORWARD PACKETS WITH MATCHING SA AND DA
(SA: 10.10.10.103, DA: 10.10.10.106)
N
REMAINING PACKETS ARE DROPPED PER THE DEFAULT ACTION (DROP)
SA: 10.10.10.103, DA: 10.10.10.107
SA: 10.10.10.103, DA: 10.10.10.108
SA: 10.10.10.192, DA: 10.10.10.16
SA: 10.10.10.155, DA: 10.10.10.21

Figure 24: Filtering Process Example

Configuration Notes

The following information describes filter implementation caveats:

- Creating a filter policy is optional.
- Associating a service with a filter policy is optional.
- When a filter policy is configured, it should be defined as having either an *exclusive* scope for one-time use, or a *template* scope meaning that the filter can be applied to multiple SAPs.
- A specific filter must be explicitly associated with a specific service in order for packets to be matched.
- Each filter policy must consist of at least one filter entry. Each entry represents a collection of filter match criteria. When packets enter the ingress or egress ports, packets are compared to the criteria specified within the entry or entries.
- When a large (complex) filter is configured, it may take a few seconds to load and initiate the filter policy configuration.
- The action keyword must be entered for the entry to be active. Any filter entry without the action keyword will be considered incomplete and be inactive.

MAC Filters

- If a MAC filter policy is created with an entry and entry action specified but the packet matching criteria is not defined, then all packets processed through this filter policy entry will pass and take the action specified. There are no default parameters defined for matching criteria.
- MAC filters cannot be applied to routable VPLS.
- MAC filters cannot be applied to network interfaces.
- Some of the MAC filter of type "normal" match criteria fields are exclusive to each other, based on the type of Ethernet frame. Use the following table to determine the exclusivity of fields.

Table 11: MAC Match Criteria Exclusivity Rules

Frame Format	Etype	LLC – Header (ssap & dsap)	SNAP-OUI	SNAP- PID
Ethernet – II	Yes	No	No	No
802.3	No	Yes	No	No
802.3 – snap	No	No [a]	Yes	Yes

a.

Note: When snap header is present, this is always set to AA-AA.

IP Filters

- IP filters are used for IPv4 traffic only. IPv6 filters are to be used for IPv6 traffic. If a filter policy is created with an entry and entry action specified but the packet matching criteria is not defined, then all packets processed through this filter policy entry will pass and take the action specified. There are no default parameters defined for matching criteria.
- An action parameter must be specified for the entry to be active. Any filter entry without an action parameter specified will be considered incomplete and be inactive.

IPv6 Filters

- If a filter policy is created with an entry and entry action specified but the packet matching criteria is not defined, then all packets processed through this filter policy entry will pass and take the action specified. There are no default parameters defined for matching criteria.
- An action parameter must be specified for the entry to be active. Any filter entry without an action parameter specified will be considered incomplete and be inactive.

Log Filter

- Summarization logging is the collection and summarization of log messages for 1 specific log-id within a period of time.
- Filter log can be applied to different filters or CPM hardware filters.
- The implementation of the feature applies to filter logs with destination syslog.
- In case of VPLS scenario both Layer 2 & Layer 3 are applicable.
  - → Layer 2: Source MAC or optionally destination MAC
  - → Layer 3: Source IPv6 or optionally destination IPv6 for Layer 3 filters.
- The summarization interval is 100 seconds.
- Upon activation of a summary, a mini-table with src/dst-address and count is created for each type (IP/IPv6/MAC).
- Every received log packet (due to filter hit) is examined for source or destination address. If the log packet (source/destination address) matches a source/destination address entry in the mini-table a packet received previously), the summary counter of the matching address is incremented.
- If source or destination address of the log messages does not match an entry already present in the table, the source/destination address is stored in a free entry in the mini-table.

- In case the mini-table has no more free entries, only total counter is incremented.
- At expiry of the summarization interval, the mini-table for each type is flushed to the syslog destination.