Distributed CPU Protection

In This Chapter

This section describes Distributed CPU Protection (DCP) configurations.

Topics in this section include:

Applicability

This Distributed CPU Protection (DCP) configuration example was created using the 7750 SR-c12 platform but is equally applicable to the following platforms: 7750 SR-7/12, 7450 ESS-6/7/12, 7750 SR-c4/c12 and 7950 XRS. DCP is not supported on the 7750 SR-1, 7450 ESS-1 or 7710 SR platforms.

DCP operates on the line cards and requires line cards with the FP2 or greater hardware (for example, IOM3-XP, IMMs and C-XMAs).

The configuration was tested on release 11.0R1.

Overview

SR OS provides several rate limiting mechanisms to protect the CPM/CFM processing resources of the router:

- CPU Protection: A centralized rate limiting function that operates on the CPM to limit traffic destined to the CPUs.
- Distributed CPU Protection: A control traffic rate limiting protection mechanism for the CPM/CFM that operates on the line cards (hence 'distributed'). CPU protection protects the CPU of the node that it is configured on from a DOS attack by limiting the amount of traffic coming in from one of its ports and destined to the CPM (to be processed by its CPU) using a combination of the configurable limits.

The goal of this example is to familiarize the reader with the configuration and use of Distributed CPU Protection. A simple and controlled setup is used to illustrate how the protection behaves and how to use the tools provided for the feature.

External testing equipment ("tester") is used to send control traffic of various protocols at various rates to the router in order to exercise DCP. Log events and show routines are examined to explain the indications that the router provides to an operator.

Configuration

The test topology is shown in Figure 299. A Gigabit Ethernet link is used between the Tester and the router.

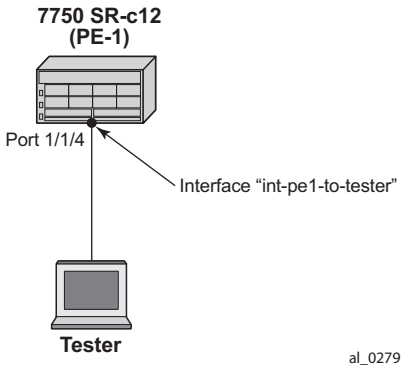

Figure 299: Test Topology

Step 1. The basic configuration of the mda, port, interface and a security event log on the router is shown below.

```
*A:PE-1# configure card 1 mda 1
*A:PE-1>config>card>mda# info
----------------------------------------------
            mda-type m5-1gb-sfp-b
            no shutdown
----------------------------------------------
*A:PE-1>config>card>mda# exit all
*A:PE-1# configure port 1/1/4
*A:PE-1>config>port# info
----------------------------------------------
        ethernet
        exit
        no shutdown
----------------------------------------------
*A:PE-1>config>port# exit all
*A:PE-1# configure router interface "int-pe1-to-tester"
*A:PE-1>config>router>if# info
----------------------------------------------
            address 192.168.10.1/24
            port 1/1/4
            no shutdown
----------------------------------------------
*A:PE-1>config>router>if# exit all
*A:PE-1# configure log log-id 15
*A:PE-1>config>log>log-id# info
----------------------------------------------
            from security
            to memory 1024
----------------------------------------------
```

This example was developed on a 7750 SR-c12 platform but it is equally applicable to other platforms such as the 7750 SR-7/12. If other platforms, such as the 7750 SR-7/12 that support centralized CPU Protection, are used to explore Distributed CPU Protection then the centralized CPU Protection should be disabled (for the purposes of this example) so that it does not interfere with reproducing the same results as described below. In a normal production network CPU Protection and DCP are complimentary and can be used together. To disable centralized CPU Protection for the purposes of reproducing the results below please ensure that:

- **protocol-protection** is disabled.
- All rates in all polices (including any default polices) are configure to **max**.

Step 2. In order to activate DCP a policy is created and assigned to the interface.

The first policy that is used in this example is used to simply count protocol packets to see that they are indeed flowing from the tester to the router and being extracted and indentified.

The *dcp-policy-count* policy is configured as follows:

```
*A:PE-1# configure system security dist-cpu-protection
*A:PE-1>config>sys>security>dist-cpu-protection# info
----------------------------------------------
                policy "dcp-policy-count" create
                    description "Static policers with rate 0 for counting packets"
                    static-policer "sp-arp" create
                        rate packets 0 within 1
                    exit
                    static-policer "sp-icmp" create
                        rate packets 0 within 1
                    exit
                    static-policer "sp-igmp" create
                        rate packets 0 within 1
                    exit
                    protocol arp create
                        enforcement static "sp-arp"
                    exit
                    protocol icmp create
                        enforcement static "sp-icmp"
                    exit
                    protocol igmp create
                        enforcement static "sp-igmp"
                    exit
                exit
```

For the *dcp-policy-count policy* configuration:

- The policy contains three static policers: s*p-arp*, *sp-icmp* and *sp-igmp*. These policers are then used by the three configured protocols that are part of the policy: a*rp*, *icmp* and *igmp*.
- The list of protocols that are applicable to DCP are as follows: arp, dhcp, http-redirect, icmp, igmp, mld, ndis, pppoe-pppoa, all-unspecified, mpls-ttl, bfd-cpm, bgp, eth-cfm, isis, ldp, ospf, pim and rsvp. The all-unspecified protocol is a special "catch-all". Please see the 7750 SR OS System Management Guide for more details.
- This policy instantiates three permanent (static) policers for every object (for example, interface) that the policy is associated with.
- The three protocols each reference their own static policer so each protocol will be independently rate limited. A single static policer can also be used to rate limit multiple protocols but that capability is not used in this example.
- The rate is set to 0 which means all packets will be considered as non-conformant to the policer. This configuration is used to provide counters of protocol packets. The DCP counters provide the count of packets exceeding the policing parameters since the given policer was previously declared as conformant or newly instantiated. A rate of zero ensures that the policer will never be declared as conformant and hence will never reset the counters.
- The exceed-action is not configured and takes the default value of *none.* The *log-events* parameter is not configured and is enabled by default. That means the policer will notify the operator when the first packet arrives but will not discard or mark any packets.

Step 3. Assign the *dcp-policy-count* to the interface:

```
*A:PE-1# configure router interface "int-pe1-to-tester"
*A:PE-1>config>router>if# dist-cpu-protection "dcp-policy-count"
```

Step 4. Examine some log and status on the router to get a baseline (no traffic is flowing from the tester to the router at this point). Notice that the cpu utilization is fairly low with an overall Idle of 96% and no task groups at more than 5% capacity usage. Future example output from this show routine will be snipped to only show relevant and interesting lines.

```
*A:PE-1# show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                         Usage
-------------------------------------------------------------------------------
BFD                                            0           0.00%           0.00%
BGP                                       28,779           0.32%           0.47%
BGP PE-CE                                      0           0.00%           0.00%
CFLOWD                                     7,384           0.08%           0.38%
Cards & Ports                             65,941           0.73%           5.35%
DHCP Server                                   55          ~0.00%          ~0.00%
ICC                                        1,195           0.01%           0.06%
IGMP/MLD                                   1,883           0.02%           0.12%
```

```
IMSI Db Appl                                120          ~0.00%          ~0.00%
IOM                                     132,522           1.47%           3.11%
IP Stack                                  7,666           0.08%           0.39%
IS-IS                                     1,415           0.01%           0.07%
ISA                                      11,988           0.13%           0.43%
LDP                                         496          ~0.00%           0.04%
Logging                                     185          ~0.00%           0.01%
MBUF                                          0           0.00%           0.00%
MPLS/RSVP                                 6,219           0.06%           0.48%
MSCP                                          0           0.00%           0.00%
MSDP                                          0           0.00%           0.00%
Management                                4,077           0.04%           0.13%
OAM                                      10,311           0.11%           0.44%
OSPF                                        661          ~0.00%           0.05%
PIM                                           0           0.00%           0.00%
RIP                                           0           0.00%           0.00%
RTM/Policies                                  0           0.00%           0.00%
Redundancy                                7,641           0.08%           0.51%
SNMP Daemon                                   0           0.00%           0.00%
Services                                  3,965           0.04%           0.09%
Stats                                         0           0.00%           0.00%
Subscriber Mgmt                           7,437           0.08%           0.44%
System                                   57,081           0.63%           3.49%
Traffic Eng                                   0           0.00%           0.00%
VRRP                                      1,918           0.02%           0.09%
WEB Redirect                                 77          ~0.00%          ~0.00%
-------------------------------------------------------------------------------
Total                                 8,965,427         100.00%
   Idle                               8,605,657          95.98%
   Usage                                359,770           4.01%
Busiest Core Utilization                134,481          13.49%
===============================================================================
```

The DCP feature is reporting no violations for interfaces on card 1.

```
*A:PE-1# tools dump security dist-cpu-protection violators enforcement interface card 1
===============================================================================
Distributed Cpu Protection Current Interface Enforcer Policer Violators
===============================================================================
Interface                       Policer/Protocol                    Hld Rem
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Violators on Slot-1 Fp-1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
[S]-Static [D]-Dynamic [M]-Monitor
-------------------------------------------------------------------------------
===============================================================================
```

There are no security log events.

```
*A:PE-1# show log log-id 15
===============================================================================
Event Log 15
===============================================================================
Description : (Not Specified)
Memory Log contents  [size=1024   next event=1  (not wrapped)]
```

The detailed DCP status for the interface shows all three policers are currently in the conform state.

```
*A:PE-1# show router interface "int-pe1-to-tester" dist-cpu-protection
===============================================================================
Interface "int-pe1-to-tester" (Router: Base)
===============================================================================
Distributed CPU Protection Policy :  dcp-policy-count
-------------------------------------------------------------------------------
Statistics/Policer-State Information
===============================================================================
-------------------------------------------------------------------------------
Static Policer
-------------------------------------------------------------------------------
Policer-Name        : sp-arp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : arp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none

Policer-Name        : sp-icmp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : icmp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none

Policer-Name        : sp-igmp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : igmp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Local-Monitoring Policer
-------------------------------------------------------------------------------
No entries found
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Dynamic-Policer (Protocol)
-------------------------------------------------------------------------------
No entries found
-------------------------------------------------------------------------------
===============================================================================
```

Step 5. Configure the tester to send ARP, ICMP and IGMP traffic to the router using the following rates:

- ARP: 2 packets per second (pps)
- ICMP: 4 pps
- IGMP: 8 pps

Here are some tips for how to configure the tester to send protocol packets that will be recognized by the router:

- ARP:
  - → Set the MAC destination address to FF-FF-FF-FF-FF-FF
  - → Use an ARP Request format
- ICMP:
  - → Use an icmp type of 8 (echo request, such as **ping**).
  - → Set the MAC destination address equal to the MAC address of the receiving port. The MAC address of port 1/1/4 can be seen in the output of show port 1/1/4 as the Configured Address.
  - → Set the IP destination address to 192.168.10.1 and the IP source address to 192.168.10.2.
- IGMP:
  - → Set the MAC destination address equal to the MAC address of the receiving port. The MAC address of port 1/1/4 can be seen in the output of show port 1/1/4 as the Configured Address.
  - → Set the IP destination address to 224.0.0.2 and the IP source address to 0.0.0.0.
  - → Set the IGMP version to 2, make the IGMP message type a Membership Query to Group 0.

Also ensure that the tester interleaves the three streams of protocol packets such that it schedules them independently in an interleaved fashion, not serially.

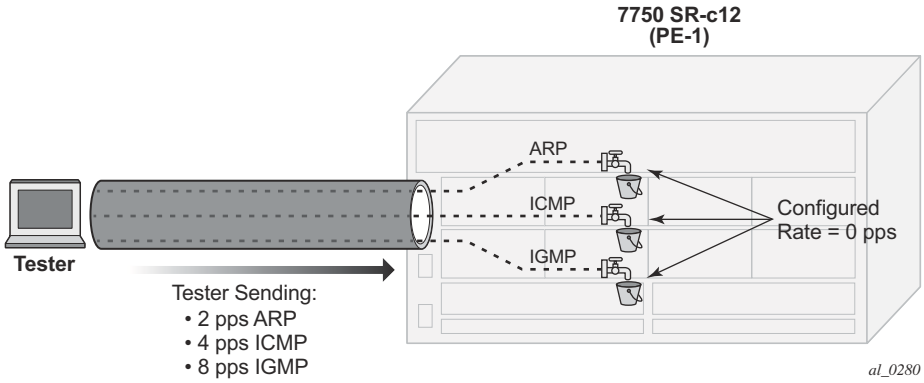

Figure 300: Count Traffic with DCP Policy Count

Step 6. Notice that DCP now reports some violations of the policy against the interface.

```
*A:PE-1# tools dump security dist-cpu-protection violators enforcement interface card 1
===============================================================================
Distributed Cpu Protection Current Interface Enforcer Policer Violators
===============================================================================
Interface                       Policer/Protocol                    Hld Rem
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Violators on Slot-1 Fp-1
-------------------------------------------------------------------------------
int-pe1-to-tester               sp-arp                           [S] none
int-pe1-to-tester               sp-icmp                          [S] none
int-pe1-to-tester               sp-igmp                          [S] none
-------------------------------------------------------------------------------
[S]-Static [D]-Dynamic [M]-Monitor
-------------------------------------------------------------------------------
===============================================================================
```

After a few seconds the DCP exceed-count values can be seen incrementing.

Note the following details:

- Exceed-Count is non-zero. This will continue incrementing and will never reset since the rate configured in the DCP policy is zero.
- The Policer-State is Exceed. The policers have detected that the protocol is non-conformant to the configured rate.
- Detec. Time Remain stays at 29 seconds. This countdown timer is automatically reset to 30 seconds every time a policer is detected as non-conformant (which will be continually when the rate is set to 0 and packets of that protocol are being received).

```
*A:PE-1# show router interface "int-pe1-to-tester" dist-cpu-protection
===============================================================================
Interface "int-pe1-to-tester" (Router: Base)
===============================================================================
Distributed CPU Protection Policy :  dcp-policy-count
-------------------------------------------------------------------------------
Statistics/Policer-State Information
===============================================================================
-------------------------------------------------------------------------------
Static Policer
-------------------------------------------------------------------------------
Policer-Name        : sp-arp
Card/FP             : 1/1               Policer-State       : Exceed
Protocols Mapped    : arp
Exceed-Count        : 72
Detec. Time Remain  : 29 seconds        Hold-Down Remain.   : none

Policer-Name        : sp-icmp
Card/FP             : 1/1               Policer-State       : Exceed
Protocols Mapped    : icmp
Exceed-Count        : 144
Detec. Time Remain  : 29 seconds        Hold-Down Remain.   : none

Policer-Name        : sp-igmp
Card/FP             : 1/1               Policer-State       : Exceed
Protocols Mapped    : igmp
Exceed-Count        : 290
Detec. Time Remain  : 29 seconds        Hold-Down Remain.   : none
-------------------------------------------------------------------------------
[snip]
```

Step 7. Keep the tester running.

Now a DCP policy that enforces protocol rates using static policers will be applied to the interface. First, the policy is created:

```
*A:PE-1# configure system security dist-cpu-protection
*A:PE-1>config>sys>security>dist-cpu-protection# policy "dcp-static-policy-1" create
                    description "Static policers for arp, icmp and igmp"
                    static-policer "sp-arp" create
                        rate packets 10 within 1
                        exceed-action discard
                    exit
                    static-policer "sp-icmp" create
                        rate packets 20 within 1
                        exceed-action discard
                    exit
                    static-policer "sp-igmp" create
                        rate packets 10 within 1
                        exceed-action discard
                    exit
                    protocol arp create
                        enforcement static "sp-arp"
                    exit
                    protocol icmp create
                        enforcement static "sp-icmp"
                    exit
                    protocol igmp create
```

```
                    enforcement static "sp-igmp"
                exit
            exit
```

For the dcp-static-policy-1 policy configuration, note that a few parameters are different than in the previously created dcp-policy-cou*nt* policy:

- The rates are set to low (but non-zero) values.
- The exceed-action is configured such that packets are dropped once the rate is exceeded.

Now assign the policy to the test interface:

```
*A:PE-1# configure router interface "int-pe1-to-tester"
*A:PE-1>config>router>if# dist-cpu-protection "dcp-static-policy-1"
*A:PE-1>config>router>if# exit all
*A:PE-1# show system security dist-cpu-protection policy "dcp-static-policy-1" association
===============================================================================
Distributed CPU Protection Policy
===============================================================================
Policy Name : dcp-static-policy-1
Description : Static policers for arp, icmp and igmp

-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------

SAP associations
-------------------------------------------------------------------------------
  None

Managed SAP associations
-------------------------------------------------------------------------------
  None

Interface associations
-------------------------------------------------------------------------------
Router-Name : Base
  int-pe1-to-tester
-------------------------------------------------------------------------------
Number of interfaces : 1
===============================================================================
```

Step 8. Increase the rate of IGMP packets that the tester is sending to 1000pps (keep ARP and ICMP at 2pps and 4pps).

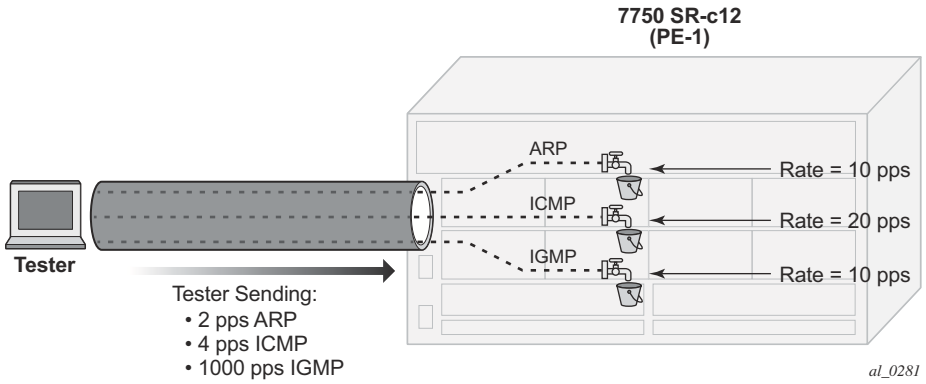

Figure 301: Limit Traffic with dcp-static-policy-1

Step 9. Notice that the system has identified a violation of the DCP rates for the igmp policer.

```
*A:PE-1# tools dump security dist-cpu-protection violators enforcement interface card 1
===============================================================================
Distributed Cpu Protection Current Interface Enforcer Policer Violators
===============================================================================
Interface                       Policer/Protocol                    Hld Rem
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Violators on Slot-1 Fp-1
-------------------------------------------------------------------------------
int-pe1-to-tester             sp-igmp                           [S] none
-------------------------------------------------------------------------------
[S]-Static [D]-Dynamic [M]-Monitor
-------------------------------------------------------------------------------
===============================================================================
```

After a few minutes the violation will be indicated as a log event. This delay is due to the design of DCP. In order to support large scale operation of DCP, and also to avoid overload conditions, a polling process is used to monitor state changes in the policers and to gather violations. This means there can be a delay between when an event occurs in the data plane and when the relevant state change or event notification occurs towards an operator, but in the meantime the policers are still operating and protecting the control plane.

```
*A:PE-1# show log log-id 15
===============================================================================
Event Log 15
===============================================================================
Description : (Not Specified)
```

```
Memory Log contents  [size=1024   next event=11  (not wrapped)]

10 2013/04/18 17:31:54.58 EDT WARNING: SECURITY #2066 Base DCPUPROT
"Non conformant network_if "int-pe1-to-tester" on fp 1/1 detected at 04/18/2013 17:31:33.
Policy "dcp-static-policy-1". Policer="sp-igmp"(static). Excd count=135"
… [snip] …
```

The status of DCP on the interface also shows the igmp policer as being in an Exceed state:

```
*A:PE-1# show router interface "int-pe1-to-tester" dist-cpu-protection
===============================================================================
Interface "int-pe1-to-tester" (Router: Base)
===============================================================================
Distributed CPU Protection Policy :  dcp-static-policy-1
-------------------------------------------------------------------------------
Statistics/Policer-State Information
===============================================================================
-------------------------------------------------------------------------------
Static Policer
-------------------------------------------------------------------------------
Policer-Name        : sp-arp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : arp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none

Policer-Name        : sp-icmp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : icmp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none

Policer-Name        : sp-igmp
Card/FP             : 1/1               Policer-State       : Exceed
Protocols Mapped    : igmp
Exceed-Count        : 19031
Detec. Time Remain  : 29 seconds        Hold-Down Remain.   : none
-------------------------------------------------------------------------------
…[snip]…
```

The CPU utilization of the IGMP task group is not impacted since DCP is discarding packets that are non-conformant to the configure rate.

```
*A:PE-1# show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                         Usage
-------------------------------------------------------------------------------
BFD                                            0           0.00%           0.00%
…[snip]…
IGMP/MLD                                   1,883           0.02%           0.12%
IMSI Db Appl                                 120          ~0.00%          ~0.00%
IOM                                      132,522           1.47%           3.11%
IP Stack                                   7,666           0.08%           0.39%
```

```
IS-IS                                      1,415           0.01%           0.07%
ISA                                       11,988           0.13%           0.43%
LDP                                          496          ~0.00%           0.04%
…[snip]…
WEB Redirect                                  77          ~0.00%          ~0.00%
-------------------------------------------------------------------------------
Total                                  8,965,427         100.00%
   Idle                                8,605,657          95.98%
   Usage                                 359,770           4.01%
Busiest Core Utilization                 134,481          13.49%
===============================================================================
```

Step 10. Remove the DCP policy from the interface and see the CPU utilization goes up for the IGMP task group.

```
*A:PE-1# configure router interface "int-pe1-to-tester"
*A:PE-1>config>router>if# no dist-cpu-protection
*A:PE-1>config>router>if# /show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                          Usage
-------------------------------------------------------------------------------
BFD                                             0           0.00%           0.00%
…[snip]…
IGMP/MLD                                   82,142           0.91%           8.14%
IMSI Db Appl                                   98          ~0.00%          ~0.00%
IOM                                       129,851           1.45%           3.15%
IP Stack                                  196,549           2.19%          19.35%
IS-IS                                       1,484           0.01%           0.07%
ISA                                        11,765           0.13%           0.42%
LDP                                           449          ~0.00%           0.04%
…[snip]…
WEB Redirect                                  102          ~0.00%           0.01%
-------------------------------------------------------------------------------
Total                                  8,948,806         100.00%
   Idle                                8,259,903          92.30%
   Usage                                 688,903           7.69%
Busiest Core Utilization                 210,435          21.16%
===============================================================================
```

Step 11. Increase the rate of IGMP traffic from the tester to 5000 pps. See the CPU utilization increase further.

```
*A:PE-1# show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                          Usage
-------------------------------------------------------------------------------
BFD                                             0           0.00%           0.00%
…[snip]…
IGMP/MLD                                  417,124           4.65%          41.78%
IMSI Db Appl                                   82          ~0.00%          ~0.00%
IOM                                       133,029           1.48%           2.92%
```

```
IP Stack                                  935,491          10.43%          93.45%
IS-IS                                       1,343           0.01%           0.06%
ISA                                        12,350           0.13%           0.45%
LDP                                           394          ~0.00%           0.03%
…[snip]…
WEB Redirect                                  116          ~0.00%           0.01%
-------------------------------------------------------------------------------
Total                                   8,966,128         100.00%
   Idle                                 6,972,962          77.77%
   Usage                                1,993,166          22.22%
Busiest Core Utilization                  484,748          48.65%
===============================================================================
```

Step 12. Reinstall the DCP policy to the interface and see the CPU utilization drop.

```
*A:PE-1# configure router interface "int-pe1-to-tester"
*A:PE-1>config>router>if# dist-cpu-protection "dcp-static-policy-1"
*A:PE-1>config>router>if# exit all
*A:PE-1# show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                           Usage
-------------------------------------------------------------------------------
BFD                                             0           0.00%           0.00%
…[snip]…
IGMP/MLD                                    2,058           0.02%           0.10%
IMSI Db Appl                                   48          ~0.00%          ~0.00%
IOM                                       135,148           1.50%           3.04%
IP Stack                                    7,851           0.08%           0.47%
IS-IS                                       1,398           0.01%           0.07%
ISA                                        11,730           0.13%           0.43%
LDP                                           299          ~0.00%           0.02%
…[snip]…
WEB Redirect                                   71          ~0.00%          ~0.00%
-------------------------------------------------------------------------------
Total                                   8,975,262         100.00%
   Idle                                 8,611,593          95.94%
   Usage                                  363,669           4.05%
Busiest Core Utilization                  136,669          13.70%
===============================================================================
```

Step 13. Stop the tester from sending packets, wait a few minutes and then note the status of the system.

There are no longer any violations of any enforcement policers on any interfaces on card 1.

```
*A:PE-1# tools dump security dist-cpu-protection violators enforcement interface card 1
===============================================================================
Distributed Cpu Protection Current Interface Enforcer Policer Violators
===============================================================================
Interface                      Policer/Protocol                   Hld Rem
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Violators on Slot-1 Fp-1
```

```
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
[S]-Static [D]-Dynamic [M]-Monitor
-------------------------------------------------------------------------------
===============================================================================
```

The IGMP policer is indicated as conformant in the log events.

```
*A:PE-1# show log log-id 15
===============================================================================
Event Log 15
===============================================================================
Description : (Not Specified)
Memory Log contents  [size=1024   next event=7  (not wrapped)]

…[snip]…

12 2013/04/18 17:42:12.43 EDT WARNING: SECURITY #2072 Base DCPUPROT
"Network_if "int-pe1-to-tester" on fp 1/1 newly conformant at 04/18/2013 17:41:57:27. Pol-
icy "dcp-static-policy-1". Policer="sp-igmp"(static). Excd count=316418"

…[snip]…
```

The interface DCP details show all policers as conformant.

```
*A:PE-1# show router interface "int-pe1-to-tester" dist-cpu-protection
===============================================================================
Interface "int-pe1-to-tester" (Router: Base)
===============================================================================
Distributed CPU Protection Policy :  dcp-static-policy-1
-------------------------------------------------------------------------------
Statistics/Policer-State Information
===============================================================================
-------------------------------------------------------------------------------
Static Policer
-------------------------------------------------------------------------------
Policer-Name        : sp-arp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : arp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none

Policer-Name        : sp-icmp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : icmp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none

Policer-Name        : sp-igmp
Card/FP             : 1/1               Policer-State       : Conform
Protocols Mapped    : igmp
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none
-------------------------------------------------------------------------------
…[snip]…
```

An optional hold-down can be used in the configuration of the exceed-action of the policers in order to apply the exceed-action for a defined period (even if the policer goes conformant again during that period). The hold-down could be used, for example, to discard all packets associated with a policer for one hour after a violation is detected. An "indefinite" period is also supported which enforces discard or marking until the operator clears the policer with the **tools perform security dist-cpu-protection release-hold-down** command.

Step 14. The next scenario explored in this example is the use of DCP dynamic enforcement.

In order to use dynamic enforcement policers, a number of dynamic policers must be allocated to the DCP pool for the particular card being used.

```
*A:PE-1# configure card 1 fp dist-cpu-protection
*A:PE-1>config>card>fp>d-cpu-prot# info
----------------------------------------------
                dynamic-enforcement-policer-pool 1000
----------------------------------------------
```

The number allocated should be greater than the maximum number of dynamic policers expected to be in use on the card at one time. A conservative (large) number could be selected at first, and then the following show command can give data to help tune the pool to a smaller size over time:

```
*A:PE-1# show card 1 fp 1 dist-cpu-protection
===============================================================================
Card : 1 Forwarding Plane(FP) : 1
===============================================================================
Dynamic Enforcement Policer Pool : 1000
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Statistics Information
-------------------------------------------------------------------------------
Dynamic-Policers Currently In Use      : 0
Hi-WaterMark Hit Count                 : 0
Hi-WaterMark Hit Time                  : 04/20/2013 08:16:24 UTC
Dynamic-Policers Allocation Fail Count : 0
-------------------------------------------------------------------------------
===============================================================================
```

If the dynamic-enforcement-policer-pool is too small then when a local-monitoring-policer detects violating traffic, the dynamic enforcement policers will not be able to be instantiated. A log event will warn the operator when the pool is nearly exhausted.

A sample dynamic enforcement policy is created as follows:

```
*A:PE-1# configure system security dist-cpu-protection
*A:PE-1>config>sys>security>dist-cpu-protection# policy "dcp-dynamic-policy-1" create
                    description "Dynamic policing policy"
                    local-monitoring-policer "local-mon" create
                        description "Monitor for arp, icmp, igmp
                                     and all-unspecified"
                        rate packets 100 within 10
                    exit
```

```
protocol arp create
    enforcement dynamic "local-mon"
    dynamic-parameters
        rate packets 20 within 10
        exceed-action discard
    exit
exit
protocol icmp create
    enforcement dynamic "local-mon"
    dynamic-parameters
        rate packets 20 within 10
        exceed-action discard
    exit
exit
protocol igmp create
    enforcement dynamic "local-mon"
    dynamic-parameters
        rate packets 20 within 10
        exceed-action discard
    exit
exit
protocol all-unspecified create
    enforcement dynamic "local-mon"
    dynamic-parameters
        rate packets 100 within 10
        exceed-action discard
    exit
exit
```

For the *dcp-dynamic-policy-1* policy configuration:

- The policy contains no static policers. Per-protocol enforcement policers will be instantiated dynamically but only if triggered by a violation of the local-monitoring-policer.
- A local-monitoring-policer is configured for the policy. The configured rate determines the rate of arriving protocol packets at which the policy will trigger the automatic instantiation of dynamic per-protocol policers for the interface.
- Four protocols are configured and they are all associated with the local-monitoring-policer. The all-unspecified protocol will include all other extracted control packets on the interface.
- Each protocol has its own configured dynamic rates that will be used by the dynamic enforcement policers if they are instantiated. Note these rates are lower than previous scenarios (the **within** parameter is 10 seconds instead of 1 second).
- When this DCP policy is associated with an interface, only a single policer (the local-monitoring-policer) will be instantiated (statically/permanently). The per-protocol dynamic policers are only instantiated when there is a violation of the local-monitoring-policer.

The policy is then associated with the interface:

```
*A:PE-1# configure router interface "int-pe1-to-tester"
*A:PE-1>config>router>if# dist-cpu-protection "dcp-dynamic-policy-1"
```

Step 15. Configure the tester to send:

- 1pps of ARP
- 4pps of ICMP
- 1000pps of IGMP

Start the tester.

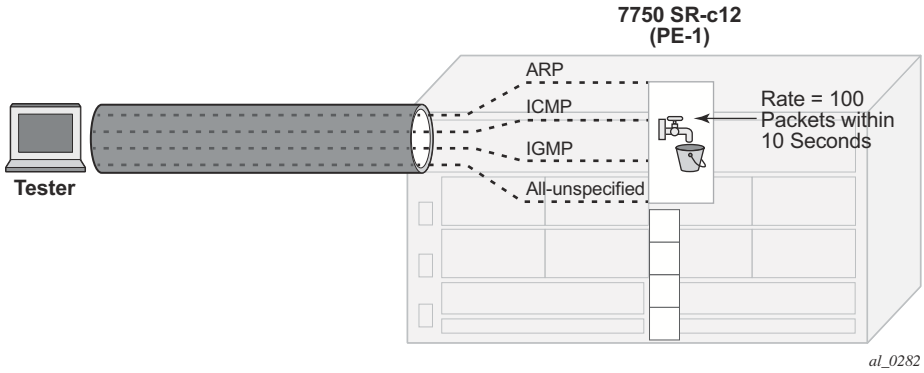

Figure 302: Dynamic Policing – Local Monitor

In Figure 302, the dynamic policers have not been instantiated yet.

Step 16. The local-monitoring-policer will become non-conforming since the aggregate arrival rate of arp+icmp+igmp+all-unspecified packets is greater than the configured local-monitoring-policer rate of 100 packets within 10 seconds. Dynamic enforcement policers will then be instantiated.

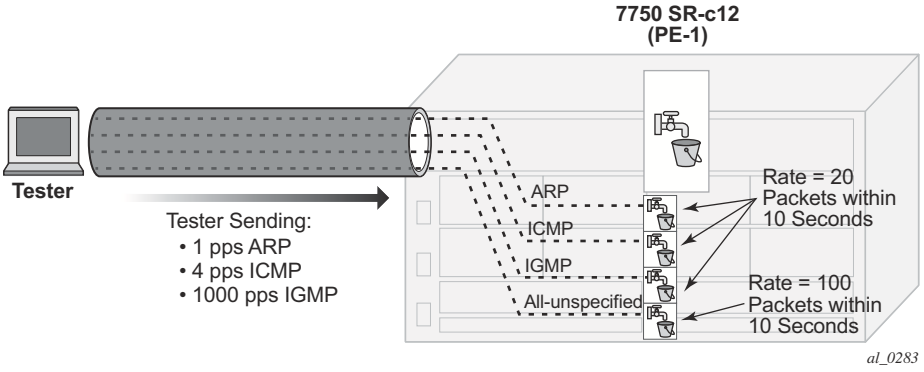

Figure 303: Dynamic Policers Instantiated

The ICMP and IGMP dynamic policers will see violations since their dynamic rates are being exceeded.

```
*A:PE-1# tools dump security dist-cpu-protection violators enforcement interface card 1
===============================================================================
Distributed Cpu Protection Current Interface Enforcer Policer Violators
===============================================================================
Interface                      Policer/Protocol                    Hld Rem
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Violators on Slot-1 Fp-1
-------------------------------------------------------------------------------
int-pe1-to-tester              icmp                              [D] none
int-pe1-to-tester              igmp                              [D] none
-------------------------------------------------------------------------------
[S]-Static [D]-Dynamic [M]-Monitor
-------------------------------------------------------------------------------
===============================================================================
```

The arp and all-unspecified dynamic policers were instantiated but will be counting down their detection time (if this show command is issued within 30 seconds of the attack starting).

```
*A:PE-1# show router interface "int-pe1-to-tester" dist-cpu-protection
===============================================================================
Interface "int-pe1-to-tester" (Router: Base)
===============================================================================
Distributed CPU Protection Policy :  dcp-dynamic-policy-1
-------------------------------------------------------------------------------
Statistics/Policer-State Information
===============================================================================
-------------------------------------------------------------------------------
Static Policer
-------------------------------------------------------------------------------
No entries found
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Local-Monitoring Policer
-------------------------------------------------------------------------------
Policer-Name        : local-mon
Card/FP             : 1/1               Policer-State       : Exceed
Protocols Mapped    : arp, icmp, igmp, all-unspecified
Exceed-Count        : 1097
All Dyn-Plcr Alloc. : True
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Dynamic-Policer (Protocol)
-------------------------------------------------------------------------------
Protocol(Dyn-Plcr)  : arp
Card/FP             : 1/1               Protocol-State      : Conform
Exceed-Count        : 0
Detec. Time Remain  : 5 seconds         Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True

Protocol(Dyn-Plcr)  : icmp
Card/FP             : 1/1               Protocol-State      : Exceed
Exceed-Count        : 31
```

```
Detec. Time Remain  : 28 seconds        Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True

Protocol(Dyn-Plcr)  : igmp
Card/FP             : 1/1               Protocol-State      : Exceed
Exceed-Count        : 23867
Detec. Time Remain  : 29 seconds        Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True

Protocol(Dyn-Plcr)  : all-unspecified
Card/FP             : 1/1               Protocol-State      : Conform
Exceed-Count        : 0
Detec. Time Remain  : 5 seconds         Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True
-------------------------------------------------------------------------------
===============================================================================
```

After 30 seconds have passed, the “Detec. Time Remain” for arp and all-unspecified will simply read 0 (zero).

After a few minutes the log events will be collected indicating a non-conformance was seen.

```
*A:PE-1# show log log-id 15
===============================================================================
Event Log 15
===============================================================================
Description : (Not Specified)
Memory Log contents  [size=1024   next event=3  (not wrapped)]

2 2013/04/20 08:56:59.37 EDT WARNING: SECURITY #2067 Base DCPUPROT
"Non conformant network_if "int-pe1-to-tester" on fp 1/1 detected at 04/20/2013 08:52:28.
Policy "dcp-dynamic-policy-1". Policer="icmp"(dynamic). Excd count=2"

1 2013/04/20 08:56:59.37 EDT WARNING: SECURITY #2067 Base DCPUPROT
"Non conformant network_if "int-pe1-to-tester" on fp 1/1 detected at 04/20/2013 08:52:22.
Policy "dcp-dynamic-policy-1". Policer="igmp"(dynamic). Excd count=27"
```

Step 17. Stop the tester.

The dynamic policer detection timers will start counting down since they are no longer seeing violating packets.

```
*A:PE-1# show router interface "int-pe1-to-tester" dist-cpu-protection
===============================================================================
Interface "int-pe1-to-tester" (Router: Base)
===============================================================================
Distributed CPU Protection Policy :  dcp-dynamic-policy-1
-------------------------------------------------------------------------------
Statistics/Policer-State Information
===============================================================================
-------------------------------------------------------------------------------
Static Policer
-------------------------------------------------------------------------------
No entries found
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```

```
Local-Monitoring Policer
-------------------------------------------------------------------------------
Policer-Name        : local-mon
Card/FP             : 1/1               Policer-State       : Exceed
Protocols Mapped    : arp, icmp, igmp, all-unspecified
Exceed-Count        : 1097
All Dyn-Plcr Alloc. : True
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Dynamic-Policer (Protocol)
-------------------------------------------------------------------------------
Protocol(Dyn-Plcr)  : arp
Card/FP             : 1/1               Protocol-State      : Conform
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True

Protocol(Dyn-Plcr)  : icmp
Card/FP             : 1/1               Protocol-State      : Exceed
Exceed-Count        : 511
Detec. Time Remain  : 14 seconds        Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True

Protocol(Dyn-Plcr)  : igmp
Card/FP             : 1/1               Protocol-State      : Exceed
Exceed-Count        : 345550
Detec. Time Remain  : 18 seconds        Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True

Protocol(Dyn-Plcr)  : all-unspecified
Card/FP             : 1/1               Protocol-State      : Conform
Exceed-Count        : 0
Detec. Time Remain  : 0 seconds         Hold-Down Remain.   : none
Dyn-Policer Alloc.  : True
-------------------------------------------------------------------------------
===============================================================================
```

After 30 seconds there are no more violators.

```
*A:PE-1# tools dump security dist-cpu-protection violators enforcement interface card 1
===============================================================================
Distributed Cpu Protection Current Interface Enforcer Policer Violators
===============================================================================
Interface                      Policer/Protocol                    Hld Rem
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Violators on Slot-1 Fp-1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
[S]-Static [D]-Dynamic [M]-Monitor
-------------------------------------------------------------------------------
===============================================================================
```

The dynamic policer pool Hi-WaterMark for card 1 fp 1 shows 4 since the highest number of dynamic policers allocated at any one time on the card/fp was 4.

```
*A:PE-1# show card 1 fp 1 dist-cpu-protection
===============================================================================
Card : 1 Forwarding Plane(FP) : 1
===============================================================================
Dynamic Enforcement Policer Pool : 1000
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Statistics Information
-------------------------------------------------------------------------------
Dynamic-Policers Currently In Use      : 0
Hi-WaterMark Hit Count                 : 4
Hi-WaterMark Hit Time                  : 04/20/2013 08:52:22 UTC
Dynamic-Policers Allocation Fail Count : 0
-------------------------------------------------------------------------------
===============================================================================
```

A few minutes later the log events indicate that the flood has ended.

```
*A:PE-1# show log log-id 15
===============================================================================
Event Log 15
===============================================================================
Description : (Not Specified)
Memory Log contents  [size=1024   next event=5  (not wrapped)]

4 2013/04/20 09:01:59.39 EDT WARNING: SECURITY #2073 Base DCPUPROT
"Network_if "int-pe1-to-tester" on fp 1/1 newly conformant at 04/20/2013 08:58:39. Policy
"dcp-dynamic-policy-1". Policer="igmp"(dynamic). Excd count=345550"

3 2013/04/20 09:01:59.39 EDT WARNING: SECURITY #2073 Base DCPUPROT
"Network_if "int-pe1-to-tester" on fp 1/1 newly conformant at 04/20/2013 08:58:35. Policy
"dcp-dynamic-policy-1". Policer="icmp"(dynamic). Excd count=511"
```

Conclusion

Distributed CPU Protection (DCP) offers a powerful rate limiting function for control protocol traffic that is extracted from the data path and sent to the CPM.

This example has demonstrated how to configure DCP on an interface and what indications SR OS provides to the operator during a potential attack or misconfiguration.

DCP can also be deployed in scenarios where per-SAP-per-protocol rate limiting is useful, such as for subscriber management in a subscriber per-vlan scenario. A DCP policy can be assigned to an MSAP policy on a Broadband Network Gateway, for example, to limit traffic related to certain protocols and to discard certain protocols. When deployed in a subscriber management scenario, DCP can help isolate SAPs (subscribers) from each other and even isolate protocols from each other within an individual SAP (subscriber). Many of the same concepts introduced in this example are applicable when DCP is deployed in a subscriber management application.