User-Based Rate Limiting in the Cisco Catalyst 6500

When the Cisco® Catalyst® 6500 Series Supervisor Engine 720 was introduced, it brought with it a number of new features that used hardware advancements found in the policy feature card 3 (PFC3). In contrast to major features such as IPv6 and Multiprotocol Label Switching (MPLS) that gained a lot of attention, user-based rate limiting (UBRL) is one feature that is not as well known.

UBRL is a form of microflow policing allowing the administrator to rate limit traffic flows, but unlike a normal microflow policer, it allows a policer to be applied to all traffic to or from a specific user. This paper will explore UBRL in more detail by providing information about how UBRL works and some examples as to how UBRL can be deployed and configured.

Microflow Policing: A Recap

The Cisco Catalyst 6500 has supported two different forms of policing since its inception: aggregate policing and microflow policing. A total of 1023 aggregate policers and 63 microflow policers are supported on the Supervisor Engine 720, and it supports just over 128,000 flows.

An aggregate policer applies a rate-limiting policy to all traffic in a VLAN or port that matches set classification criteria (defined using an access control list [ACL]). For example, if an aggregate policer defining a rate of 50 Mb were applied to a VLAN containing 10 physical ports, then all traffic matching the classification criteria (set by the ACL) entering those ports in that VLAN would be subject to that policing rule and would not exceed 50 Mb.

The microflow policer differs in that it applies a rate-limiting policy on a per-flow basis. Whereas the aggregate policer limits the total amount of traffic entering that VLAN, the same microflow policer would only limit each flow to the stated rate. If a microflow policer were applied to the same VLAN enforcing a policing rule of 2 Mb, then no one flow entering any port in that VLAN could exceed 2 Mb. It is worth noting that although a microflow policer limits traffic for specific flows, it does not limit the number of flows that can be active in that VLAN.

Another area of differentiation between the microflow and aggregate is where the policer can be applied. On a PFC3x-based supervisor, the aggregate policer can be applied on ingress or egress, whereas the microflow policer can only be applied on ingress. Microflow policing is also limited to a single instance of the token leaky bucket whereas the aggregate policer has the option of using a dual token leaky bucket.

Netflow and the Flow Entry

UBRL differs from a normal microflow policer in the way it views a flow. Within the Supervisor Engine 720, it is the NetFlow process that is responsible for collecting information about flows that exist in the switch and in part defines how a microflow policer and a UBRL view and act on a flow. NetFlow stores information about each flow in memory located on the PFC. The definition of a flow and what is stored in a flow entry in the NetFlow table are determined by what is called the flow mask. The flow mask identifies fields in the packet header that are used to perform a lookup into

the NetFlow table and is what defines a flow. In the PFC1 and PFC2, the flow mask can use the following fields to define a flow:

- All flows with a unique destination only IP address
- All flows with a unique source and destination IP address
- Full flow (source and destination IP address, protocol, source and destination port number)

When a microflow policer is defined on a PFC1 or PFC2, it requires the use of a full flow mask. As the PFC1 and PFC2 can only use a single flow mask at any one time, it means that when a microflow policer is enabled, other processes that use the flow mask also have to use the same full flow mask. Other processes that use flow entries in the NetFlow table (and flow masks) include Network Address Translation (NAT), Port Address Translation (PAT), TCP intercept, NetFlow Data Export, Web Cache Communication Protocol (WCCP), content-based access control (CBAC), Cisco IOS® Software SLB: Server Load Balancing, and reflexive ACLs. Using a full flow mask is the most specific mask and results in more flow entries being consumed into the NetFlow table. This has implications for the utilization and capacity of the NetFlow table.

With the introduction of the Supervisor Engine 720 came the PFC3. The PFC3 incorporates a number of hardware enhancements over older PFCs, one of which is the ability to store more than one flow mask. In total, the Supervisor Engine 720 supports four flow masks in hardware. This is available in the PFC3a, PFC3B, and PFC3BXL. There is one flow mask that is reserved for multicast, and a second is reserved for system use. That leaves two flow masks that are available for normal system use, and these are the masks that UBRL can utilize.

The PFC3x also introduces a number of new flow masks. The complete supported mask list includes the options shown in Table 1.

Table 1. Flow Masks Available on the PFC3x

Flow Mask Type	Description
Source Only	A less-specific flow mask. The PFC maintains one entry for each source IP address. All flows from a given source IP address use this entry.
Destination Only	A less-specific flow mask. The PFC maintains one entry for each destination IP address. All flows to a given destination IP address use this entry.
Destination-Source	A more-specific flow mask. The PFC maintains one entry for each source and destination IP address pair. All flows between same source and destination IP addresses use this entry.
Destination-Source-Interface	A more-specific flow mask. Adds the source VLAN Simple Network Management Protocol (SNMP) ifIndex to the information in the destination-source flow mask.
Full	The PFC creates and maintains a separate cache entry for each IP flow. A full entry includes the source IP address, destination IP address, protocol, and protocol interfaces.
Full-Interface	The most specific flow mask. The flow mask adds the source VLAN SNMP ifIndex to the information in the full-flow mask.

Flow Masks

As discussed in the previous section, the flow mask defines the way in which the system views a flow. It defines the pieces of information in the packet header that identify the flow. More importantly, it defines which traffic UBRL will see that constitutes the flow. In the context of UBRL, three of the flow masks identified in the previous section are used. Those flow masks are source only, destination only, and full flow masks. The quality-of-service (QoS) ACL that inspects traffic on ingress to the switch initially determines the flow mask that is used by the UBRL process. The full flow mask uses the source and destination IP address as well as the source and destination port numbers. An example of this is shown in Figure 1.

Figure 1. Full Flow Mask Fields

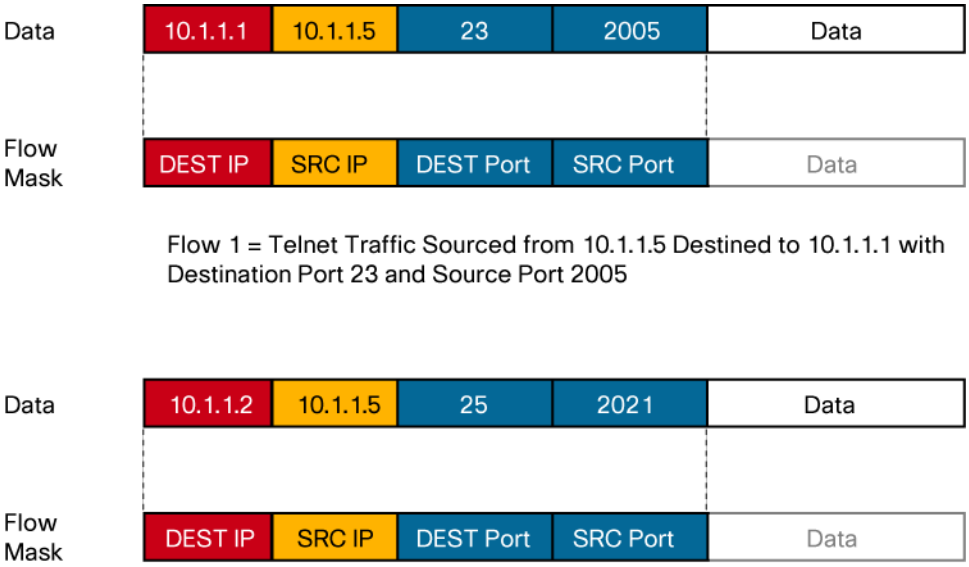

The full flow diagram in Figure 1 uses the source and destination IP addresses as well as the source and destination port numbers to identify each flow. A given user who initiates a Telnet session and accesses an e-mail server would initiate two separate flows. Because Telnet and e-mail use discrete port numbers, they would be recognized as discrete flows.

The source-only IP flow mask uses the source IP address in its packet header. An example of how this is used is shown in Figure 2.

Figure 2. Source IP Only Flow Mask Fields

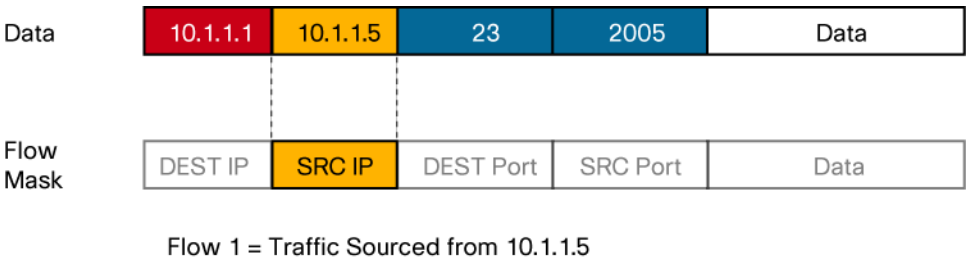

In our earlier example, the same user who initiated a Telnet and e-mail session would now be seen as initiating a single flow. The reason for this is that each flow shares a common source IP address, and it is only the source address that the flow mask is using to identify unique flows.

The last flow mask used by UBRL is the destination IP only flow mask. This mask uses the destination IP address field in the packet's header to identify a unique flow. An example of this is shown in Figure 3.

Figure 3. Destination IP Only Flow Mask Fields

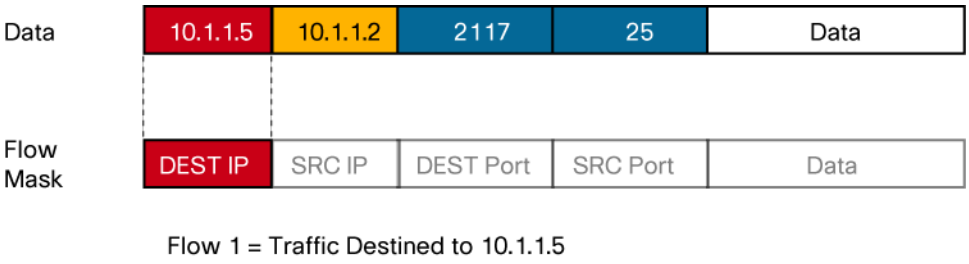

This flow mask uses the destination IP address to determine what constitutes a flow. In Figure 3, traffic is destined (returning) to host 10.1.1.5. Traffic has arrived from an e-mail server (10.1.1.2) and an FTP server (10.2.1.7). In both cases, traffic from each server is considered to be part of the same flow as the mask is, only using the destination address as the unique flow identifier. The destination IP only flow mask, however, is generally used for inbound (or return) rather than outbound traffic. It also tends to be used in conjunction with the source IP only flow mask. Examples of this will be shown later in this paper.

Netflow Table

Flow entries are stored in the NetFlow table on the PFC. To facilitate high-speed lookups for flow entries, a special piece of high-speed lookup memory called ternary content addressable memory (TCAM), also located on the PFC, is used. On the Supervisor Engine 720, three PFC3x options are available, each of which dictates what the supervisor model is. The capacities of each PFC3x with respect to the number of flows that can be stored in the NetFlow table are shown in Table 2.

Table 2. Supervisor Engine 720 NetFlow Capacities

Supervisor	PFC Type	Total Capacity	Effective Capacity	Hash Efficiency
Supervisor Engine 720	PFC3a	128 K	64 K	50%
Supervisor Engine 720-3B	PFC3B	128 K	115 K	90%
Supervisor Engine 720-3BXL	PFC3BXL	256 K	230 K	90%

It is worthwhile noting that although the PFC3B has the same NetFlow table size as the PFC3a, it uses an improved hash algorithm, which increases the effective utilization of the NetFlow table.

The PFC uses a hash algorithm to locate and store flow entries in the NetFlow table. The flow mask identifies the fields of interest in the packet's header. Those interesting fields are used as input to the hash algorithm. The hash algorithm process will point to a TCAM location, which contains a key. They key provides the index into the NetFlow table, which contains the actual NetFlow entry. This process is shown in Figure 4.

Figure 4. NetFlow Hash Operation

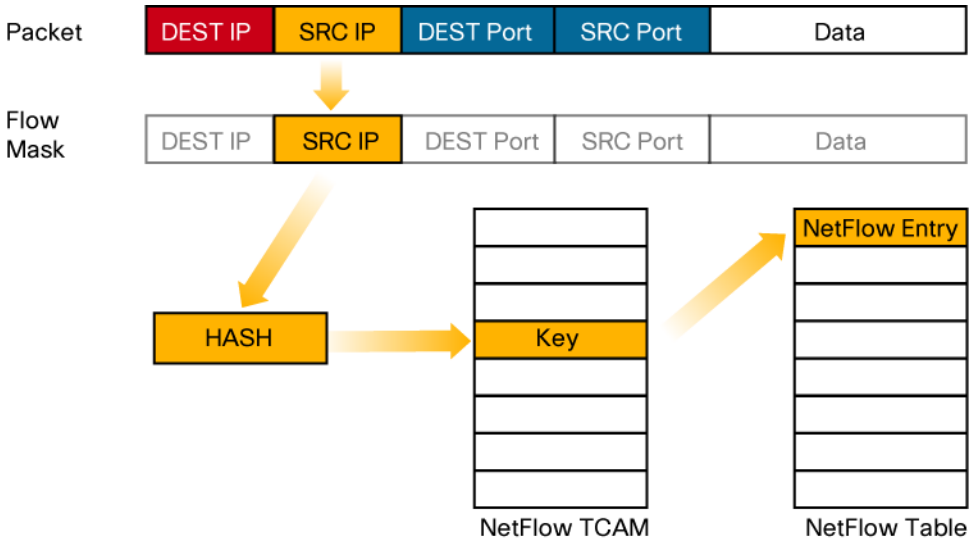

Configuring UBRL: Simple Example

Consider the following scenario. A service provider is about to bring on a new customer: ABZ Corp. This customer shares a 100-Mb Ethernet connection into its office building with other consumers. The customer premises equipment connecting into the customer's network uses PAT to translate all outbound communications using a single source IP address. The customer has paid for 10 Mb of access into the service provider network, so the service provider needs to rate limit the customer to this agreed rate. (See Figure 5.)

Figure 5. UBRL Configuration Example for ABZ Corp

UBRL is configured in a similar way to a microflow policer. The configuration starts with the definition of a set of classification criteria using an ACL. The customer's IP address is 201.10.1.5, so the ACL is initially defined as follows:

```
6500(config)# access-list 102 permit ip host 201.10.1.5 any
6500(config)# class-map identify-ABZ-traffic
6500(config-cmap)# match access-group 102
```

The configuration just described uses an ACL (102) to identify traffic from host 201.10.1.5 to any destination. This ACL is referenced within a class map called **identify-ABZ-traffic**. Next a policy map is created with the policer included. Within the policy map will be a policer statement. Normally, a microflow policer is identified by the use of the keyword flow. UBRL uses this keyword, but it also uses a **flow** mask keyword to set the flow mask required for this operation. The following is the full configuration:

```
6500(config)# policy-map police-customer-traffic
6500(config-pmap)# class identify-ABZ-traffic
6500(config-pmap-c)# police flow mask src-only 10000000 5000
conform-action transmit exceed action drop
6500(config-pmap-c)# interface gig3/1
6500(config-if)# service-policy input police-customer-traffic
```

A policy map called **police-customer-traffic** has been created. Within the policy map, the class map of **identify-ABZ-traffic** is used. This class map includes the classification criteria set by the ACL in the previous step. Within the class map, the UBRL policer is created. The keyword **flow** is underlined to emphasize that this is a microflow policer. Following the flow keyword is the mask keyword, indicating the flow mask value to be used for this operation. With UBRL, three different mask values can be set. These are shown in the following command-line interface example:

```
6500(config-pmap-c)# police flow mask ?
 dest-only  destination-address only flow mask
 full-flow  full flow mask
 src-only   source-address only flow mask
```

This configuration example uses the source IP address of the customer as the flow identifier. For this reason, the **src-only** keyword is used to define the mask for this UBRL operation. The **src-only** mask tells UBRL that any traffic with the same source IP address should be considered as part of the same flow. In this regard, we now have the means to identify all traffic from that user and apply a single policer to that traffic to rate limit it to 10 Mb.

Finally, the policy would need to be mapped to Gigabit Ethernet interface 3/1 using the service-policy keyword. Once applied to the interface, the policy is now active.

Configuring UBRL: Multiple Sources

This next example takes the previous scenario and adds an additional two customers in the same building (for a total of three customers in this building). Each customer has been assigned a unique IP address with which that customer is viewed by the outside world. Each customer has the network credentials and requirements shown in Table 3.

Table 3. Customer IP Addresses and Requested Bandwidths

Customer Name	Source IP Address	Requested Bandwidth
ABZ Corp	201.10.1.5	10 Mb
What's Happnin Corp (WHC)	201.10.1.6	5 Mb
Just Cruzin Corp (JCC)	201.10.1.7	1 Mb

The network is set out as shown in Figure 6.

Figure 6. UBRL Configuration Example for Multiple Sources

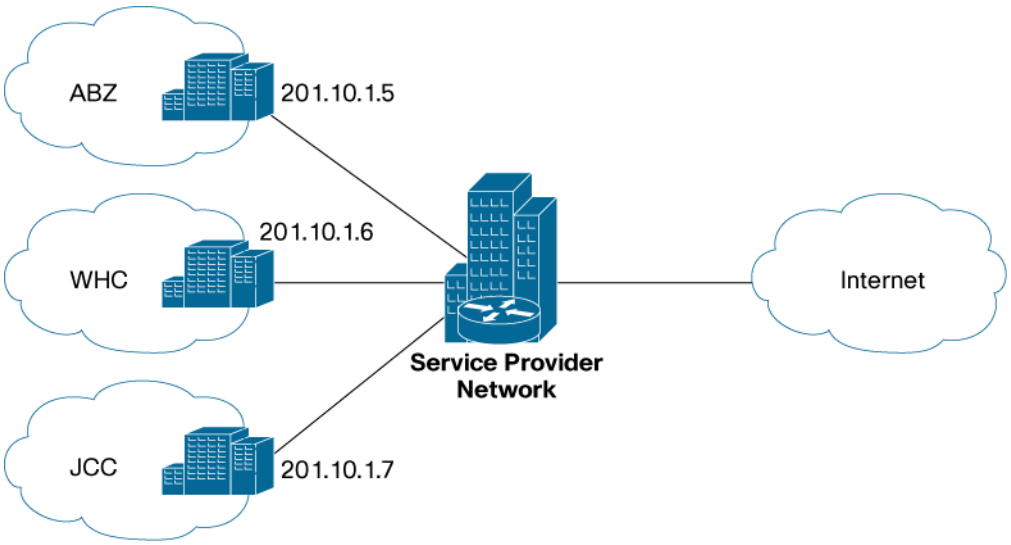

Based on these credentials, we can configure the network up as follows:

```
6500(config)# access-list 102 permit ip host 201.10.1.5 any
6500(config)# access-list 103 permit ip host 201.10.1.6 any
6500(config)# access-list 104 permit ip host 201.10.1.7 any

6500(config)# class-map identify-ABZ-traffic
6500(config-cmap)# match access-group 102
6500(config)# class-map identify-WHC-traffic
```

```
6500(config-cmap)# match access-group 103
6500(config)# class-map identify-JCC-traffic
6500(config-cmap)# match access-group 104

6500(config)# policy-map police-customer-traffic
6500(config-pmap)# class identify-ABZ-traffic
6500(config-pmap-c)# police flow mask src-only 10000000 5000
conform-action transmit exceed action drop
6500(config-pmap)# class identify-WHC-traffic
6500(config-pmap-c)# police flow mask src-only 5000000 2500
conform-action transmit exceed action drop
6500(config-pmap)# class identify-JCC-traffic
6500(config-pmap-c)# police flow mask src-only 1000000 1000
conform-action transmit exceed action drop
6500(config-pmap-c)# interface gig3/1
6500(config-if)# service-policy input police-customer-traffic
```

In this example, we have three different classes defined for each customer connecting to the service provider network. Each class defines a separate policer rate to meet the criteria for that customer. Each policer is using the source only flow mask, making sure that all traffic originating from each customer is treated as a single flow.

Configuring UBRL: Bidirectional UBRL

UBRL is well suited to university requirements, where student traffic can compromise available bandwidth on the campus. The sheer volume of peer-to-peer (P2P) traffic created by a host of file-sharing applications is usually the main culprit. Many universities seek ways to limit this traffic to a more manageable level. UBRL can be applied to rate limit both outbound and inbound traffic to users of the network.

Figure 7. Bidirectional UBRL Scenario Example

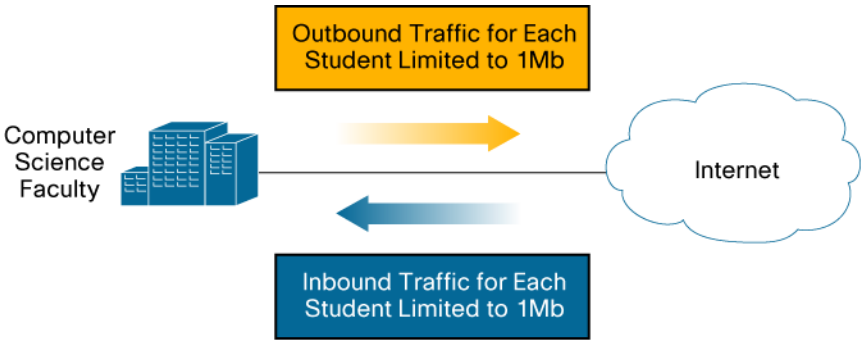

In this example (Figure 7), two flow masks will be combined to limit traffic to and from users in the computer science faculty. Each user (student) will be limited to uploading or downloading no more than 1 Mb of data. The computer science faculty uses the 202.25.1.0/24 subnet. In order to achieve this, the following configuration could be applied:

```
6500(config)# access-list 142 permit ip 202.25.1.0 0.0.0.255 any
6500(config)# access-list 143 permit ip any 202.25.1.0 0.0.0.255

6500(config)# class-map identify-outbound-student
6500(config-cmap)# match access-group 142
6500(config)# class-map identify-inbound-student
6500(config-cmap)# match access-group 143

6500(config)# policy-map police-student-traffic-outbound
6500(config-pmap)# class identify-outbound-student
6500(config-pmap-c)# police flow mask src-only 1000000 1000
conform-action transmit exceed action drop
6500(config)# policy-map police-student-traffic-inbound
6500(config-pmap)# class identify-inbound-student
6500(config-pmap-c)# police flow mask dest-only 1000000 1000
conform-action transmit exceed action drop

6500(config-pmap-c)# interface gig8/22
6500(config-if)# service-policy input police-student-traffic-inbound
6500(config-pmap-c)# interface gig8/23
6500(config-if)# service-policy input police-student-traffic-outbound
```

In this example, two separate ACLs are defined: one to classify outbound traffic by focusing on the source subnet and the second to classify inbound traffic focusing on the destination subnet (this is return traffic back to the faculty). Two separate policers are configured: one for outbound traffic and one for return traffic. Each policer uses a different flow mask to match on interesting traffic to or from the faculty. For outbound traffic, the policer uses a source only flow mask to match on traffic originating from the faculty. Each unique user will be limited to 1 Mb of upstream bandwidth. Return traffic matching on the inbound policer sues the destination-only IP flow mask. This matches on faculty user addresses and limits their download bandwidth also to 1 Mb.

Summary

UBRL provides a way to apply a rate-limiting policy for all traffic sourced from or destined to a given user. This enhanced policing capability is embedded in the Supervisor Engine 720 (more specifically the PFC3x) and is made possible by its support for multiple flow masks. The ability of UBRL to extend the functionality of the base microflow policing feature provides customer networks with additional policing deployment options. More importantly, it allows the definition of specific rate-limiting policies to limit all traffic from a given user, something not available with earlier supervisor models.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA

C11-384405-00 12/06