



7450 Ethernet Service Switch
7750 Service Router
7950 Extensible Routing System
Virtualized Service Router
Releases up to 25.7.R2

System Management Advanced Configuration Guide for Classic CLI

3HE 20808 AAAE TQZZA
Edition: 01
October 2025

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Table of contents

List of figures.....4

Preface..... 5

Distributed CPU Protection..... 6

Event Handling System..... 27

SR OS NETCONF Server Basics..... 40

List of figures

Figure 1: Test Topology.....7

Figure 2: Count Traffic with DCP Policy Count..... 12

Figure 3: Limit Traffic with dcp-static-policy-1..... 15

Figure 4: Dynamic Policing – Local Monitor..... 22

Figure 5: Dynamic Policers Instantiated..... 22

Figure 6: Example topology..... 28

Figure 7: NETCONF client-server communication.....41

Preface

About This Guide

Each Advanced Configuration Guide is organized alphabetically and provides feature and configuration explanations, CLI descriptions, and overall solutions. The Advanced Configuration Guide chapters are written for and based on several Releases, up to 25.7.R2. The Applicability section in each chapter specifies on which release the configuration is based.

The Advanced Configuration Guides supplement the user configuration guides listed in the *7450 ESS*, *7750 SR*, and *7950 XRS Guide to Documentation*.

Audience

This manual is intended for network administrators who are responsible for configuring the routers. It is assumed that the network administrators have a detailed understanding of networking principles and configurations.

Distributed CPU Protection

This chapter describes Distributed CPU Protection (DCP) configurations.

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

This chapter was originally written for SR OS Release 11.0R1. The CLI in the current edition corresponds to Release 15.0.R1.

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 chapter is to familiarize the reader with the configuration and use of DCP. 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 1: Test Topology](#). A 10Gb Ethernet link is used between the tester and the router.

Figure 1: Test Topology



al_0279

1. The basic configuration of the MDA, port, interface and a security event log on the router is as follows.

```
configure
  card 4
    card-type iom3-xp-b
    mda 1
      mda-type m4-10gb-xp-xfp
      no shutdown
    exit
  exit
exit
```

```
configure
  port 4/1/1
    ethernet
    exit
    no shutdown
  exit
exit
```

```
configure
  router Base
    interface "int-R1-T1.4.4"
      address 192.168.40.1/24
      description "to tester T1, port 4.4"
      port 4/1/1
      no shutdown
    exit
  exit
exit
```

```
configure
  log
    log-id 15
      from security
      to memory 1024
      no shutdown
    exit
  exit
exit
```

This chapter was originally 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 DCP then the centralized CPU Protection should be disabled (for the purposes of this chapter) 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 configured to *max*.

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

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

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

```
configure
  system
    security
      dist-cpu-protection
        policy "dcp-policy-count" create
          description "Static policers with rate 0 for counting packets"
          static-policer "sp-arp" create
            description "static policer for ARP"
            rate packets 0 within 1
          exit
          static-policer "sp-icmp" create
            description "static policer for ICMP"
            rate packets 0 within 1
          exit
          static-policer "sp-igmp" create
            description "static policer for IGMP"
            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
      exit
    exit
  exit
```

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

- The policy contains three static policers: *sp-arp*, *sp-icmp* and *sp-igmp*. These policers are then used by the three configured protocols that are part of the policy: *arp*, *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". See the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR 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 chapter.
- 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.

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

```
*A:R1# configure router interface "int-R1-T1.4.4"
*A:R1>config>router>if# dist-cpu-protection "dcp-policy-count"
```

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 94% and no task groups at more than 2.5% capacity usage. Future example output from this show routine will be snipped to only show relevant and interesting lines.

```
*A:R1# show system cpu
```

```
=====
CPU Utilization (Sample period: 1 second)
=====
```

Name	CPU Time (uSec)	CPU Usage	Capacity Usage
-----	-----	-----	-----
BFD	60	~0.00%	~0.00%
BGP	30,892	0.34%	0.62%
BGP PE-CE	0	0.00%	0.00%
CALLTRACE	5,210	0.05%	0.51%
CFLOWD	5,128	0.05%	0.51%
Cards & Ports	39,591	0.44%	0.95%
DHCP Server	35	~0.00%	~0.00%
ETH-CFM	4,584	0.05%	0.46%
HQoS Algorithm	0	0.00%	0.00%
HQoS Statistics	0	0.00%	0.00%
ICC	1,225	0.01%	0.12%
IGMP/MLD	1,080	0.01%	0.10%
IMSI Db Appl	258	~0.00%	0.01%
IOM	0	0.00%	0.00%
IP Stack	56,965	0.63%	0.51%
IS-IS	51,342	0.57%	0.60%
ISA	16,173	0.18%	0.55%
LDP	31,118	0.34%	0.55%
Logging	53	~0.00%	~0.00%
MBUF	0	0.00%	0.00%
MCS	536	~0.00%	0.04%
MPLS/RSVP	8,915	0.09%	0.57%
MSCP	0	0.00%	0.00%
MSDP	0	0.00%	0.00%
Management	18,039	0.20%	0.73%

OAM	12,422	0.13%	0.48%
OSPF	118,279	1.32%	0.58%
OpenFlow	1,037	0.01%	0.01%
PIM/L2Mcast	0	0.00%	0.00%
PKI	272	~0.00%	0.02%
PTP	71	~0.00%	~0.00%
RIP	0	0.00%	0.00%
RTM/Policies	0	0.00%	0.00%
Redundancy	10,321	0.11%	0.63%
SNMP Daemon	0	0.00%	0.00%
Security	0	0.00%	0.00%
Services	8,775	0.09%	0.52%
Stats	0	0.00%	0.00%
Subscriber Mgmt	6,439	0.07%	0.14%
System	94,262	1.05%	2.28%
Traffic Eng	0	0.00%	0.00%
VRRP	1,942	0.02%	0.12%
WEB Redirect	95	~0.00%	~0.00%

Total	8,936,439	100.00%	
Idle	8,411,320	94.12%	
Usage	525,119	5.87%	
Busiest Core Utilization	79,550	8.01%	
=====			
*A:R1#			

The DCP feature is reporting no violations for interfaces on card 4.

```
*A:R1# tools dump security dist-cpu-protection violators enforcement interface card 4
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface                Policer/Protocol          Hld Rem
-----
Violators on Slot-4 Fp-1
-----
[S]-Static [D]-Dynamic [M]-Monitor
=====
*A:R1#
```

There are no security log events.

```
*A:R1# show log log-id 15
=====
Event Log 15
=====
Description : (Not Specified)
Memory Log contents [size=1024  next event=1  (not wrapped)]
*A:R1#
```

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

```
*A:R1# show router interface "int-R1-T1.4.4" dist-cpu-protection
=====
Interface "int-R1-T1.4.4" (Router: Base)
=====
Distributed CPU Protection Policy :  dcp-policy-count
```

```

-----
Statistics/Policer-State Information
=====
-----
Static Policer
-----
Policer-Name      : sp-arp
Card/FP           : 4/1          Policer-State      : Conform
Protocols Mapped  : arp
Exceed-Count      : 0
Detec. Time Remain : 0 seconds   Hold-Down Remain.  : none

Policer-Name      : sp-icmp
Card/FP           : 4/1          Policer-State      : Conform
Protocols Mapped  : icmp
Exceed-Count      : 0
Detec. Time Remain : 0 seconds   Hold-Down Remain.  : none

Policer-Name      : sp-igmp
Card/FP           : 4/1          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
-----
=====
*A:R1#

```

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 4/1/1 can be seen in the output of show port 4/1/1 as the configured address.
 - Set the IP destination address to 192.168.40.1 and the IP source address to 192.168.40.2.
- IGMP:

- Set the MAC destination address equal to the MAC address of the receiving port. The MAC address of port 4/1/1 can be seen in the output of show port 4/1/1 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 2: Count Traffic with DCP Policy Count



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

```
*A:R1# tools dump security dist-cpu-protection violators enforcement interface card 4
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface                               Policer/Protocol                        Hld Rem
-----
Violators on Slot-4 Fp-1
-----
int-R1-T1.4.4                           sp-arp                                [S] none
int-R1-T1.4.4                           sp-icmp                              [S] none
int-R1-T1.4.4                           sp-igmp                              [S] none
-----
[S]-Static [D]-Dynamic [M]-Monitor
=====
*A:R1#
```

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:R1# show router interface "int-R1-T1.4.4" dist-cpu-protection
```

```

=====
Interface "int-R1-T1.4.4" (Router: Base)
=====
Distributed CPU Protection Policy : dcp-policy-count

-----
Statistics/Policer-State Information
=====
-----
Static Policer
-----
Policer-Name      : sp-arp
Card/FP           : 4/1
Protocols Mapped  : arp
Exceed-Count      : 263
Detec. Time Remain : 29 seconds
Policer-State     : Exceed
Hold-Down Remain. : none

Policer-Name      : sp-icmp
Card/FP           : 4/1
Protocols Mapped  : icmp
Exceed-Count      : 525
Detec. Time Remain : 29 seconds
Policer-State     : Exceed
Hold-Down Remain. : none

Policer-Name      : sp-igmp
Card/FP           : 4/1
Protocols Mapped  : igmp
Exceed-Count      : 1050
Detec. Time Remain : 29 seconds
Policer-State     : Exceed
Hold-Down Remain. : none
-----
Local-Monitoring Policer
-----
No entries found
-----
Dynamic-Policer (Protocol)
-----
No entries found
=====
*A:R1#

```

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:

```

configure
  system
    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
                                exit
                                exit
                                exit
                                exit

```

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

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

Now assign the policy to the test interface:

```

*A:R1# configure router interface "int-R1-T1.4.4"
                                dist-cpu-protection "dcp-static-policy-1"

```

```

*A:R1# 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-R1-T1.4.4
-----
Number of interfaces : 1
=====
*A:R1#

```

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

Figure 3: Limit Traffic with dcp-static-policy-1



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

```
*A:R1# tools dump security dist-cpu-protection violators enforcement interface card 4
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface                               Policer/Protocol                        Hld Rem
-----
Violators on Slot-4 Fp-1
-----
int-R1-T1.4.4                           sp-igmp                                [S] none
-----
[S]-Static [D]-Dynamic [M]-Monitor
=====
*A:R1#
```

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:R1# show log log-id 15
=====
Event Log 15
=====
Description : (Not Specified)
Memory Log contents [size=1024 next event=2 (not wrapped)]

1 2017/04/27 09:47:53.21 CEST WARNING: SECURITY #2066 Base DCPUPROT
"Non conformant network_if "int-R1-T1.4.4" on fp 4/1 detected at 04/27/2017 09:47:07.
Policy "dcp-static-policy-1". Policer="sp-igmp"(static). Excd count=411"
*A:R1#
```

```
*A:R1# show router interface "int-R1-T1.4.4" dist-cpu-protection
=====
Interface "int-R1-T1.4.4" (Router: Base)
=====
```

Distributed CPU Protection Policy : dcp-static-policy-1

Statistics/Policer-State Information
=====

Static Policer

Policer-Name	: sp-arp	Policer-State	: Conform
Card/FP	: 4/1		
Protocols Mapped	: arp		
Exceed-Count	: 0		
Detec. Time Remain	: 0 seconds	Hold-Down Remain.	: none
Policer-Name	: sp-icmp	Policer-State	: Conform
Card/FP	: 4/1		
Protocols Mapped	: icmp		
Exceed-Count	: 0		
Detec. Time Remain	: 0 seconds	Hold-Down Remain.	: none
Policer-Name	: sp-igmp	Policer-State	: Exceed
Card/FP	: 4/1		
Protocols Mapped	: igmp		
Exceed-Count	: 640151		
Detec. Time Remain	: 29 seconds	Hold-Down Remain.	: none

Local-Monitoring Policer

No entries found

Dynamic-Policer (Protocol)

No entries found

=====

*A:R1#

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

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

*A:R1# show system cpu

=====

CPU Utilization (Sample period: 1 second)

Name	CPU Time (uSec)	CPU Usage	Capacity Usage
BFD	160	~0.00%	0.01%
---snip---			
IGMP/MLD	1,795	0.02%	0.09%
---snip---			
Stats	0	0.00%	0.00%
Subscriber Mgmt	5,661	0.06%	0.09%
System	77,189	0.86%	1.23%
Traffic Eng	0	0.00%	0.00%
VRRP	1,679	0.01%	0.09%
WEB Redirect	83	~0.00%	~0.00%

```
Total                8,936,280        100.00%
  Idle                8,420,519        94.22%
  Usage               515,761         5.77%
Busiest Core Utilization 78,908         7.94%
=====
```

*A:R1#

10. Remove the DCP policy from the interface and see the capacity usage going up for the IGMP task group.

```
*A:R1# configure router interface "int-R1-T1.4.4" no dist-cpu-protection
*A:R1#
```

```
*A:R1# show system cpu
```

```
=====
CPU Utilization (Sample period: 1 second)
=====
Name                                CPU Time      CPU Usage      Capacity
                                (uSec)                               Usage
-----
BFD                                191           ~0.00%         0.01%
---snip---
ICC                                1,332         0.01%         0.13%
IGMP/MLD                        78,184      0.87%       7.78%
IMSI Db Appl                       249          ~0.00%        ~0.00%
IOM                                 0            0.00%         0.00%
IP Stack                          183,448       2.05%         9.16%
IS-IS                             49,866       0.55%         0.56%
---snip---
Subscriber Mgmt                     8,153         0.09%         0.14%
System                           144,738       1.62%         6.11%
Traffic Eng                         0            0.00%         0.00%
VRRP                               2,396         0.02%         0.15%
WEB Redirect                        83           ~0.00%        ~0.00%
-----
Total                            8,925,786    100.00%
  Idle                            8,123,698    91.01%
  Usage                           802,088      8.98%
Busiest Core Utilization          170,131     17.15%
=====
```

*A:R1#

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

```
*A:R1# show system cpu
```

```
=====
CPU Utilization (Sample period: 1 second)
=====
Name                                CPU Time      CPU Usage      Capacity
                                (uSec)                               Usage
-----
BFD                                158           ~0.00%         0.01%
---snip---
ICC                                1,061         0.01%         0.10%
IGMP/MLD                        398,106     4.44%       39.99%
IMSI Db Appl                       142          ~0.00%        ~0.00%
IOM                                 0            0.00%         0.00%
IP Stack                          648,378       7.24%         43.68%
IS-IS                             59,623       0.66%         0.65%
---snip---
```

```
Subscriber Mgmt          7,308          0.08%          0.13%
System                  364,124        4.06%          28.73%
Traffic Eng              0           0.00%          0.00%
VRRP                    2,156          0.02%          0.12%
WEB Redirect            117          ~0.00%          0.01%
-----
Total                    8,951,453      100.00%
  Idle                  7,114,732       79.48%
  Usage                 1,836,721       20.51%
Busiest Core Utilization 590,342       59.35%
=====
*A:R1#
```

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

```
*A:R1# configure router interface "int-R1-T1.4.4"          dist-cpu-
protection "dcp-static-policy-1"
```

```
*A:R1# show system cpu

=====
CPU Utilization (Sample period: 1 second)
=====
Name                                CPU Time      CPU Usage      Capacity
                                (uSec)                               Usage
-----
BFD                                72            ~0.00%          ~0.00%
---snip---
ICC                                1,000          0.01%          0.10%
IGMP/MLD                           2,149          0.02%          0.12%
IMSI Db Appl                        166           ~0.00%          ~0.00%
IOM                                  0             0.00%          0.00%
IP Stack                           60,407         0.67%          0.55%
IS-IS                              50,966         0.57%          0.58%
---snip---
Subscriber Mgmt                     5,847          0.06%          0.09%
System                             96,233          1.07%          2.23%
Traffic Eng                          0             0.00%          0.00%
VRRP                                2,338          0.02%          0.15%
WEB Redirect                         94            ~0.00%          ~0.00%
-----
Total                              8,925,256      100.00%
  Idle                             8,396,016       94.07%
  Usage                            529,240         5.92%
Busiest Core Utilization            81,620         8.22%
=====
*A:R1#
```

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:R1# tools dump security dist-cpu-protection violators enforcement interface card 4
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface                        Policer/Protocol                                Hld Rem
-----
Violators on Slot-4 Fp-1
-----
```

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

```
*A:R1#
```

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

```
*A:R1# show log log-id 15
```

```
Event Log 15
```

```
Description : (Not Specified)
```

```
Memory Log contents [size=1024 next event=4 (not wrapped)]
```

```
3 2017/04/27 10:02:53.25 CEST WARNING: SECURITY #2072 Base DCPUPROT
"Network_if "int-R1-T1.4.4" on fp 4/1 newly conformant at 04/27/2017 10:02:04. Policy
"dcp-static-policy-1". Policer="sp-igmp"(static). Excd count=227756"
```

```
---snip---
```

```
*A:R1#
```

The interface DCP details show all policers as conformant.

```
*A:R1# show router interface "int-R1-T1.4.4" dist-cpu-protection
```

```
Interface "int-R1-T1.4.4" (Router: Base)
```

```
Distributed CPU Protection Policy : dcp-static-policy-1
```

```
Statistics/Policer-State Information
```

```
Static Policer
```

```
Policer-Name      : sp-arp
Card/FP           : 4/1          Policer-State      : Conform
Protocols Mapped  : arp
Exceed-Count      : 0
Detec. Time Remain : 0 seconds   Hold-Down Remain. : none
```

```
Policer-Name      : sp-icmp
Card/FP           : 4/1          Policer-State      : Conform
Protocols Mapped  : icmp
Exceed-Count      : 0
Detec. Time Remain : 0 seconds   Hold-Down Remain. : none
```

```
Policer-Name      : sp-igmp
Card/FP           : 4/1          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
-----
```

```
*A:R1#
```

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, 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.

14. The next scenario explored in this chapter 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:R1# configure card 4 fp dist-cpu-protection dynamic-enforcement-policer-pool 1000
*A:R1#
```

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:R1# show card 4 fp 1 dist-cpu-protection
```

```
=====
Card : 4 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                  : 04/27/2017 10:08:24 UTC
Dynamic-Policers Allocation Fail Count : 0
-----
```

```
*A:R1#
```

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:

```
configure
  system
    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
        exit
      exit
    exit
  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
exit
exit
exit
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:R1# configure router interface "int-R1-T1.4.4"
      dist-cpu-protection "dcp-dynamic-policy-1"
*A:R1#
```

15. Configure the tester to send:

- 1pps of ARP

- 4pps of ICMP
 - 1000pps of IGMP
- Start the tester.


Figure 4: Dynamic Policing – Local Monitor



In [Figure 4: Dynamic Policing – Local Monitor](#), the dynamic policers have not been instantiated yet.

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 5: Dynamic Policers Instantiated



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

```
*A:R1# tools dump security dist-cpu-protection violators enforcement interface card 4
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface          Policer/Protocol      Hld Rem
-----
Violators on Slot-4 Fp-1
int-R1-T1.4.4      icmp                 [D] none
int-R1-T1.4.4      igmp                 [D] none
[S]-Static [D]-Dynamic [M]-Monitor
```

```
=====
*A:R1#
```

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:R1# show router interface "int-R1-T1.4.4" dist-cpu-protection
```

```
=====
Interface "int-R1-T1.4.4" (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            : 4/1          Policer-State       : Exceed
Protocols Mapped   : arp, icmp, igmp, all-unspecified
Exceed-Count       : 1249
All Dyn-Plcr Alloc. : True
-----
```

```
Dynamic-Policer (Protocol)
```

```
-----
Protocol(Dyn-Plcr) : arp
Card/FP            : 4/1          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            : 4/1          Protocol-State    : Exceed
Exceed-Count       : 72
Detec. Time Remain : 26 seconds    Hold-Down Remain. : none
Dyn-Policer Alloc. : True
-----
```

```
Protocol(Dyn-Plcr) : igmp
Card/FP            : 4/1          Protocol-State    : Exceed
Exceed-Count       : 56190
Detec. Time Remain : 29 seconds    Hold-Down Remain. : none
Dyn-Policer Alloc. : True
-----
```

```
Protocol(Dyn-Plcr) : all-unspecified
Card/FP            : 4/1          Protocol-State    : Conform
Exceed-Count       : 0
Detec. Time Remain : 0 seconds    Hold-Down Remain. : none
Dyn-Policer Alloc. : True
-----
```

```
=====
*A:R1#
```

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:R1# show log log-id 15

=====
Event Log 15
=====
Description : (Not Specified)
Memory Log contents [size=1024  next event=10  (not wrapped)]

9 2017/04/27 10:22:53.32 CEST WARNING: SECURITY #2067 Base DCPUPROT
"Non conformant network_if "int-R1-T1.4.4" on fp 4/1 detected at 04/27/2017 10:18:37.
Policy "dcp-dynamic-policy-1". Policer="icmp"(dynamic). Excd count=2"

8 2017/04/27 10:22:53.32 CEST WARNING: SECURITY #2067 Base DCPUPROT
"Non conformant network_if "int-R1-T1.4.4" on fp 4/1 detected at 04/27/2017 10:18:30.
Policy "dcp-dynamic-policy-1". Policer="igmp"(dynamic). Excd count=80"

---snip---

*A:R1#
```

17. Stop the tester.

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

```
*A:R1# show router interface "int-R1-T1.4.4" dist-cpu-protection

=====
Interface "int-R1-T1.4.4" (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            : 4/1          Policer-State       : Exceed
Protocols Mapped   : arp, icmp, igmp, all-unspecified
Exceed-Count       : 5072
All Dyn-Plcr Alloc. : True
-----
Dynamic-Policer (Protocol)
-----
Protocol(Dyn-Plcr) : arp
Card/FP            : 4/1          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           : 4/1           Protocol-State      : Exceed
Exceed-Count      : 482
Detec. Time Remain : 0 seconds    Hold-Down Remain.   : none
Dyn-Policer Alloc. : True

Protocol(Dyn-Plcr) : igmp
Card/FP           : 4/1           Protocol-State      : Exceed
Exceed-Count      : 326409
Detec. Time Remain : 0 seconds    Hold-Down Remain.   : none
Dyn-Policer Alloc. : True

Protocol(Dyn-Plcr) : all-unspecified
Card/FP           : 4/1           Protocol-State      : Conform
Exceed-Count      : 0
Detec. Time Remain : 0 seconds    Hold-Down Remain.   : none
Dyn-Policer Alloc. : True
=====

```

*A:R1#

After 30 seconds there are no more violators.

```

*A:R1# tools dump security dist-cpu-protection violators enforcement interface card 4
=====
Distributed Cpu Protection Current Interface Enforcer Policer Violators
=====
Interface                Policer/Protocol          Hld Rem
-----
Violators on Slot-4 Fp-1
-----
[S]-Static [D]-Dynamic [M]-Monitor
=====

```

*A:R1#

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:R1# show card 4 fp 1 dist-cpu-protection
=====
Card : 4 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                  : 04/27/2017 10:10:34 UTC
Dynamic-Policers Allocation Fail Count : 0
=====

```

*A:R1#

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

```
*A:R1# show log log-id 15

=====
Event Log 15
=====
Description : (Not Specified)
Memory Log contents [size=1024  next event=12  (not wrapped)]

11 2017/04/27 10:27:53.34 CEST WARNING: SECURITY #2073 Base DCPUPROT
"Network_if "int-R1-T1.4.4" on fp 4/1 newly conformant at 04/27/2017 10:24:27. Policy
"dcp-dynamic-policy-1". Policer="igmp"(dynamic). Excd count=326409"

10 2017/04/27 10:27:53.34 CEST WARNING: SECURITY #2073 Base DCPUPROT
"Network_if "int-R1-T1.4.4" on fp 4/1 newly conformant at 04/27/2017 10:24:27. Policy
"dcp-dynamic-policy-1". Policer="icmp"(dynamic). Excd count=482"

---snip---

*A:R1#
```

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 chapter 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 chapter are applicable when DCP is deployed in a subscriber management application.

Event Handling System

This chapter provides information about event handling systems (EHS).

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

This chapter was initially written for SR OS Release 13.0.R3. The CLI in the current edition is based on SR OS Release 23.7.R2.

SR OS Release 13.0.R1 introduced event handling system (EHS).

SR OS Release 14.0.R4 introduced EHS script enhanced capabilities, such as static variables, advanced syntax (shell scripting commands), and so on. The examples in this chapter do not include these enhancements,

Overview

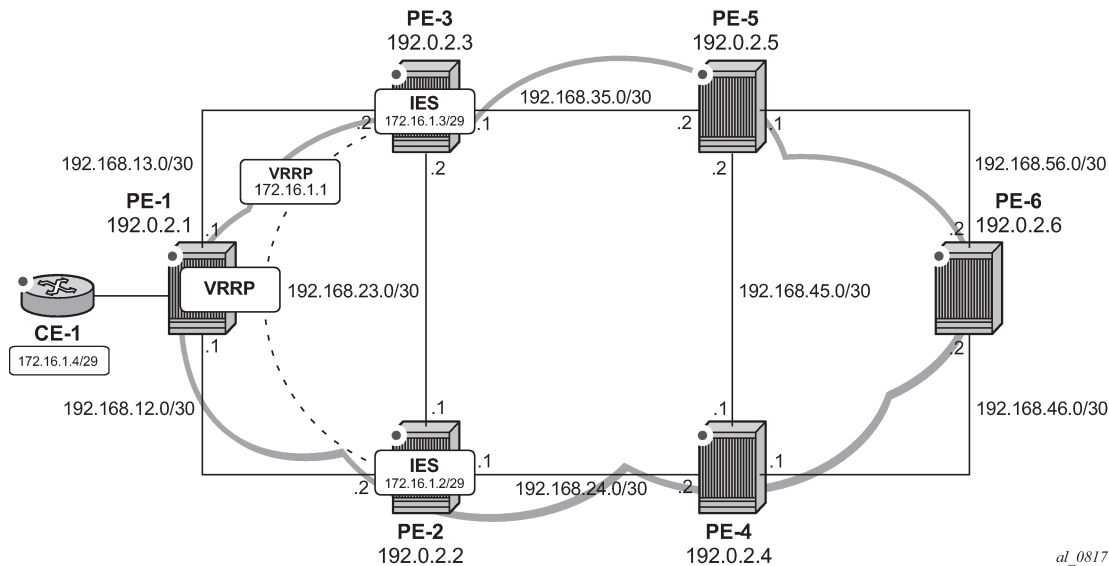
The event handling system (EHS) in SR OS allows operators to configure user-defined actions defined in CLI scripts that the router executes in response to an event. The event is referred to as the trigger, where the trigger can be all or part of any event message generated by the event-control framework. The user-defined action is controlled by the script-control function. This script-control function references one or more scripts that are able to execute any command available in CLI when the trigger event occurs.

This feature allows for customized automated event management based on specific operator requirements.

Configuration

The topology shown in [Figure 6: Example topology](#) provides an example of an EHS configuration. All routers within the example topology participate in the same IS-IS level-2 area and run LDP. All routers are BGP speakers and form part of autonomous system 64496, exchanging routes for IPv4 address family only.

Figure 6: Example topology



PE-1 has a CE router connected (CE-1) that is indexed into a VPLS service. This VPLS has spoke-SDPs to an IES instance on both PE-2 and PE-3, which provide a redundant default gateway to CE-1 using the virtual router redundancy protocol (VRRP). The subnet used for this redundant gateway connectivity between PE-2 and PE-3 is 172.16.1.0/29. The configuration at PE-3 is shown in the following output. The configuration at PE-2 is similar; the exception being IP addressing and VRRP priority, which is 254.

```
# on PE-3:
configure
service
  ies 1 name "IES-1" customer 1 create
  interface "redundant-interface" create
    address 172.16.1.3/29
    ip-mtu 1500
    vrrp 1
      backup 172.16.1.1
      priority 253
      ping-reply
    exit
  spoke-sdp 31:1 create
    no shutdown
  exit
exit
no shutdown
exit
```

The objective is to ensure that both upstream and downstream traffic are always routed through the same PE router. That is, if PE-3 is VRRP primary, it will attract upstream traffic from CE-1 using the VRRP virtual IP/MAC. At the same time, PE-3 should also attract the downstream traffic destined toward CE-1. Having both upstream and downstream traffic transit through the same PE router, simplifies troubleshooting, QoS configuration, and reconciliation of ingress/egress statistics.

In normal operation, PE-2 is the VRRP master and advertises the BGP prefix 172.16.1.0/29 with a local preference of 100 (default value). Similarly, PE-3 is the VRRP backup and advertises the BGP prefix 172.16.1.0/29 with a local preference of 50, using the BGP export policy "redundant-interface":

```
# on PE-3:
configure
router Base
  policy-options
  begin
    prefix-list "172.16.1.0/29"
    prefix 172.16.1.0/29 exact
  exit
  policy-statement "redundant-interface"
  entry 10
    from
    prefix-list "172.16.1.0/29"
  exit
  to
    protocol bgp
  exit
  action accept
  origin igp
  local-preference 50
  exit
exit
exit
commit
```

Therefore, upstream and downstream traffic normally transit through PE-2. The following shows that the VRRP instance on "redundant-interface" on PE-3 is backup.

```
*A:PE-3# show router vrrp instance
```

```
=====
VRRP Instances
=====
```

Interface Name	VR Id IP	Own No	Adm Opr	State Pol Id	Base Pri InUse Pri	Msg Int Inh Int
redundant-interface	1 IPv4	No	Up Up	Backup n/a	253 253	1 No

```
Backup Addr: 172.16.1.1
-----
Instances : 1
=====
```

When PE-3 is backup, it advertises the prefix 172.16.1.0/29 with a local preference of 50, as follows:

```
*A:PE-3# show router bgp routes 172.16.1.0/29 hunt | match expression "Network|Nexthop|To|Local
Pref"
```

Network	: 172.16.1.0/29	
Nexthop	: 192.0.2.2	
Res. Nexthop	: 192.168.23.1	
Local Pref.	: 100	Interface Name : int-PE-3-PE-2
Network	: 172.16.1.0/29	
Nexthop	: 192.0.2.3	
To	: 192.0.2.6	
Res. Nexthop	: n/a	
Local Pref.	: 50	Interface Name : NotAvailable

When PE-3 transitions from backup to primary, it must modify its local preference attribute for prefix 172.16.1.0/29 to a value of 150 to attract downstream traffic destined toward CE-1. Similarly, when PE-3 reverts to backup, it must advertise the prefix with a local preference of 50.

Script control

The first step in configuring event handling is to configure a script containing the CLI commands to be executed when the event is triggered. This script can be stored locally on the compact flash, or it can be stored off-node at a defined remote URL, where it can be accessed using FTP or TFTP. When the script is stored locally on the compact flash and the router is equipped with redundant CPMs, the script must be manually saved on the same compact flash on both CPMs, because it is not synchronized automatically.

The first requirement is to modify the local preference of the prefix 172.16.1.0/29 to 150 on transition to VRRP master. The script, which in this example is held locally on CF3:/, therefore contains the following commands (where the policy-statement, redundant-interface, is the name of the export policy used to advertise the 172.16.1.0/29 prefix):

```
*A:PE-3>file cf3:\ # type cf3:vrrp-master.txt
File: vrrp-master.txt
-----
exit all
configure router policy-options
begin
policy-statement redundant-interface
entry 10
action accept
local-preference 150
exit
exit
exit
commit
exit all

=====
```

There is no syntax checking when the script file is created; instead, the script will fail with a command error. Also, transactional CLI (for example the **edit** command) cannot be used in the script, and will fail with a command error.

Within the **system script-control** context, the script is assigned a name and reference is made to its location. It is then put in the no shutdown state. When the script has been defined, a **script-policy** is configured that calls the previously configured script. The script-policy also specifies a location and filename for a results file that records the successful or unsuccessful conclusion of each script run and each command executed during that run. Each time the script is run, the results are recorded in a file with the name specified for results, followed by an underscore and the date and time when the script was run. A results file must be specified in order for the script to successfully run. The results file can be on the local compact flash, or a remote URL can be specified. As with the script, the script-policy must also be put in the no shutdown state.

```
# on PE-3:
configure
  system
    script-control
      script "vrrp-master-script"
```

```

        location "cf3:/vrrp-master.txt"
        no shutdown
    exit
    script-policy "vrrp-master-policy"
        results "cf3:/script-results.txt"
        script "vrrp-master-script"
        max-completed 4
        expire-time 3600
        lifetime forever
        no shutdown
    exit
exit

```

The optional **lifetime** command specifies the maximum time that the script may run. The **max-completed** command specifies the maximum number of script run history status entries to be retained. An optional **expire-time** command specifies the maximum time that the system keeps the run history status (default is 1 h). The system maintains the script run history table, which has a maximum size of 255 entries. Entries are removed from this table when the max-completed or expire-time thresholds are crossed. If the table reaches the maximum value, subsequent script launch requests are not run until older run history entries expire (due to expire-time), or entries are manually cleared. To manually clear entries, the following command is used:

```
clear system script-control script-policy completed <script-policy-name>
```

The script run history status information can be viewed using the following command (in this case, after one successful run of the corresponding script) :

```
*A:PE-3# show system script-control script-policy "vrrp-master-policy"
```

Script-policy Information

```

Script-policy          : vrrp-master-policy
Script-policy Owner    : TiMOS CLI
Administrative status  : enabled
Operational status    : enabled
Script                : vrrp-master-script
Script owner          : TiMOS CLI
Python script         : N/A
Source location        : cf3:/vrrp-master.txt
Results location       : cf3:/script-results.txt
Max running allowed    : 1
Max completed run histories : 4
Max lifetime allowed   : 248d 13:13:56 (21474836 seconds)
Completed run histories : 1
Executing run histories : 0
Initializing run histories : 0
Max time run history saved : 0d 01:00:00 (3600 seconds)
Script start error     : N/A
Python script start error : N/A
Last change            : 2023/09/13 07:43:55 UTC
Max row expire time    : never
Last application       : event-script
Last auth. user account : not-specified

```

Script Run History Status Information

Script Run #1

```
Start time : 2023/09/13 07:45:35 UTC
```

```

End time       : 2023/09/13 07:45:35 UTC
Elapsed time   : 0d 00:00:00           Lifetime      : 0d 00:00:00
State          : terminated            Run exit code : noError
Result time    : 2023/09/13 07:45:35 UTC
Keep history   : 0d 00:59:29
Error time     : never
Source file    : cf3:/vrrp-master.txt
Results file   : cf3:/script-results.txt_20230913-074534-UTC.833059.out
Run exit       : Success
Error          : N/A
Application    : event-script          Auth. user ac*: not-specified
* indicates that the corresponding row element may have been truncated.
=====

```

Event handler

The second step in configuring event handling is to assign actions to be performed as a result of the trigger event. These actions are typically one or more configured scripts defined as entries in an action list. In the following output, the event handler is assigned the name event-handler-1, and the action list consists of a single entry. This entry calls the previously configured script policy vrrp-master-policy (which in turn references the previously defined script vrrp-master-script). If multiple actions are required based on a single event trigger, they can be configured in the action list with subsequent entries, which are run in sequence (up to 1500 action list entries are supported).

For this example, only a single entry is required; therefore, there is a one to one relationship between the event handler and the action list entry. Both the entry within the action list and the handler should be put in the no shutdown state.

```

# on PE-3:
configure
  log
    event-handling
      handler "event-handler-1"
        action-list
          entry 10
            script-policy "vrrp-master-policy"
            no shutdown
          exit
        exit
      no shutdown
    exit
  exit

```

Event trigger

The final step in configuring event handling is to configure the event trigger. The event trigger defines the event that triggers the running of the script. The event trigger is based on any event generated by the event-control framework, and can match against the application and event number (event_id). Log filters can also be used to match against specific events using the subject and/or message fields. Regular expressions can be used where required. EHS will not use any message that is suppressed through event-control configuration, or any event message that is throttled.

The general format for an event in an event log is as follows:

```

nnnn YYYY/MM/DD HH:MM:SS.SS Zone <severity>:<application> #

```



```
<event_id> <router-name> <subject> description
```

Where:

nnnn	The log entry sequence number
YYYY/MM/DD	The UTC date stamp for the log entry:
	YYYY - Year
	MM - Month
	DD - Date
HH:MM:SS.SS	The UTC time stamp for the event
	HH - Hours (24 hour format)
	MM - Minutes
	SS.SS - Seconds
TZONE	The timezone
<severity>	The severity level name of the event
<application>	The application generating the log message
<event_id>	The application's event ID number for the event
<subject>	The subject/affected object for the event
<message>	A textual description of the event

In the example, the following event message is generated when PE-3 becomes VRRP primary:

```
152 2023/09/13 07:44:50.432 UTC MINOR: VRRP #2001 Base Becoming Master
"VRRP virtual router instance 1 on interface redundant-interface
(primary address 172.16.1.3) changed state to master"
```

Therefore, the event-trigger configuration is based on an application of VRRP and an event number of 2001 (vrrptrapNewMaster). In the following snippet, vrrp 2001 is configured as the event. The trigger entry is defined as 1, and in this example, there is only one trigger event. Up to 1500 trigger entries can be included, each of which can act as a potential trigger event. The trigger entry also references the previously configured event-handler-1. (Recall that the event handler references the script control, which in turn references the script that should be run.)

```
# on PE-3:
configure
  log
    event-trigger
      event "vrrp" 2001
        trigger-entry 1
          event-handler "event-handler-1"
          log-filter 1
          no shutdown
        exit
      no shutdown
    exit
  exit
```

Finally, there is a reference to log-filter 1. Without more explicit filtering, event handling will be triggered on any event with the application of VRRP and event number 2001. There may be multiple VRRP instances running on this router, but the requirement is that event handling should only be triggered when the VRRP instance running on redundant-interface transitions to master at PE-3. Therefore, log filter 1 is used to define a more explicit match using the message field, which contains an explicit reference to the interface. Both the trigger entry and the event handler should be put in the no shutdown state.

```
configure
  log
    filter 1 name "itf 172.31.1.3 becomes primary"
      default-action drop
```

```

        entry 10 name "newPrimary"
            action forward
            match
                message eq pattern "interface redundant-interface
                                   (primary address 172.16.1.3) changed state to master"
            exit
        exit
    exit

```

The configuration of the example event handling for the failure event (PE-3 transitions to VRRP primary) is now complete. By shutting down the spoke-SDP between PE-1 and PE-2, it is possible to simulate a failure event where the VRRP message path is broken. Therefore, four events are generated.

- The first indicates that PE-3 has become VRRP master for the interface named redundant-interface.
- The second indicates that EHS handler event-handler-1 was invoked by a CLI user.
- The third indicates that a script file has initiated an attempt to execute CLI commands contained in script file vrrp-master.txt.
- The fourth indicates that the attempt to execute those CLI commands was successful.

```

154 2023/09/13 07:45:34.832 UTC MINOR: VRRP #2001 Base Becoming Master
"VRRP virtual router instance 1 on interface redundant-interface
(primary address 172.16.1.3) changed state to master"

155 2023/09/13 07:45:34.832 UTC MINOR: SYSTEM #2069 Base EHS script
"Ehs handler : "event-handler-1" with the description : "" was invoked by the
cli-user account "not-specified".

156 2023/09/13 07:45:34.836 UTC MAJOR: SYSTEM #2052 Base CLI 'exec'
"A CLI user has initiated an 'exec' operation to process the commands in the SROS CLI
file cf3:/vrrp-master.txt"

157 2023/09/13 07:45:34.841 UTC MAJOR: SYSTEM #2053 Base CLI 'exec'
"The CLI user initiated 'exec' operation to process the commands in the SROS CLI
file cf3:/vrrp-master.txt has completed with the result of success"

```

A successful script run shows the commands contained in the script, followed by an indication that the commands were executed.

```

*A:PE-3>file cf3:\ # type script-results.txt_20230913-074534-UTC.833059.out
File: script-results.txt_20230913-074534-UTC.833059.out
-----
exit all
configure router policy-options
begin
policy-statement redundant-interface
entry 10
action accept
local-preference 150
exit
exit
exit
commit
exit all
Executed 14 lines in 0.0 seconds from file "cf3:/vrrp-master.txt"

=====

```

The following output confirms that PE-3 is VRRP primary:

```
*A:PE-3# show router vrrp instance

=====
VRRP Instances
=====
Interface Name          VR Id Own Adm State      Base Pri  Msg Int
                        IP      Opr Pol Id      InUse Pri  Inh Int
-----
redundant-interface    1      No  Up  Master    253      1
                        IPv4    Up  n/a      253      No
Backup Addr: 172.16.1.1
-----
Instances : 1
=====
```

Also, the local preference attribute for prefix 172.16.1.0/29 has changed to a value of 150. The result of this action is that PE-3 will now be the transit router for both upstream and downstream traffic.

```
*A:PE-3# show router bgp routes 172.16.1.0/29 hunt | match expression "Network|Nexthop|To|Local
Pref"
Network      : 172.16.1.0/29
Nexthop      : 192.0.2.3
Res. Nexthop : Unresolved
Local Pref. : 150                      Interface Name : NotAvailable
---snip---
Network      : 172.16.1.0/29
Nexthop      : 192.0.2.3
To           : 192.0.2.6
Res. Nexthop : n/a
Local Pref. : 150                      Interface Name : NotAvailable
```

The event handler indicates that the referenced script was triggered and run using the command shown in the following output. The Handler Action-List Entry Execution Statistics window provides statistics on the number of times an action (script) was queued to run, and the number of times an error was experienced, both during launch and due to a non-operational admin status. The remainder of the fields in the output are self-explanatory.

```
*A:PE-3# show log event-handling handler "event-handler-1"

=====
Event Handling System - Handlers
=====

Handler          : event-handler-1
=====
Description      : (Not Specified)
Admin State      : up                      Oper State : up
-----

Handler Execution Statistics
Success          : 1
Err No Entry     : 0
Err Adm Status   : 0
Total            : 1
-----

Handler Action-List Entry
```

```
-----
Entry-id       : 10
Description    : (Not Specified)
Admin State    : up                               Oper State : up
Script
  Policy Name   : vrrp-master-policy
  Policy Owner  : TiMOS CLI
Min Delay      : 0
Last Exec      : 09/13/23 07:45:35 UTC
-----
Handler Action-List Entry Execution Statistics
Success        : 1
Err Min Delay  : 0
Err Launch     : 0
Err Adm Status : 0
Total          : 1
=====
```

The example includes an event trigger and script to meet the requirements of a fail-forward where PE-3 becomes VRRP primary. Now, configuration is needed for when PE-3 reverts to VRRP backup. Without another event trigger and script, PE-3 will continue to advertise the prefix 172.16.1.0/29 with a local preference of 150 and upstream/downstream traffic will be asymmetric through PE-2/PE-3 respectively.

As before, a script is required. Because PE-2 advertises the prefix with a local preference of 100 (default), PE-3 needs to advertise the same prefix with a lower value (50 in the following output), so that PE-2 is the preferred next hop.

```
*A:PE-3>file cf3:\ # type cf3:vrrp-backup.txt
File: vrrp-backup.txt
-----
exit all
configure router policy-options
begin
policy-statement redundant-interface
entry 10
action accept
local-preference 50
exit
exit
exit
commit
exit all
=====
```

The script must then be configured within the script-control context, and subsequently referenced in a script policy as vrrp-backup-policy.

```
# on PE-3:
configure
  system
    script-control
      script "vrrp-backup-script"
        location "cf3:/vrrp-backup.txt"
        no shutdown
      exit
    script-policy "vrrp-backup-policy"
      results "cf3:/script-revert-results.txt"
      script "vrrp-backup-script"
      max-completed 4
      lifetime forever
```

```
no shutdown
exit
```

The event handler acts as the interface between the configured script policy and event trigger. Therefore, a second event handler is configured with an action list consisting of a single entry referencing the newly configured vrrp-backup-policy.

```
# on PE-3:
configure
log
    event-handling
        handler "event-handler-2"
        action-list
            entry 10
                script-policy "vrrp-backup-policy"
                no shutdown
            exit
        exit
    exit
no shutdown
exit
```

Finally, the event trigger is configured. To revert to VRRP Backup, the application is VRRP and the event number is 2006 (tmnxVrrpBecameBackup). The configuration is filtered on the message field, as before, using log filter 2, so that it is specific to the interface named redundant-interface.

```
# on PE-3:
configure
log
    filter 2 name "itf 172.16.1.3 state becomes backup"
    default-action drop
    entry 10 name "becameBackup"
        action forward
        match
            message eq pattern "interface redundant-interface changed
                                state to backup"
        exit
    exit
exit
```

```
# on PE-3:
configure
log
    event-trigger
        event "vrrp" 2006
        trigger-entry 1
            event-handler "event-handler-2"
            log-filter 2
            no shutdown
        exit
    no shutdown
    exit
exit
```

The configuration of the example event handling for the revertive failure event (PE-3 transitions to VRRP backup) is now complete. By re-enabling the spoke-SDP between PE-1 and PE-2, the VRRP message path is restored, and PE-2 again becomes the VRRP master. The following events are generated:

- The first indicates that PE-3 has become VRRP backup for the interface named redundant-interface.
- The second indicates that EHS handler event-handler-2 was invoked by a CLI user.

- The third indicates that a script file has initiated an attempt to execute CLI commands contained in script file vrrp-backup.txt.
- The fourth indicates that the attempt to execute those CLI commands was successful.

```
158 2023/09/13 07:58:24.686 UTC MINOR: VRRP #2006 Base Becoming Backup
"VRRP virtual router instance 1 on interface redundant-interface changed state to
backup - current master is 172.16.1.2"

159 2023/09/13 07:58:24.686 UTC MINOR: SYSTEM #2069 Base EHS script
"Ehs handler : "event-handler-2" with the description : "" was invoked by the cli-user
account "not-specified". "

160 2023/09/13 07:58:24.691 UTC MAJOR: SYSTEM #2052 Base CLI 'exec'
"A CLI user has initiated an 'exec' operation to process the commands in the SROS CLI
file cf3:/vrrp-backup.txt"

161 2023/09/13 07:58:24.696 UTC MAJOR: SYSTEM #2053 Base CLI 'exec'
"The CLI user initiated 'exec' operation to process the commands in the SROS CLI
file cf3:/vrrp-backup.txt has completed with the result of success"
```

The following outputs confirm that PE-3 is VRRP backup, and that the local preference attribute for prefix 172.16.1.0/29 has changed to a value of 50. The result of this action is that PE-2 will now be the transit router for both upstream and downstream traffic.

```
*A:PE-3# show router vrrp instance
```

```
=====
VRRP Instances
=====
```

Interface Name	VR Id IP	Own	Adm Opr	State Pol Id	Base Pri InUse Pri	Msg Int Inh Int
redundant-interface	1 IPv4	No	Up Up	Backup n/a	253 253	1 No

```
Backup Addr: 172.16.1.1
```

```
-----
Instances : 1
=====
```

```
*A:PE-3# show router bgp routes 172.16.1.0/29 hunt | match expression "Network|Nexthop|To|Local
Pref"
```

Network	: 172.16.1.0/29	
Nexthop	: 192.0.2.2	
Res. Nexthop	: 192.168.23.1	
Local Pref.	: 100	Interface Name : int-PE-3-PE-2
Network	: 172.16.1.0/29	
Nexthop	: 192.0.2.3	
To	: 192.0.2.6	
Res. Nexthop	: n/a	
Local Pref.	: 50	Interface Name : NotAvailable

Conclusion

EHS allows operators to configure user-defined actions on the router when an event occurs. The event trigger can be anything that is generated by the event-control framework, and explicit filtering is possible

using regular expressions. A user-defined action typically runs a script that allows any CLI commands to be executed. Multiple actions are permitted, running multiple scripts if required.

SR OS NETCONF Server Basics

This chapter provides information about SR OS NETCONF server basics.

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)


Applicability

This chapter was initially written for SR OS Release 16.0.R4, but the MD-CLI in the current edition is based on SR OS Release 21.5.R2.

Overview

The SR OS Network Configuration Protocol (NETCONF) server can communicate with a NETCONF client, that is, exchange hello messages, receive requests, and reply with responses. Before communicating with the SR OS NETCONF server, some SR OS configurations are prerequisites, and others are optional. This chapter describes the basic configurations needed for a seamless interaction with the SR OS NETCONF server. [Figure 7: NETCONF client-server communication](#) shows the NETCONF client-server communication between the controller and the SR OS node.

Figure 7: NETCONF client-server communication



28626

Configuration

The following steps describe the procedure to configure a NETCONF server on SR OS.

- Because NETCONF uses SSH for transport, enable the SSH server in SR OS:

```
configure system security ssh no server-shutdown
```

- Enable the NETCONF server:

```
configure system netconf no shutdown
```

- Enable the YANG modules to use with NETCONF; for example, the Nokia modules:

```
configure
  system
    management-interface
      yang-modules
        no nokia-combined-modules
        nokia-submodules
    exit
```



Note:

The Nokia combined modules and the Nokia submodules cannot both be set to true at the same time.

- Configure an "nc_user" user with administrative privileges (**access netconf**):

```
configure
  system
```

```
security
  user "nc-user"
    password <password>
    access netconf
    console
      member "administrative"
    exit
  exit
```

- Optionally, enable NETCONF auto-config-save, which auto-saves the data (that is, makes it persistent) after each successful NETCONF commit:

```
configure system netconf auto-config-save
```

- Optionally, grant the NETCONF user permission to **lock** a datastore through NETCONF:

```
configure system security profile "administrative" netconf base-op-authorization lock
```

- Optionally, grant the NETCONF user permission to kill an open NETCONF session:

```
configure system security profile "administrative" netconf base-op-authorization kill-session
```

- Save the configuration:

```
admin save
```

Conclusion

This chapter describes general SR OS NETCONF server configurations.

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)