Onboarding and Security Enablement

Onboarding and security enablement

Onboarding and security enablement requires the DPU to be in a powered-on state on the Smart Switches. By default, the DPU is in a powered-off state. When the DPU is in a powered-off state, the switch works in a network-only mode. When the DPU is in a powered-on state, the switch works in both network and DPU security modes.

Note Enabling the DPU security services requires registering the Smart Switch with the Hypershield Controller. Because of this you should configure the Hypershield Controller first and then onboard the Smart Switch with the Hypershield Controller.

The Smart Switch, when used in DPU security mode, provides a distributed, stateful Layer-4 segmentation and DPU security enforcement directly within the switch using embedded Data Processing Units (DPUs).

For more information about licenses, refer to the Cisco NX-OS Licensing Options Guide. For more information about Smart Licensing Using Policy (SLP) and setup, refer to the specific articles that help you learn, deploy, and troubleshoot SLP Cisco NX-OS Licensing collection page.

For troubleshooting information, refer to Cisco N9300 Series Smart Switch Troubleshooting Guide for DPU Security Mode.

Configure Hypershield Controller

Configure the Hypershield Controller before onboarding and security enablement.

The Hypershield Controller, by default, uses 10.20.0.0/16 and 10.30.0.0/16 IP subnets for setting up internal network connectivity within the Hypershield Controller Kubernetes cluster. If these subnets overlap with any routable networks in your data center environment and you want to avoid routing issues with the controller, update the `bootstrap.yaml` on the Hypershield Controller. The other option is to deploy the Hypershield Controller with a different set of subnets of /16 size instead of the default. The fields to be updated (with example overrides) in the `bootstrap.yaml` are:

```
# Override the bootstrap defaults of 10.20.0.0/16 and 10.30.0.0/16.
# 10.0.0.0/12 is not Cisco-internal routable (e.g. 10.0.0.0/16 through 10.15.0.0/16).
podCIDR: "10.8.0.0/16"
serviceCIDR: "10.9.0.0/16"
```

Workflow for onboarding and enabling security

Perform these steps to enable the DPU security mode on N9300 Smart Switches.

1. Bringup the switch by installing the software on the switch. NX-OS manages software images for both the Smart Switch and its DPUs by bundling them together.

   For more information about Software images, refer to the Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide.

2. Apply networking configurations similar to that of rest of the N9000 Series Switches.

   For more information, refer to Configuration Guides for Cisco NX-OS 10.6(x). For example, to configure VRF contexts and Layer 3 interfaces with VRF members and configure VLANs and interfaces pertaining to those VLANs, refer to Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide, and Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide.

3. Enable the service acceleration feature.

4. Configure a loopback interface for Hypershield source-interface for the Hypershield Agent to Hypershield Controller connectivity; refer to Create a loopback interface for Hypershield connectivity.

5. Add the configuration to connect to the Hypershield Controller; refer to Configure Hypershield connectivity.

6. Request and apply the token from Hypershield Controller; refer to Register the N9300 Smart switch with Hypershield.

7. Configure traffic inspection for traffic in the required VRFs and VLANs; refer to Configure Traffic redirection to the firewall service.

8. Enable the firewall service functionality; refer to Enable service firewall.

Enable Service Acceleration feature

The **feature service-acceleration** command is used to enable the DPU security mode or the DPU functionality. When enabled, NX-OS powers up the DPUs, which takes some time. As a result, NX-OS prevents the execution of no feature service-acceleration command until the DPUs are fully powered up and reach a terminal state.

The feature service-acceleration command cannot be disabled until all DPUs in the system have been powered on. Feature service-acceleration can be re-enabled after it has been disabled, without a switch reboot.

Note If you use the configure replace feature (refer to Performing Configuration Replace), the success of configure replace may depend on the timing of the configuration relative to when service-acceleration was enabled or disabled.

Perform this task to enable the service acceleration feature and power up the DPUs in the switch.

Procedure

Enable the **feature service-acceleration** command to power up the DPUs.

Example:

```
switch(config)# feature service-acceleration
```

If the **feature service-acceleration** command is *not* configured, the DPUs are powered off. The switch functions as a NXOS switch.

Note
Enabling feature service-acceleration powers up the DPUs; however, to finalize the configuration, you must define which VRFs traffic should be redirected to the DPU. See Configure traffic redirection to the firewall service, on page 8.

Verify service acceleration

Perfrom this task to verify service acceleration enablement.

Procedure

Step 1 Verify the service acceleration status using the **show run service-acceleration | grep feature** command.

Example:

```
switch# show run service-acceleration | grep feature

feature service-acceleration
```

Step 2 Use the **show interfaces brief** command to view the status of the interfaces.

Example:

```
switch# show interface brief
..!
--------------------------------------------------------------------------------
Service         VLAN    Type Mode   Status  Reason                 Speed     Port
Ethernet                                                                     Ch #
--------------------------------------------------------------------------------
SEth1/1         --      eth  routed up      none                   200G(D) --
SEth1/2         --      eth  routed up      none                   200G(D) --
SEth1/3         --      eth  routed up      none                   200G(D) --
SEth1/4         --      eth  routed up      none                   200G(D) --
```

All DPUs power up when the feature service-acceleration is enabled.

Step 3 Use the **show module** command to verify if the DPUs are powered up and online.

Example:

```
switch# show module
Mod Ports                  Module-Type                           Model           Status
--- ----- ------------------------------------------------ --------------------- --------
1    56   48x25G/2x100G/6x400G Ethernet Module             N9348Y2C6D-SE1U       ok
27   0    Virtual Supervisor Module                        N9348Y2C6D-SE1U       active *

Mod  Sw                       Hw    Slot
---  ---------------------- ------ ----
1    10.6(3r)                 0.1060 NA
27   10.6(3r)                 0.1060 VSUP

Mod  MAC-Address(es)                          Serial-Num
---  ---------------------------------------  ----------
1    fc-58-9a-1d-04-d4 to fc-58-9a-1d-05-b8  FDO29090ZW6
27   fc-58-9a-1d-04-d4 to fc-58-9a-1d-05-b8  FDO29090ZW6

Mod  Online Diag Status
---  ------------------
1    Pass
27   Pass

* this terminal session
Mod DPU   Module-Type           Model         Status
--- ---  ------------- ---------------------- --------
1   1    DPU           N9348Y2C6D-SE1U-DPU   ok
1   2    DPU           N9348Y2C6D-SE1U-DPU   ok

Mod DPU     Sw            Hw      Serial-Num   Online Diag Status
--- ---  ------------- ---------- ------------ ------------------
1   1     1.162.2      ES         NA           Pass
1   2     1.162.2      ES         NA           Pass
```

Disable service-acceleration feature

Procedure

Disable service-acceleration feature and Layer 4-7 services on the switch, for it to function in DPU powered-off state, using the using the **no feature service-acceleration** command.

Example:

```
switch(config)# no feature service-acceleration
```

Create a loopback interface for Hypershield connectivity

Perform this task to create a loopback interface to be used to connect to the Hypershield Controller. To associate this loopback interface with the Hypershield agent in the service system Hypershield configuration, see Configure Hypershield connectivity, on page 6.

To ensure that the loopback address is reachable from a front-panel interface without any routing configuration on the external management network you can configure the interface as follows:

Procedure

Step 1 Create a loopback interface for the Hypershield source-interface using the **interface loopback** *instance* command.

Example:

```
switch(config)# interface loopback 100
switch(config-if)#
```

Step 2 Configure an IP address for the interface using the **ip address** *ip-address/length* command.

Example:

```
switch(config-if)# ip address 10.16.0.140/32
```

For more information about IP addresses, see the Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide.

ip-address/length: Sets an IP address for the loopback

Note

- You have to configure /32 as length of the IP address for Hypershield Agent.
- Management traffic from Hypershield Agent to Hypershield system is *not* supported over IPv6.
- The loopback interface must be configured in the default VRF.

Step 3 Configure the physical interface and assign it the IP address of the loopback interface using the **interface** *interface-type slot/port*, **medium p2p**, and **ip unnumbered loopback** *loopback instance* commands. This means the physical interface borrows the IP address of the loopback interface, allowing the loopback IP to be reachable through this physical interface without requiring additional IP addresses.

Example:

```
switch(config)# interface eth1/3/3
switch(config-if)# medium p2p
switch(config-if)# ip unnumbered loopback 1
```

Step 4 Add static route for subnet through the physical interface with the next hop. This static route directs traffic destined for the subnet to go out through the interface with the next-hop IP, ensuring proper routing of packets to and from the loopback address through the physical interface. To configure, use the **ip route** *ip-address/subnet-mask interface-type slot/port next-hop-ip-address* command.

Example:

```
switch(config)# ip route 10.16.0.0/24 eth1/3/3 10.16.0.1
```

Step 5 Use the **show run** command to verify the loopback interface configuration.

Example:

```
switch# show run
!
interface loopback1
  ip address 10.16.0.140/32
!
interface eth1/3/3
  medium p2p
  ip unnumbered loopback1
!
ip route 10.16.0.0/24 eth1/3/3 10.16.0.1
```

Configure Hypershield connectivity

Hypershield management traffic operates within the default VRF. A loopback interface must be configured for this communication. This loopback is used as the source interface in the service system hypershield configuration, its IP address is exclusively reserved for communication between the agent and the Hypershield and it can also be used for HA. The loopback IP address cannot be reused for any control protocols running on the switch. Any attempt to test reachability to other destinations using ICMP echo from the NX-OS command line interface fails, if the loopback IP address is specified as the source.

Connectivity from the Hypershield agent to the Hypershield controller must occur over a routed front panel interfaces within the default VRF and is not supported over SVIs on vPCs.

Management traffic from the Hypershield Agent to the Hypershield Controller is not supported over IPv6.

Even if the default VRF is used for the connectivity to the Hypershield Controller, this VRF can also be used for traffic forwarding and it is possible to configure the default VRF for traffic inspection.

Use the troubleshooting commands to verify the connectivity between the Hypershield Agent and the Hypershield Controller. For more information, see Cisco N9300 Series Smart Switches DPU Security Mode Troubleshooting Guide.

Perform this task to establish connectivity between the Hypershield Agent in the N9300 Smart switch and the Hypershield Controller.

Procedure

Step 1 Enable the **service system hypershield** command to set up the Hypershield instance.

Example:

```
switch(config)# service system hypershield
```

Step 2 Configure the **source-interface** command to assign to the Hypershield Agent the IP address of the loopback interface that you previously defined.

Example:

```
switch(config-svc-sys)# source-interface loopback 100
```

The loopback interface must be configured in the default VRF. See Create a loopback interface for Hypershield connectivity, on page 5.

The example shows service acceleration feature configuration.

```
switch# show run service-acceleration
!
feature service-acceleration
service system hypershield register 34C58A…
!
service system hypershield
source-interface loopback 100
...!
service firewall
  in-service
```

Register the Smart Switch with Hypershield

Registering the Smart Switch with Hypershield requires obtaining a token from Hypershield.

This token must be entered on the switch to establish communication between the switch and Hypershield. For more information, refer to Cisco Hypershield User Documentation.

Procedure

Establish communication between the switch and Hypershield with the obtained token using the **service system hypershield register** *otp* command.

Example:

```
Switch# service system hypershield register 34C58A…
```

Execute this command at the EXEC level.

otp: Indicates the token string (maximum size 4094). Type the token without any quotes.

Verify Hypershield connection status

Procedure

Verify the status of the connectivity using the **show service-acceleration status details** command.

Example:

```
switch# show service-acceleration status details

Service System: hypershield
Source Interface: loopback101 (12.0.0.1)
```

```
DPU Load-balance Mode: symmetric-hash
Agent Status: firewall-ready,redirect-installed
Agent Health Status: ok
Controller Connection Status: success
Agent HA Status: ready
Services:
Firewall: in-service
[...]
```

Configure traffic redirection to the firewall service

You can configure VRFs or VLANs for traffic inspection through the DPUs. In this document we refer to these VRFs and VLANs as service VRFs and service VLANs respectively.

Note The firewall service does not inspect non-ip traffic, multicast IP traffic, broadcast IP traffic, traffic destined for the local switch, traffic originating from the supervisor, BFD echo packets, management VRF traffic, and traffic pertaining to any VRFs or VLANs that are not added under service firewall. All other IPv4 and IPv6 traffic pertaining to the VRFs and VLANs for service firewall are sent to the DPU firewall services for inspection. The DPU inspects IPv4 and IPv6 traffic with unicast IP addresses regardless of the MAC addresses being multicast or broadcast.

Configure load-balance mode

By default, traffic in service VRFs and VLANs is load-balanced using a symmetric hardware-based hash to all available DPUs. Routed and bridged traffic for VRFs and VLANs that are configured for filtering is distributed across the DPUs using a hardware-based symmetric hash mechanism. Users have the option to enable a pinning load-balance mode to selectively pin traffic for entire VRFs or VLANs to a specific DPU.

Note Only symmetric hashing mode is supported for Inter-VRF filtering.

The load-balance mode may only be modified when the service firewall is not in the in-service state or before the service firewall is configured.

You can configure the switch to selectively pin traffic pertaining to VRFs and VLANs to specific DPUs or allow the system to select specific DPUs for specific VRFs and VLANs.

Note This is an optional step and as of NX-OS 10.6(3s)F, the **service dpu load-balancing pinning** configuration is not for production deployments.

Procedure

Step 1 Configure the load-balance mode to pinning mode.

Example:

```
switch(config)# service dpu load-balance pinning
```

Step 2 Revert the configuration to symmetric-hash mode.

Example:

```
switch(config)# service dpu load-balance symmetric-hash
```

OR

```
switch(config)# no service dpu load-balance pinning
```

The example shows service acceleration with pinning load-balance mode configured.

```
switch# show run service-acceleration
!
feature service-acceleration
service dpu load-balance pinning
!
```

Configure Service VRFs

If a VRF is configured under service firewall, this means that this VRF is subject to traffic redirection. Unicast routed IPv4 and IPv6 traffic arriving on Layer 3 physical interfaces and port-channels, physical and port-channel subinterfaces, SVIs and VXLAN fabric interfaces is inspected if they pertain to this VRF. Traffic arriving on SVIs over Layer 2 access and trunk interfaces, port-channels, and vPCs is inspected if the SVIs pertain to this VRF.

If the VRF is subject to traffic redirection, and if the service firewall is not `in-service`, the unicast traffic routed by this VRF is dropped until the service firewall is `in-service` and until a security rule to allow that traffic is in place in the DPU. Traffic redirection doesn't apply to control protocol traffic: if the VRFs are configured for service-acceleration, and the traffic is destined to the supervisor of the N9300 Smart switch, this is not inspected. Hence it is not dropped regardless of whether service firewall is in-service or not, and regardless of whether the security rules are configured.

When the service firewall is not in `in-service` state, and the DPU is not ready to inspect traffic, the Layer 3 routing protocols ensure graceful insertion and removal (GIR) behavior. To achieve this, the service-acceleration VRFs isolate the switch from the network by altering the route advertisement behavior. Protocols supporting this functionality include Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and Routing Information Protocol (RIP). The protocols resume regular route advertisement behavior once the service firewall is ready for use and is in `in-service` state.

The default VRF may be configured as a VRF under service firewall for traffic inspection. Traffic from the Hypershield Agent and Controller in the default VRF are excluded from inspection. The management VRF cannot be configured for traffic inspection under service firewall.

Perform this task to specify the VRFs whose traffic must be firewalled.

Procedure

Step 1 Enable the **service firewall** command.

Example:

```
switch(config-svc-sys)# service firewall
```

Step 2 Configure the VRF under the service firewall to redirect the traffic in the VRF for inspection by the firewall service using the **vrf** *vrf-name* command.

Example:

```
switch(config-svc-sys-fw)# vrf red module-affinity dynamic
```

By default, the switch uses the hardware-based symmetric hash to distribute the traffic across the DPUs. The VRF context and other required networking configuration is entered as usual on the switch.

Step 3 (Optional) Configure a specific DPU number to indicate that traffic in the VRF must be inspected by the firewall service in that DPU using the **module-affinity** command when the load-balance mode of pinning is configured.

Note
Explicitly configuring module-affinity in the 10.6(3s)F release is only for testing, not for production deployments.

Example:

```
switch(config-svc-sys-fw)# vrf blue module-affinity 1
```

Alternatively in the pinning load-balance mode, if module-affinity dynamic is configured, the system selects a DPU to inspect all traffic in the VRF.

When the system is not using the pinning load-balance mode, module-affinity can only be configured as dynamic.

This is a sample configuration example of a VRF for traffic redirection to the DPU.

```
switch(config-svc-sys)# service firewall
switch(config-svc-sys-fw)# vrf red module-affinity dynamic
switch(config-svc-sys-fw)# vrf blue module-affinity 1
```

Traffic in the VRF blue is inspected by DPU1.

This example shows service acceleration with firewalling enabled, with pinning load-balance mode

```
switch# show run service-acceleration
!
feature service-acceleration
service dpu load-balance mode pinning
!
service system hypershield
source-interface loopback 100
service firewall
   vrf blue module-affinity 1
   vrf red module-affinity dynamic
```

The example shows service acceleration with firewalling enabled, with symmetric-hash load-balance mode.

```
switch# show run service-acceleration
!
feature service-acceleration
service system hypershield register 34C58A…
!
service system hypershield
source-interface loopback 100
service firewall
vrf blue module-affinity dynamic
vrf red module-affinity dynamic
```

Configure Service VLANs

Bridged traffic within VLANs (enabled on Layer 2 access and trunk interfaces, port-channels, and vPCs) configured for service firewall undergo inspection.

VLANs configured for service firewall might have SVIs that do not belong to service firewall VRFs or might not have SVIs configured on the switch. In such scenarios, only bridged traffic within those VLANs are inspected while routed traffic is not.

VLANs might not be configured for service firewall but could have SVIs associated with service firewall VRFs. In such cases, only routed traffic within those VLANs are inspected and bridged traffic is not.

You can enable both routed and bridged traffic for a VLAN for inspection. To do so, configure the VLAN for the service firewall and configure the VRF related to its SVI for the service firewall. Note that the users or systems can choose different DPUs for inspecting routed and bridged traffic within a VLAN.

Just like for service VRFs, if a VLAN is configured under the service firewall configuration (indicating that the VLAN is subject to traffic filtering), and if the service firewall is not `in-service`, the IPv4 and IPv6 traffic bridged for this VLAN is dropped until the service firewall is `in-service` and a security rule allowing that traffic is established in the DPU.

Perform this task to specify the VLANs whose bridged traffic must be firewalled.

Procedure

Step 1 Enable the **service firewall** command.

Example:

```
switch(config-svc-sys)# service firewall
```

Step 2 Configure the VLAN under the service firewall to redirect the bridged traffic in the VLAN for inspection by the firewall service using the **vlan id** *{ vlan-id | vlan-range }* command.

Example:

```
switch(config-svc-sys-fw)# vlan id 100 bridged-traffic module-affinity dynamic
```

Example:

```
switch(config-svc-sys-fw)# vlan id 105-110 bridged-traffic module-affinity dynamic
```

By default, the switch uses the hardware-based symmetric hash to distribute the traffic across the DPUs.

The VLAN definition and other required networking configuration is entered as usual on the switch.

Step 3 (Optional) If the pinning load-balance mode has been configured, you may configure a specific DPU number to indicate that traffic in the VLAN must be inspected by the firewall service in that DPU using the **module-affinity** command.

Note
Explicitly configuring module-affinity in the 10.6(3s)F release is only for testing, not for production deployments.

Example:

```
switch(config-svc-sys-fw)# vlan id 200 bridged-traffic module-affinity 2
```

Alternatively, in the pinning load-balance mode, if module-affinity dynamic is configured, the system selects a DPU to inspect all traffic in the VLAN.

When the system is not using the pinning load-balance mode, module-affinity can only be configured as dynamic.

This is a sample configuration example of a VLAN for traffic redirection to the DPU.

```
switch(config-svc-sys)# service firewall
switch(config-svc-sys-fw)# vlan id 100 bridged-traffoc module-affinity dynamic
switch(config-svc-sys-fw)# vlan id 200 bridged-traffic module-affinity 2
```

Bridged traffic in VLAN 200 is inspected by DPU2.

This example shows service acceleration with firewalling enabled, with pinning load-balance mode.

```
switch# show run service-acceleration
!
feature service-acceleration
service dpu load-balance mode pinning
!
service system hypershield
source-interface loopback1
 service firewall
   vrf blue module-affinity 1
   vrf red module-affinity dynamic
   vlan id 100 bridged-traffoc module-affinity dynamic
   vlan id 200 bridged-traffic module-affinity 2
```

This example shows service acceleration with firewalling enabled, with symmetric-hash load-balance mode.

```
switch# show run service-acceleration

feature service-acceleration
service dpu load-balance mode pinning

service system hypershield
  source-interface loopback 100
  service firewall
    vrf blue module-affinity 1
    vrf red module-affinity dynamic
    vlan id 100 bridged-traffoc module-affinity dynamic
    vlan id 200 bridged-traffic module-affinity dynamic
    vlan id 100,200 bridged-traffic module-affinity dynamic
```

Enable service firewall

The **in-service** command enables the service firewall functionality, and allows traffic inspection by the DPUs.

You can also use the **no in-service** command to trigger maintenance mode for the firewall functionality and, to disable traffic redirection to DPUs for any reason or modify the load-balance mode.

Procedure

Use the **in-service** command to enable the service firewall to redirect specific traffic.

Example:

```
switch(config-svc-sys-fw)# in-service
```

Verify the service firewall and redirection status

You can verify the status of the service firewall and redirection when the status is out-of-service or in-service. This section provides examples for both the statuses.

- Use the **show service-acceleration status details** command to view the firewall service when is in out-of-service state.

```
switch# show service-acceleration status details
Service System: hypershield
   Source Interface: loopback48 (10.1.1.1)
   DPU Load-balance Mode: symmetric-hash
   Agent Status: firewall-disable
   Agent Health Status: ok
   Controller Connection Status: init
   Agent HA Status: not-initialized
   Services:
     Firewall: out-of-service

VRF             Operational State       Affinity(DPU)
======================================================================
blue            isolated                n/a
red             isolated                n/a

VLAN            Operational State       Affinity(DPU)
======================================================================
100             VLAN set to drop        n/a
200             VLAN set to drop        n/a
```

- Use the **show service-acceleration status details** command to view the status of the firewall service when it is in `in-service` state.

  This example shows the firewall service in the `in-service` state and in pinning mode.

```
switch# show service-acceleration status details
  Service System: hypershield
    Source Interface: loopback48 (10.1.1.1)
    DPU Load-balance Mode: pinning
    Agent Status: firewall-ready,redirect-installed
    Agent Health Status: ok
```

```
    Controller Connection Status: success
    Agent HA Status: not-initialized Services:
      Firewall: in-service

      VRF                               Operational State       Affinity(DPU)
      ================================================================================
      blue                              forwarding ready        1
      red                               forwarding ready        2

      VLAN                              Operational State       Affinity(DPU)
      ================================================================================
      100                               forwarding ready        1
      200                               forwarding ready        2
```

This example shows the firewall service in the in-service state and in symmetric-hash mode.

```
show service-acceleration status details
  Service System: hypershield
    Source Interface: loopback48 (10.1.1.1)
    DPU Load-balance Mode: symmetric-hash
    Agent Status: firewall-ready,redirect-installed
    Agent Health Status: ok
    Controller Connection Status: success
    Agent HA Status: not-initialized
    Services:
      Firewall: in-service

      VRF                               Operational State       Affinity(DPU)
      ================================================================================
      blue                              forwarding ready        1-4
      red                               forwarding ready        1-4

      VLAN                              Operational State       Affinity(DPU)
      ================================================================================
      100                               forwarding ready        1-4
      200                               forwarding ready        1-4
```

- Use the command **show service-acceleration redirect-policy brief** to identify the service port-channel subinterfaces for redirecting traffic from the VRFs or VLANs to the DPUs.

  This example shows the output when symmetric-hash load-balance mode is in use.

```
switch# show service-acceleration redirect-policy brief

VRF                            AF Type Interface[Status]    Affinity   Redirect Status

=============================================================================================
red                              IPv4    Spo512.10[UP]        1-4        Enabled

blue                             IPv4    Spo512.11[UP]        1-4        Enabled

red                              IPv6    Spo512.10[UP]        1-4        Enabled

blue                             IPv6    Spo512.11[UP]        1-4        Enabled


VLAN                           AF Type Interface[Status]    Affinity   Redirect Status

=============================================================================================
100                              IPv4    Spo512[UP]           1-4        Enabled

200                              IPv4    Spo512[UP]           1-4        Enabled

100                              IPv4    Spo512[UP]           1-4        Enabled
```

```
200                              IPv4    Spo512[UP]           1-4        Enabled
```

This example shows the output when pinning load-balance mode is in use.

```
switch# show service-acceleration redirect-policy brief

VRF                            AF Type Interface[Status]    Affinity   Redirect Status

=============================================================================================
red                              IPv4    Spo2.10[UP]        1-4        Enabled

blue                             IPv4    Spo1.11[UP]        1-4        Enabled

red                              IPv6    Spo2.10[UP]        1-4        Enabled

blue                             IPv6    Spo1.11[UP]        1-4        Enabled


VLAN                           AF Type Interface[Status]    Affinity   Redirect Status

=============================================================================================
100                              IPv4    Spo1[UP]           1-4        Enabled

200                              IPv4    Spo2[UP]           1-4        Enabled
100                              IPv4    Spo1[UP]           1-4        Enabled

200                              IPv4    Spo2[UP]           1-4        Enabled
```