



Cisco Nexus 9000 Series NX-OS ePBR Configuration Guide, Release 10.2(x)

First Published: 2021-12-16

Last Modified: 2022-04-26

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and-if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021–2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

Trademarks ?

P R E F A C E

Preface v

Audience v

Document Conventions v

Related Documentation for Cisco Nexus 9000 Series Switches vi

Documentation Feedback vi

Communications, Services, and Additional Information vi

C H A P T E R 1

New and Changed Information 1

New and Changed Information 1

C H A P T E R 2

Overview 3

Licensing Requirements 3

Supported Platforms 3

C H A P T E R 3

Configuring ePBR L3 5

Information About ePBR L3 5

Licensing Requirements 5

Configuring ePBR Service and Policy 5

Applying ePBR to an Interface 6

Creating Bucket and Load Balancing 6

ePBR Object Tracking, Health Monitoring, and Fail-Action 6

ePBR Session-based Configuration 7

ePBR Multi-Site 8

ACL Refresh 8

Guidelines and Limitations for ePBR L3	8
Configuring ePBR L3	11
Configuring ePBR Service, Policy, and Associating to an Interface	11
Modifying a Service Using ePBR Session	13
Modifying a Policy Using ePBR Session	14
Updating the Access-list Used by ePBR Policies	16
ePBR Show Commands	16
Verifying ePBR Configuration	17
Configuration Examples for ePBR L3	18
Additional References	26
Related Documents	26
Standards	26

CHAPTER 4**Configuring ePBR L2** 27

Information About ePBR L2	27
Configuring ePBR Service and Policy	27
Applying ePBR to an L2 Interface	28
Enabling Production Interfaces as Access Port	28
Enabling Production Interfaces as Trunk Ports	28
Creating Bucket and Load Balancing	28
ePBR Object Tracking, Health Monitoring, and Fail-Action	28
ePBR Session-based Configuration	29
ACL Refresh	29
Guidelines and Limitations for ePBR L2	30
Configuring ePBR Service, Policy, and Associating to an Interface	32
Modifying a Service Using ePBR Session	34
Modifying a Policy Using ePBR Session	35
Updating the Access-list Used by ePBR Policies	37
ePBR Show Commands	37
Verifying ePBR Configuration	38
Configuration Examples for ePBR	39



Preface

This preface includes the following sections:

- [Audience, on page v](#)
- [Document Conventions, on page v](#)
- [Related Documentation for Cisco Nexus 9000 Series Switches, on page vi](#)
- [Documentation Feedback, on page vi](#)
- [Communications, Services, and Additional Information, on page vi](#)

Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which you supply the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x y]	Square brackets enclosing keywords or arguments that are separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments that are separated by a vertical bar indicate a required choice.
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.

Convention	Description
variable	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string includes the quotation marks.

Examples use the following conventions:

Convention	Description
screen font	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information that you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
<>	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Related Documentation for Cisco Nexus 9000 Series Switches

The entire Cisco Nexus 9000 Series switch documentation set is available at the following URL:

http://www.cisco.com/en/US/products/ps13386/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus9k-docfeedback@cisco.com. We appreciate your feedback.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you’re looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

New and Changed Information

- New and Changed Information, on page 1

New and Changed Information

Table 1: New and Changed Features for Cisco NX-OS Release 10.2(x)

Feature	Description	Changed in Release	Where Documented
ESR: ePBR L2	Added Enhanced Policy Based Routing (ePBR) support for L2 service type which provides traffic redirection and service chaining across the standalone and fabric topologies	10.2(3)F	Configuring ePBR L2, on page 27
ESR: ePBR L3	Added Enhanced Policy Based Routing (ePBR) support for L3 service type which provides traffic redirection and service chaining across the standalone and fabric topologies.	10.2(3)F	Configuring ePBR L3, on page 5
Support for ePBR L2 on access port	Added support for Production interfaces as Access Port	10.2(3)F	
Multiple back-up hops	Added guidelines and limitations to inter-VRF service chaining	10.2(2)F	Guidelines and Limitations for ePBR L3, on page 8
CTP probes	Added guidelines and limitations for EPBR layer-2 services	10.2(2)F	Guidelines and Limitations for ePBR L3, on page 8

New and Changed Information

Feature	Description	Changed in Release	Where Documented
ePBR Multi-site	Service-chaining in a VXLAN multisite fabric	10.2(1)F	ePBR Multi-Site , on page 8
Inter-VRF service chaining	Added guidelines and limitations to inter-VRF service chaining	10.2(1)F	Guidelines and Limitations for ePBR L3, on page 8
Configuring ePBR Service and Policy	ePBR Service and Policy	10.2(1)F	Configuring ePBR Service and Policy, on page 5



CHAPTER 2

Overview

- Licensing Requirements, on page 3
- Supported Platforms, on page 3

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide* and the *Cisco NX-OS Licensing Options Guide*.

Supported Platforms

Starting with Cisco NX-OS release 7.0(3)I7(1), use the [Nexus Switch Platform Support Matrix](#) to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.



CHAPTER 3

Configuring ePBR L3

This chapter describes how to configure Enhanced Policy-based Redirect (ePBR) on Cisco NX-OS devices.

- [Information About ePBR L3, on page 5](#)
- [Guidelines and Limitations for ePBR L3, on page 8](#)
- [Configuring ePBR L3, on page 11](#)
- [Configuration Examples for ePBR L3, on page 18](#)
- [Additional References, on page 26](#)

Information About ePBR L3

Enhanced Policy-based Redirect (ePBR) in Elastic Services Re-direction (ESR) provides traffic redirection and service chaining across the NX-OS and fabric topologies by leveraging policy-based redirect solution and achieves service chaining without adding extra headers, and avoids latency in using extra headers.

ePBR enables application-based routing and provides a flexible, device-agnostic policy-based redirect solution without impacting application performance. The ePBR service flow includes the following tasks:

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the [Cisco NX-OS Licensing Guide](#) and the [Cisco NX-OS Licensing Options Guide](#).

Configuring ePBR Service and Policy

You must first create an ePBR service which defines the attributes of service end points. Service end points are the service appliances such as firewall, IPS, etc., that can be associated with switches. You can also define probes to monitor the health of the service end points and can define the forward and reverse interfaces where the traffic policies are applied. ePBR also supports load balancing along with service chaining. ePBR allows you to configure multiple service end points as a part of the service configuration.

Beginning with Cisco NX-OS Release 10.2(1)F, the VRF of every service in a chain may either be unique or may be exactly identical. The service endpoints and interfaces defined for a service, should pertain to the VRF defined for the service.

Applying ePBR to an Interface

Service end-point interfaces having an existing IPv4 PBR policy cannot be used inside an IPv4 ePBR service. Similarly service end-point interfaces having an existing ipv6 PBR policy cannot be used inside an IPv6 ePBR service.

After creating the ePBR service, you must create an ePBR policy. The ePBR policy allows you to define traffic selection, redirection of traffic to the service end point and various fail-action mechanisms on the end point health failure. You may use IP access-list end points with permit access control entries (ACE) to define the traffic of interest to match and take the appropriate action.

The ePBR policy supports multiple ACL match definitions. A match can have multiple services in a chain which can be sequenced by a sequence number. This allows flexibility to add, insert, and modify elements in a chain in a single service policy. In every service sequence, you can define the fail action method such as drop, forward, and bypass. The ePBR policy allows you to specify source or destination-based load balancing and bucket counts in order to have granular load balancing of traffic.

Applying ePBR to an Interface

After creating the ePBR policy you need to apply the policy on an interface. This allows you to define the interface at which the traffic ingresses into the NX-OS or Nexus fabric. You can also apply the policy in both the forward and reverse directions. There may only be two IPv4/IPv6 policies applied to the interface, one in the forward and one in the reverse direction.

Beginning with Cisco NX-OS Release 10.2(1)F, ePBR supports policy application on layer-3 port-channel sub-interfaces

Beginning with Cisco NX-OS Release 10.2(1)F, the interface on which the ePBR policy is applied may be on a different VRF than the VRF of the services in the chain.

ePBR IPv4 policies cannot be applied to an interface on which an IPv4 PBR policy is already applied. ePBR IPv6 policies cannot be applied to an interface on which an IPv6 PBR policy is already applied.

Creating Bucket and Load Balancing

ePBR computes the number of traffic buckets based on the service that has maximum number of service-end-points in the chain. If you configure the load balance buckets, your configuration will have the precedence. ePBR supports load balancing methods of source IP and destination IP but does not support L4-based source or destination load balancing methods.

ePBR Object Tracking, Health Monitoring, and Fail-Action

ePBR creates SLA and Track objects based on the probe types configured in the service and supports various probes and timers such as ICMP, TCP, UDP, DNS, and HTTP. ePBR also supports user defined tracks, which allows you to create tracks with various parameters including milli second probes in associating with ePBR.

ePBR monitors the health of the end points by provisioning IP SLA probes and object tracks to track the IP SLA reachability when you apply the ePBR probe configuration.

You can configure the ePBR probe options for a service or for each of the forward or reverse end points. You can also configure frequency, timeout, retry up and down counts, and source loopback interface so that they can be used for source IP of an IP SLA session. You can define any type of tracks and associate them with the forward or the reverse end points. The same track objects is re-used for all policies using the same ePBR service.

You can define tracks separately and assign the track ID to each service-end point in ePBR. If you do not assign any user-defined track to an endpoint, ePBR will create a track using probe method for the end point. If no probe method is defined at the end point level, the probe method configured for the service level will be used.

ePBR supports the following fail-action mechanisms for its service chain sequences:

- Bypass
- Drop on Fail
- Forward

Bypass of a service sequence indicates that the traffic must be redirected to the next service sequence when there is a failure of the current sequence.

Drop on fail of a service sequence indicates that the traffic must be dropped when all the service-end-points of the service become unreachable.

Forward is the default option and indicates that upon failure of the current service, traffic should use the regular routing tables. This is the default fail-action mechanism.



Note Symmetry is maintained when fail-action bypass is configured for all the services in the service chain. In other fail-action scenarios, when there are one or more failed services, symmetry is not maintained in the forward and the reverse flow.

ePBR Session-based Configuration

ePBR sessions allow addition, deletion or modification of the following aspects of in-service services or policies. The in-service refers to a service that is associated with a policy that has been applied to an active interface or a policy that is being modified and currently configured on an active interface.

- Service endpoints with their interfaces and probes
- Reverse endpoints and probes
- Matches under policies
- Load-balance methods for matches
- Match sequences and fail-action



Note In ePBR Sessions, you cannot move interfaces from one service to another service in the same session. To move interfaces from one service to another service, perform the following steps:

1. Use a session operation to first remove it from the existing service.
2. Use a second session operation to add it to the existing service.

ePBR Multi-Site

Beginning with Cisco NX-OS Release 10.2(1)F, service-chaining in a VXLAN multisite fabric can be achieved by using the following configuration and topology guidelines.

- Endpoints in a service or services in the chain may be distributed across different leaf switches, in the same or different site.
- Every service should be in its unique VRF, which is different from the tenant VRF context in which the ePBR policy is applied.
- To segregate traffic for different tenant VRFs, the VLANs used for the services would be required to be segregated and new services and policies would need to be defined.
- Tenant VRF routes should be leaked to each of the service VRFs on every leaf switch hosting the services, to allow traffic to be routed back at the end of the service chain to its destination, in the tenant VRF.
- VNIs should be symmetrically allocated across different leaf switches and sites.
- The ePBR policy should be enabled on all layer-3 VNIs of the service VRFs being used, on all leaf switches hosting services and on the border leaf or border gateway switches, if it is acting as transit for multi-site.
- The service chain may be isolated to one site entirely, with traffic arriving from different sites. Although this scenario doesn't involve multi-site distribution of service devices, the layer-3 VNIs of the service VRFs on the border gateways or border leafs should only be treated as multi-site transit and the ePBR policy should be applied on them. The ePBR policy should be also applied on the host or tenant facing interfaces in the remote sites where the traffic is arriving from.

ACL Refresh

ePBR session ACL refresh allows you to update the policy generated ACLs, when the user-provided ACL gets modified or added or deleted with ACEs. On the refresh trigger, ePBR will identify the policies that are impacted by this change and create or delete or modify the buckets' generated ACLs for those policies.

For ePBR scale values, see [Cisco Nexus 9000 Series NX-OS Verified Scalability Guide](#).

Guidelines and Limitations for ePBR L3

ePBR has the following guidelines and limitations:

- Beginning with Cisco Nexus NX-OS Release 10.1(2), ePBR with IPv4 and IPv6 is supported on N9K-C93108TC-FX3P switch.
- Beginning with Cisco NX-OS Release 10.1(1) each match statement under ePBR policy can support three action types - redirect, drop, and exclude. There can be only one drop and/or exclude match statement per policy. The ACE rules for the traffic, which needs to be excluded or dropped in the forward as well as the reverse directions, should be manually added to the match access-list that is used with the action of exclude or drop. The statistics for the exclude and drop match access-list may display traffic hit counters for both directions.
- ePBR policies require at least one match with redirect action.

- Beginning with Cisco NX-OS Release 10.1(1), ePBR with IPv4, IPv6 and ePBR over VXLAN are supported on below platform switches: N9K-C9316D-GX, N9K-C93600CD-GX, N9K-C9364C-GX, N9K-C93180YC-FX3S, N9K-C93360YC-FX3, and N9K-C93108TC-FX3P.
- When fail-action is specified in any match statement, probe is mandatory in the configuration.
- Whenever there is OTM track changes ePBR statistics is reset due to RPM reprogramming.
- Do not share the same user defined ACL across multiple match statements in the ePBR configuration.
- Symmetry in traffic is maintained only when fail-action bypass is configured for ePBR Service. For the other fail-actions such as forward/drop in the service chain, symmetry is not maintained for the forward and reverse flow of traffic.
- Unique layer-4 source and destination port parameters should be specified for the match filters if traffic is required to match any source and any destination IP as per the match access-list definition, and is required to be redirected to devices distributed in a VXLAN environment in both forward and reverse directions or service-chained through one-arm devices.
- Feature ePBR and feature ITD cannot co-exist with the same ingress interface.
- With scaled ePBR configuration, it is recommended to remove the policies before you use the **no feature epbr** command.
- It is recommended that you classify probe traffic in a separate CoPP class. Otherwise, probe traffic will go in the default CoPP class and might be dropped causing IP SLA bouncing for probe traffic. For information on CoPP configuration for IP SLA, see [Configuring CoPP for IP SLA Packets](#).
- ePBR is supported on the Cisco Nexus 9500 and Cisco Nexus 9300 platform switches with EX, FX, and FX2 line cards.
- ePBRv4 over VXLAN and NX-OS ePBR are supported on Cisco Nexus 9500 series switches.
- ePBRv6 over VXLAN is not supported on Cisco Nexus 9500 series switches.
- Beginning with Cisco NX-OS Release 9.3(5) Catena feature is deprecated.
- If you want to remove the ePBR service endpoint which is configured to a port-channel that is removed from the system, perform the following steps:
 - Delete the existing ePBR policy.
 - Delete the existing ePBR service.
 - Reconfigure the ePBR service endpoint to the required port-channel.
- Please do not modify the dynamically created access-list entries of ePBR that begin with the name “epbr_”. These access-lists are reserved for ePBR internal use.



Note Modifying these prefix strings can cause the ePBR to not function properly and would impact ISSU.

- Router ACLs may be enabled alongside layer-3 ePBR policies on supported layer-3 interfaces, only when statistics is not enabled for either ePBR policies or the router ACLs. See [Guidelines and Limitations](#)

for Policy-Based Routing in the Policy-based routing chapter of *Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide* for details on this limitation.

- On Cisco Nexus N9K-C9316D-GX, N9K-C93600CD-GX, and N9K-C9364C-GX switches, before performing ISSD from Cisco NX-OS, Release 10.2 and later releases to Release 10.1 disable ePBR policies and proceed with downgrade.
- ePBR policy definition can be applied to a maximum of 32 interfaces of supported interface types across forward and reverse directions.
- Configuration rollback and configuration replace are supported only when the ePBR policy is not associated with any interfaces and the ePBR service definitions are not used in any active ePBR policy in both the source and target configurations. However, configuration rollback and configuration replace do not support policy to interface association and disassociation.

The following guidelines and limitations apply to ePBR over VXLAN feature:

- In VXLAN fabric, service chaining cannot be done to devices within same VLAN. All devices must be present in separate VLANs.
- When every service in the chain is in the same VRF, ePBR is only supported in a single site in a VXLAN multisite fabric.
- When every service in the chain is in the same VRF:
 - Active/Standby chain is supported with two service nodes with no restrictions.
 - Active/Standby chain with three or more service nodes in chain requires no two nodes of different type behind same service leaf.
- In VXLAN fabric you cannot stitch traffic from one service in a leaf and come back later to the same leaf.



Note These restrictions are not applicable if every service in the chain is in a different VRF context.

- The ePBR policy should always be originally applied on host or tenant facing interfaces. The ePBR policy should be applied on Layer-3 VNI interfaces pertaining to the tenant or service VRFs only as the transit interfaces.

The following guidelines and limitations apply to the match ACL feature:

- Only ACEs with the permit method are supported in the ACL. ACEs with any other method (such as deny or remark) are ignored.
- A maximum of 256 permit ACEs are supported in one ACL.
- Layer-4 ACE rules with port operations other than port equality operations are not supported.
- ACEs with object-groups specified as address-groups or port-groups in either source or destination parameters are not supported.

The following guidelines and limitations apply to inter-VRF service chaining:

- Beginning with Cisco NX-OS 10.2(1)F Release, every service in a chain should either exist in the same VRF or completely unique VRFs.
- In version 10.2(1)F, fail-action bypass mechanism is not supported when every service in the chain exists in a unique VRF.
- From Cisco NX-OS 10.2(2)F Release, fail-action bypass is supported when the services in the chain are in unique VRFs.
- If the services are in a different VRF than the VRF context of the interface on which the ePBR policy is applied, the user should ensure that the tenant routes are leaked to every service VRF, in order to ensure that the traffic is able to route back to the tenant VRF, at the end of the service chain.
- From Cisco NX-OS Release 10.2(2)F, PBR allows multiple backup next-hops related to different VRFs to be configured for a route-map sequence. This allows ePBR to enable fail-action bypass from service pertaining to one VRF to another effectively.
- Beginning with Cisco NX-OS Release 10.2(3)F, to minimize traffic disruptions during session operations of endpoint additions, service sequence additions, deletions, and modifications, it is recommended to have load-balance buckets configured ahead and avoid modification to the load-balance configuration. Ensure that the configured buckets for load-balance are greater than the number of endpoints configured in services for every sequence in the chain.

Configuring ePBR L3

Before you begin

Make sure you have configured IP SLA and PBR features before configuring the ePBR feature.

Configuring ePBR Service, Policy, and Associating to an Interface

The following section provides information about configuring the ePBR Service, ePBR Policy, and associating the policy on to an interface.

SUMMARY STEPS

1. **configure terminal**
2. **epbr service *service-name***
3. **vrf *vrf-name***
4. **service-endpoint {ip *ipv4 address* | ipv6 *ipv6 address*} [interface *interface-name interface-number*]**
5. **probe track *track ID***
6. **reverse ip *ip address* interface *interface-name interface-number***
7. **exit**
8. **epbr policy *policy-name***
9. **match { [ip address *ipv4 acl-name*] | [ipv6 address *ipv6 acl-name*] } [redirect | drop | exclude]**
10. **[no] load-balance [method { src-ip | dst-ip }] [buckets *sequence-number*]**
11. ***sequence-number* set service *service-name* [fail-action { bypass | drop | forward }]**
12. **interface *interface-name interface-number***
13. **epbr { ip | ipv6} policy *policy-name* [reverse]**

14. exit

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	epbr service service-name Example: switch(config)# epbr service firewall	Creates a new ePBR service.
Step 3	vrf vrf-name Example: switch(config)# vrf tenant_A	Specifies the VRF for the ePBR service.
Step 4	service-endpoint {ip ipv4 address ipv6 ipv6 address} [interface interface-name interface-number] Example: switch(config-vrf)# service-endpoint ip 172.16.1.200 interface VLAN100	Configures service endpoint for the ePBR service. You can repeat steps 2 to 5 to configure another ePBR service.
Step 5	probe track track ID Example: switch(config-vrf)# probe track 30	Defines a track separately and assign an existing track ID to each service-endpoint in ePBR. You can assign track ID to each endpoint.
Step 6	reverse ip ip address interface interface-name interface-number Example: switch(config-vrf)# reverse ip 172.16.30.200 interface VLAN201	Defines the reverse IP and interfaces where the traffic policies are applied.
Step 7	exit Example: switch(config-vrf)# exit	Exits VRF configuration mode and enters global configuration mode.
Step 8	epbr policy policy-name Example: switch(config)# epbr policy Tenant_A-Redirect	Configures the ePBR policy.
Step 9	match { [ip address ipv4 acl-name] [ipv6 address ipv6 acl-name] } [redirect drop exclude] Example: switch(config)# match ip address WEB	Matches an IPv4 or IPv6 address against an IP or IPv6 ACLs. Redirect is the default action for a match traffic. Drop is used when the traffic needs to be dropped on the incoming interface. Exclude option is used to exclude certain traffic from service-chaining on the incoming interface.

	Command or Action	Purpose
		You can repeat this step to match multiple ACLs based on the requirement.
Step 10	[no] load-balance [method { src-ip dst-ip}] [buckets sequence-number] Example: switch(config)# load-balance method src-ip mask-position 3	Computes the load balance method and the number of buckets to be used by the ePBR service.
Step 11	sequence-number set service service-name [fail-action { bypass drop forward}] Example: switch(config)# set service firewall fail-action drop	Computes the fail-action mechanism.
Step 12	interface interface-name interface-number Example: switch(config)# interface vlan 2010	Configures an interface and enters interface configuration mode.
Step 13	epbr { ip ipv6} policy policy-name [reverse] Example: switch(config-if)# epbr ip policy Tenant_A-Redirect	An interface may be associated at any time with one or more of the following: <ul style="list-style-type: none">• an IPV4 policy in the forward direction• an IPv4 policy in the reverse direction• an IPv6 policy in the forward direction• an IPv6 policy in the reverse direction
Step 14	exit Example: switch(config-if)# end	Exits interface configuration mode and returns to global configuration mode.

Modifying a Service Using ePBR Session

The following steps explain how to modify a service using ePBR session.

SUMMARY STEPS

1. **epbr session**
2. **epbr service service-name**
3. **[no] service-endpoint {ip ipv4 address | ipv6 ipv6 address} [interface interface-name interface-number]**
4. **service-endpoint {ip ipv4 address | ipv6 ipv6 address} [interface interface-name interface-number]**
5. **reverse ip address interface interface-name interface-number**
6. **commit**
7. **abort**

Modifying a Policy Using ePBR Session

DETAILED STEPS

	Command or Action	Purpose
Step 1	epbr session Example: switch(config)# epbr session	Enters ePBR session mode.
Step 2	epbr service service-name Example: switch(config-epbr-sess)# epbr service TCP_OPTIMIZER	Specifies the configured ePBR service in the ePBR session mode.
Step 3	[no] service-endpoint {ip ipv4 address ipv6 ipv6 address} [interface interface-name interface-number] Example: switch(config-epbr-sess-svc)# no service-end-point ip 172.16.20.200 interface VLAN200	Disables the configured service endpoint for the ePBR service.
Step 4	service-endpoint {ip ipv4 address ipv6 ipv6 address} [interface interface-name interface-number] Example: switch(config-epbr-sess-svc)#service-end-point ip 172.16.25.200 interface VLAN200	Modifies the service endpoint and replaces the IP for the ePBR service.
Step 5	reverse ip ip address interface interface-name interface-number Example: switch(config-epbr-sess-svc-ep)# reverse ip 172.16.30.200 interface VLAN201	Defines the reverse IP and interfaces where the traffic policies are applied.
Step 6	commit Example: switch(config-epbr-sess)# commit	Completes the modification of the ePBR service using the ePBR session. Note Restart the ePBR session after you complete this step.
Step 7	abort Example: switch(config-epbr-sess)# abort	Aborts the session and clears or resets the current configuration under the session. Use this command to abandon the current session configuration in case of errors or unsupported configuration identified during commits. Note Restart a new ePBR session after this with the rectified configuration.

Modifying a Policy Using ePBR Session

The following steps explain how to modify a policy using ePBR Session.

SUMMARY STEPS

1. **epbr session**
2. **epbr policy *policy-name***
3. [no] **match { [ip address *ipv4 acl-name*] | [ipv6 address *ipv6 acl-name*] [l2 address *ipv6 acl-name*] } vlan {vlan | vlan range | all} [redirect | drop | exclude]** }
4. **match { [ip address *ipv4 acl-name*] | [ipv6 address *ipv6 acl-name*] [l2 address *ipv6 acl-name*] } vlan {vlan | vlan range | all} [redirect | drop | exclude]** }
5. **sequence-number set service *service-name* [fail-action { bypass | drop | forward}]**
6. [no] **load-balance [method { src-ip | dst-ip}] [buckets *sequence-number*]**
7. **commit**
8. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	epbr session Example: switch(config)# epbr session	Enters ePBR session mode.
Step 2	epbr policy <i>policy-name</i> Example: switch(config-epbr-sess)# epbr policy Tenant_A-Redirect	Specifies the configured ePBR policy in the ePBR session mode.
Step 3	[no] match { [ip address <i>ipv4 acl-name</i>] [ipv6 address <i>ipv6 acl-name</i>] [l2 address <i>ipv6 acl-name</i>] } vlan {vlan vlan range all} [redirect drop exclude] } Example: switch(config-epbr-sess-pol)# no match ip address WEB	Disables the IP address matching against the IP or IPv6 ACLs.
Step 4	match { [ip address <i>ipv4 acl-name</i>] [ipv6 address <i>ipv6 acl-name</i>] [l2 address <i>ipv6 acl-name</i>] } vlan {vlan vlan range all} [redirect drop exclude] } Example: switch(config-epbr-sess-pol)# match ip address HR	Modifies the IP address matching against the IP or IPv6 ACLs.
Step 5	sequence-number set service <i>service-name</i> [fail-action { bypass drop forward}] Example: switch(config-epbr-sess-pol-match)# set service firewall fail-action drop	Adds, modifies, or deletes sequences for a match, or modifies the fail-action for an existing sequence.
Step 6	[no] load-balance [method { src-ip dst-ip}] [buckets <i>sequence-number</i>] Example:	Computes the load balance method and the number of buckets to be used by the ePBR service.

Updating the Access-list Used by ePBR Policies

	Command or Action	Purpose
	switch(config-epbr-sess-pol-match) # load-balance method src-ip mask-position 3	Note On omitting this configuration in the session context while modifying the service-chain for an existing match, the load-balance configuration for the match will be reset to default.
Step 7	commit Example: switch(config-epbr-sess) #commit	Completes the modification of the ePBR policy using the ePBR session.
Step 8	end Example: switch(config-epbr-sess) #end	Exits the ePBR session mode.

Updating the Access-list Used by ePBR Policies

The following steps explain how to update the access-list used by ePBR policies:

SUMMARY STEPS

1. **epbr session access-list *acl-name* refresh**
2. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	epbr session access-list <i>acl-name</i> refresh Example: switch(config) # epbr session access-list WEB refresh	Updates or refreshes the policy generated ACLs.
Step 2	end Example: switch(config) # end	Exits the global configuration mode.

ePBR Show Commands

The following list provides the show commands associated with ePBR.

SUMMARY STEPS

1. **show epbr policy *policy-name* [reverse]**
2. **show epbr statistics *policy-name* [reverse]**
3. **show tech-support epbr**

4. **show running-config epbr**
5. **show startup-config epbr**

DETAILED STEPS

	Command or Action	Purpose
Step 1	show epbr policy <i>policy-name</i> [reverse] Example: switch# show epbr policy Tenant_A-Redirect	Displays information on the ePBR policy applied in forward or reverse direction.
Step 2	show epbr statistics <i>policy-name</i> [reverse] Example: switch# show epbr statistics policy pol2	Displays the ePBR policy statistics.
Step 3	show tech-support epbr Example: switch# show tech-support epbr	Displays the technical support information for ePBR.
Step 4	show running-config epbr Example: switch# show running-config epbr	Displays the running configuration for ePBR.
Step 5	show startup-config epbr Example: switch# show startup-config epbr	Displays the startup configuration for ePBR

Verifying ePBR Configuration

To verify the ePBR configuration, use the following commands:

Command	Purpose
show ip/ipv6 policy vrf <context>	Displays the IPv4/IPv6 route-map policies created for the Layer-3 ePBR policy, at the interfaces at which service chain is applied and the relevant end-point interfaces of the service-chain.
show route-map dynamic <route-map name>	Displays the next-hops configured for traffic re-direction for specific bucket access-lists, used for forwarding traffic at every point in the service chain.
show ip/ipv6 access-list <access-list name> dynamic	Displays the traffic match criteria for a bucket access-list.
show ip sla configuration dynamic	Displays the IP SLA configuration generated by ePBR, for the service-end-points in the chain, when probes are enabled.

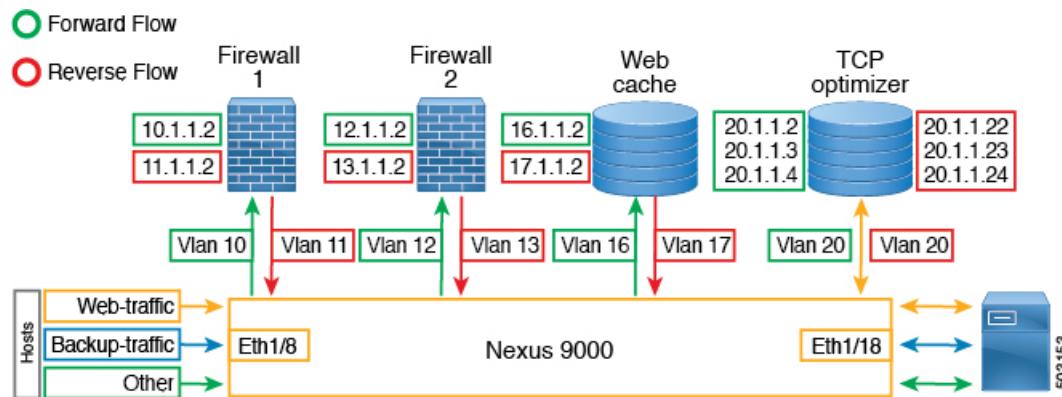
Command	Purpose
show track dynamic	Displays the tracks generated by ePBR, for the service-end-points in the chain, when probes are enabled.

Configuration Examples for ePBR L3

Example: ePBR NX-OS Configuration

The following topology illustrates ePBR NX-OS configuration.

Figure 1: ePBR NX-OS Configuration



Example: Use-Case: Create a Service Chain for Web Traffic in Forward Direction Only

The following configuration example shows how to create a service chain for web traffic in forward direction only.

```

IP access list web_traffic
    10 permit tcp any any eq www

ePBR service FW1
    service-end-point ip 10.1.1.2 interface Vlan10
    reverse interface Vlan11

ePBR service FW2
    service-end-point ip 12.1.1.2 interface Vlan12
    reverse interface Vlan13

ePBR service Web_cache
    service-end-point ip 16.1.1.2 interface Vlan16
    reverse interface Vlan17

ePBR policy tenant_1
    match ip address web_traffic
        10 set service FW1
        20 set service FW2
        30 set service Web_cache

interface Eth1/8
    ePBR ip policy tenant_1
  
```

The following example shows how to verify the configuration of service chain creation for web traffic in forward direction.

```
switch# show ePBR policy tenant_1

Policy-map : tenant_1
  Match clause:
    ip address (access-lists): web-traffic
  Service chain:
    service FW1, sequence 10, fail-action No fail-action
      IP 10.1.1.2
    service FW2, sequence 20, fail-action No fail-action
      IP 12.1.1.2
    service Web_cache, sequence 30, fail-action No fail-action
      IP 16.1.1.2
  Policy Interfaces:
    Eth1/8
```

Example: Use-Case : Load Balance TCP Traffic Using ePBR in Forward Direction Only

The following configuration example shows how to load balance TCP traffic using ePBR in forward direction only.

```
IP access list tcp_traffic
  10 permit tcp any any

ePBR service TCP_Optimizer
  service-interface Vlan20
  service-end-point ip 20.1.1.2
  service-end-point ip 20.1.1.3
  service-end-point ip 20.1.1.4

ePBR policy tenant_1
  match ip address tcp_traffic
    10 set service TCP_Optimizer

interface Eth1/8
  ePBR ip policy tenant_1
```

The following example shows how to verify the configuration of load balance TCP traffic using EPBR in forward direction.

```
switch# show ePBR policy tenant_1

Policy-map : tenant_1
  Match clause:
    ip address (access-lists): tcp_traffic
  Service chain:
    service TCP_Optimizer, sequence 10, fail-action No fail-action
      IP 20.1.1.2
      IP 20.1.1.3
      IP 20.1.1.4
  Policy Interfaces:
    Eth1/8
```

Example: Use-Case: Create a Service Chain for Web Traffic in Both Directions

The following configuration example shows how to create a service chain for web traffic in both forward and reverse directions.

```
IP access list web_traffic
  10 permit tcp any any eq www

ePBR service FW1
```

Configuration Examples for ePBR L3

```

service-end-point ip 10.1.1.2 interface Vlan10
reverse ip 11.1.1.2 interface Vlan11

ePBR service FW2
service-end-point ip 12.1.1.2 interface Vlan12
reverse ip 13.1.1.2 interface Vlan13

ePBR service Web_cache
service-end-point ip 16.1.1.2 interface Vlan16
reverse ip 17.1.1.2 interface Vlan17

ePBR policy tenant_1
match ip address web-traffic
10 set service FW1
20 set service FW2
30 set service Web_cache

interface Eth1/8
ePBR ip policy tenant_1

interface Eth1/18
ePBR ip policy tenant_1 reverse

```

The following example shows how to verify the configuration of service chain creation for web traffic in both forward and reverse directions.

```

switch# show ePBR policy tenant_1

Policy-map : tenant_1
Match clause:
  ip address (access-lists): web-traffic
Service chain:
  service FW1, sequence 10, fail-action No fail-action
    IP 10.1.1.2
  service FW2, sequence 20, fail-action No fail-action
    IP 12.1.1.2
  service Web_cache, sequence 30, fail-action No fail-action
    IP 16.1.1.2
Policy Interfaces:
  Eth1/8

switch# show ePBR policy tenant_1 reverse

Policy-map : tenant_1
Match clause:
  ip address (access-lists): web-traffic
Service chain:
  service Web_cache, sequence 30, fail-action No fail-action
    IP 17.1.1.2
  service FW2, sequence 20, fail-action No fail-action
    IP 13.1.1.2
  service FW1, sequence 10, fail-action No fail-action
    IP 11.1.1.2
Policy Interfaces:
  Eth1/18

```

Example: Use-Case: Load Balance TCP Traffic Using ePBR in Both Directions

The following configuration example shows how to load balance TCP traffic using ePBR in both forward and reverse directions.

```

ePBR service TCP_Optimizer
service-interface Vlan20
service-end-point ip 20.1.1.2
reverse ip 20.1.1.22

```

```

service-end-point ip 20.1.1.3
    reverse ip 20.1.1.23
service-end-point ip 20.1.1.4
    reverse ip 20.1.1.24

ePBR policy tenant_1
    match ip address tcp_traffic
        10 set service TCP_Optimizer

interface Eth1/8
    ePBR ip policy tenant_1

interface Eth1/18
    ePBR ip policy tenant_1 reverse

```

The following example shows how to verify the configuration of load balance TCP traffic using ePBR in both directions.

```

switch# show ePBR policy tenant_1

Policy-map : tenant_1
  Match clause:
    ip address (access-lists): tcp_traffic
  Service chain:
    service TCP_Optimizer, sequence 10, fail-action No fail-action
      IP 20.1.1.2
      IP 20.1.1.3
      IP 20.1.1.4
  Policy Interfaces:
    Eth1/8

switch# show ePBR policy tenant_1 reverse

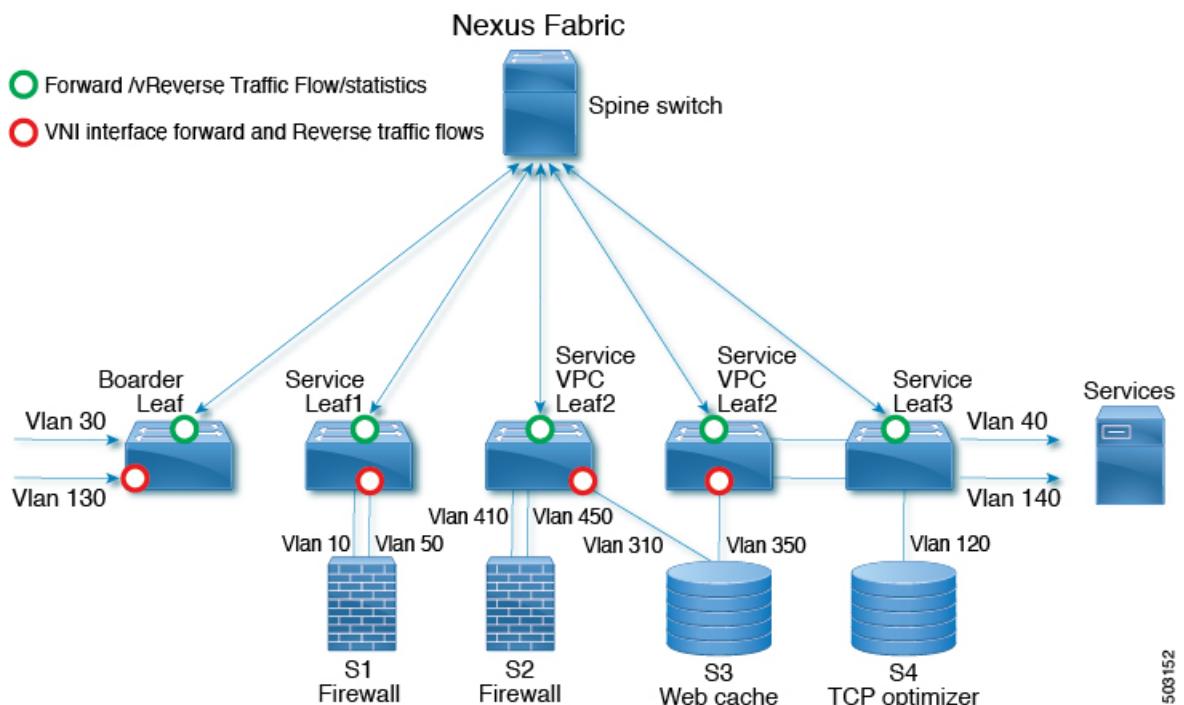
Policy-map : tenant_1
  Match clause:
    ip address (access-lists): tcp_traffic
  Service chain:
    service TCP_Optimizer, sequence 10, fail-action No fail-action
      IP 20.1.1.22
      IP 20.1.1.23
      IP 20.1.1.24
  Policy Interfaces:
    Eth1/18

```

Example: ePBR Policy Creation with VXLAN Fabric

The following example/topology shows how to configure ePBR over VXLAN fabric.

Figure 2: Configuring ePBR over VXLAN Fabric



50152

```

ip access-list acl1
 10 permit ip 30.1.1.0/25 40.1.1.0/25
 20 permit ip 30.1.1.128/25 40.1.1.128/25
ip access-list acl2
 10 permit ip 130.1.1.0/25 140.1.1.0/25
 20 permit ip 130.1.1.128/25 140.1.1.128/25

epbr service s1
  vrf vrf1
    service-end-point ip 10.1.1.2 interface Vlan10
      probe icmp frequency 4 retry-down-count 1 retry-up-count 1 timeout 2 source-interface
      loopback9
      reverse ip 50.1.1.2 interface Vlan50
      probe icmp frequency 4 retry-down-count 1 retry-up-count 1 timeout 2 source-interface
      loopback10

    epbr service s2
      vrf vrf1
        service-end-point ip 41.1.1.2 interface Vlan410
          probe icmp source-interface loopback9
          reverse ip 45.1.1.2 interface Vlan450
          probe icmp source-interface loopback10

    epbr service s3
      vrf vrf1
        service-end-point ip 31.1.1.2 interface Vlan310
          probe http get index.html source-interface loopback9
          reverse ip 35.1.1.2 interface Vlan350
          probe http get index.html source-interface loopback10
  
```

```

epbr service s4
    service-interface Vlan120
    vrf vrf1
        probe udp 6900 control enable source-interface loopback9
        service-end-point ip 120.1.1.2

        reverse ip 120.1.1.2

epbr policy p1
statistics
match ip address acl1
    load-balance buckets 16 method src-ip
    10 set service s1 fail-action drop
    20 set service s2 fail-action drop
    30 set service s4 fail-action bypass
match ip address acl2
    load-balance buckets 8 method dst-ip
    10 set service s1 fail-action drop
    20 set service s3 fail-action forward
    30 set service s4 fail-action bypass
interface Vlan100 - Vxlan L3vni interface to which the policy is applied on all service leafs
    epbr ip policy p1
    epbr ip policy p1 reverse

```

Apply forward policy on ingress interface in border leaf where traffic coming in needs to be service-chained:

```

interface Vlan30 - Traffic matching acl1
    epbr ip policy p1
    int vlan 130 - Traffic matching acl2
    epbr ip policy p1

```

Apply the reverse policy On leaf connected to server if reverse traffic flow needs to be enabled:

```

int vlan 130 - Traffic matching reverse flow for acl1
epbr ip policy p1 rev
int vlan 140 - Traffic matching reverse flow for acl1
epbr ip policy p1 rev

```

Example: Configuring ePBR Service

The following example shows how to configure ePBR service.

```

epbr service FIREWALL
    probe icmp
    vrf TENANT_A
    service-endpoint ip 172.16.1.200 interface VLAN100
        reverse ip 172.16.2.200 interface VLAN101
    service-endpoint ip 172.16.1.201 interface VLAN100
        reverse ip 172.16.2.201 interface VLAN101

epbr service TCP_Optimizer
    probe icmp
    vrf TENANT_A
    service-endpoint ip 172.16.20.200 interface VLAN200
        reverse ip 172.16.30.200 interface VLAN201

```

Example: Configuring ePBR Policy

The following example shows how to configure ePBR Policy.

Configuration Examples for ePBR L3

```

epbr service FIREWALL
  probe icmp
    service-end-point ip 1.1.1.1 interface Ethernet1/1
      reverse ip 1.1.1.2 interface Ethernet1/2
epbr service TCP_Optimizer
  probe icmp
    service-end-point ip 1.1.1.1 interface Ethernet1/3
      reverse ip 1.1.1.4 interface Ethernet1/4
epbr policy Tenant_A-Redirect
match ip address WEB
  load-balance method src-ip
  10 set service FIREWALL fail-action drop
  20 set service TCP_Optimizer fail-action bypass
match ip address APP
  10 set service FIREWALL fail-action drop
match ip address exclude_acl exclude
match ip address drop_acl drop

```

The following example shows the output of show ePBR Policy command with fail-action drop information.

```

switch(config-if)# show epbr policy Tenant_A-Redirect

Policy-map : Tenant_A-Redirect
  Match clause:
    ip address (access-lists): WEB
  action:Redirect
    service FIREWALL, sequence 10, fail-action Drop
      IP 1.1.1.1 track 1 [INACTIVE]
    service TCP_Optimizer, sequence 20, fail-action Bypass
      IP 1.1.1.1 track 2 [INACTIVE]
  Match clause:
    ip address (access-lists): APP
  action:Redirect
    service FIREWALL, sequence 10, fail-action Drop
      IP 1.1.1.1 track 1 [INACTIVE]
  Match clause:
    ip address (access-lists): exclude_acl
  action:Deny
  Match clause:
    ip address (access-lists): drop_acl
  action:Drop
  Policy Interfaces:
    Eth1/4

```

Example: Associating an Interface with ePBR Policy

The following example shows how to configure ePBR Policy.

```

interface vlan 2010
  epbr ip policy Tenant_A-Redirect

interface vlan 2011
  epbr ip policy Tenant_A-Redirect reverse

```

Example: ePBR Policy applied in forward direction

The following example shows the sample Output for policy applied in forward direction.

```

show epbr policy Tenant_A-Redirect
policy-map Tenant_A-Redirect
  Match clause:
    ip address (access-lists): WEB
  Service chain:
    service FIREWALL , sequence 10 , fail-action drop

```

```

ip 172.16.1.200 track 10 [ UP ]
ip 172.16.1.201 track 11 [ DOWN ]
    service TCP_Optimizer, sequence 20 , fail-action bypass
ip 172.16.20.200 track 12 [ UP ]

Match clause:
ip address (access-lists): APP
Service chain:
service FIREWALL , sequence 10 , fail-action drop
ip 172.16.1.200 track 10 [ UP ]
ip 172.16.1.201 track 11 [ DOWN ]

Policy Interfaces:
Vlan 2010

```

Example: ePBR Policy applied in reverse direction

The following example shows the sample Output for policy applied in reverse direction.

```

show epbr policy Tenant_A-Redirect reverse
policy-map Tenant_A-Redirect
Match clause:
ip address (access-lists): WEB

Service chain:
service TCP_Optimizer, sequence 20 , fail-action bypass
ip 172.16.30.200 track 15 [ UP ]

service FIREWALL , sequence 10 , fail-action drop
ip 172.16.2.200 track 13 [ UP ]
ip 172.16.2.201 track 14 [ DOWN ]

Match clause:
ip address (access-lists): APP

Service chain:
service FIREWALL , sequence 10 , fail-action drop
ip 172.16.2.200 track 13 [ UP ]
ip 172.16.2.201 track 14 [ DOWN ]

Policy Interfaces:
Vlan 2011

```

Example: User-defined Track

The following example shows to assign track ID to each end point.

```

epbr service FIREWALL
probe icmp
service-end-point ip 1.1.1.2 interface Ethernet1/21
probe track 30
reverse ip 1.1.1.3 interface Ethernet1/22
probe track 40
service-end-point ip 1.1.1.4 interface Ethernet1/23
reverse ip 1.1.1.5 interface Ethernet1/24

```

Example: Modifying ePBR Service Using ePBR Session

The following example shows to replace the IP of ePBR service and add another service end point.

```

switch(config)#epbr session
switch(config-epbr-sess)#epbr service TCP_OPTIMIZER
switch(config-epbr-sess-svc)# no service-end-point ip 172.16.20.200 interface VLAN200
switch(config-epbr-sess-svc)#service-end-point ip 172.16.25.200 interface VLAN200

```

Additional References

```
switch(config-epbr-sess-svc-ep)# reverse ip 172.16.30.200 interface VLAN201
switch(config-epbr-sess)#commit
```

Example: Modifying ePBR Policy Using EPBR Session

The following example shows to replace the IP of ePBR policy and add a service chain for the modified policy traffic.

```
switch(config)#epbr session
switch(config-epbr-sess)#epbr policy Tenant_A-Redirect
switch(config-epbr-sess-pol)# no match ip address WEB
switch(config-epbr-sess-pol)#match ip address WEB
switch(config-epbr-sess-pol-match)# 10 set service Web-FW fail-action drop load-balance
method src-ip
switch(config-epbr-sess-pol-match)# 20 set service TCP_Optimizer fail-action bypass
switch(config-epbr-sess-pol)#match ip address HR
switch(config-epbr-sess-pol-match)# 10 set service Web-FW
switch(config-epbr-sess-pol-match)# 20 set service TCP_Optimizer
switch(config-epbr-sess)#commit
```

Example: Displaying ePBR Statistics Policy

The following example shows the display of ePBR statistics policy.

```
switch# show epbr statistics policy pol2

Policy-map pol2, match testv6acl

Bucket count: 2

traffic match : epbr_pol2_1_fwd_bucket_1
    two : 0
traffic match : epbr_pol2_1_fwd_bucket_2
    two : 0
```

Additional References

For additional information related to configuring ePBR, see the following sections:

Related Documents

Related Topic	Document Title
Configuring CoPP for IP SLA Packets	<i>Cisco Nexus 9000 Series NX-OS IP SLAs Configuration Guide 9.3(x)</i>
ePBR Licensing	<i>Cisco NX-OS Licensing Guide</i>
ePBR Scale Values	<i>Cisco Nexus 9000 Series NX-OS Verified Scalability Guide</i>

Standards

Standards
No new or modified standards are supported by this feature, and support for existing standards has not been modified by feature.



CHAPTER 4

Configuring ePBR L2

- [Information About ePBR L2, on page 27](#)
- [Guidelines and Limitations for ePBR L2, on page 30](#)
- [Configuring ePBR Service, Policy, and Associating to an Interface, on page 32](#)
- [Modifying a Service Using ePBR Session, on page 34](#)
- [Modifying a Policy Using ePBR Session, on page 35](#)
- [Updating the Access-list Used by ePBR Policies, on page 37](#)
- [ePBR Show Commands, on page 37](#)
- [Verifying ePBR Configuration, on page 38](#)
- [Configuration Examples for ePBR, on page 39](#)

Information About ePBR L2

Enhanced Policy-based Redirect Layer2 (ePBR) in Elastic Services Re-direction (ESR) provides transparent service redirection and service chaining of Layer1/ Layer2 service appliances by leveraging Port ACL and VLAN translation. This action helps achieve service chaining and load-balancing capabilities without adding extra headers and avoids latency in using extra headers.

ePBR enables application-based routing and provides a flexible, device-agnostic policy-based redirect solution without impacting application performance. The ePBR service flow includes the following tasks:

Configuring ePBR Service and Policy

You must first create an ePBR service which defines the attributes of service end points. Service end points are the service appliances such as firewall, IPS, etc., that can be associated with switches. You can also define probes to monitor the health of the service end points and can define the forward and reverse interfaces where the traffic policies are applied. ePBR also supports load balancing along with service chaining. ePBR allows you to configure multiple service end points as a part of the service configuration.

After creating the ePBR service, you must create an ePBR policy. The ePBR policy allows you to define traffic selection, redirection of traffic to the service end point and various fail-action mechanisms on the end point health failure. You may use IP access-list end points with permit access control entries (ACE) to define the traffic of interest to match and take the appropriate action.

The ePBR policy supports multiple ACL match definitions. A match can have multiple services in a chain which can be sequenced by a sequence number. This allows flexibility to add, insert, and modify elements in a chain in a single service policy. In every service sequence, you can define the fail action method such as

drop, forward, and bypass. The ePBR policy allows you to specify source or destination-based load balancing and bucket counts in order to have granular load balancing of traffic.

Applying ePBR to an L2 Interface

After creating the ePBR policy you need to apply the policy on an interface. This allows you to define the interface at which the traffic ingresses the NX-OS switch and the interface through which traffic needs to exit the switch after redirection or service-chaining. You can also apply the policy in both the forward and reverse directions into the NX-OS switch.

Enabling Production Interfaces as Access Port

If the service-chaining switch is inserted in between the two L3 routers for traffic redirection, the production interfaces are enabled as access port with the following limitations:

- You must use the VLAN of the port as part of the match configuration.
- It is limited to mac-learn disable mode.

Enabling Production Interfaces as Trunk Ports

Production interfaces may be configured as trunk ports. The VLANs of the incoming traffic that needs to be service-chained that is trunked by the interfaces must be configured as part of the match configuration.

Alternatively, using 'vlan all' in the match configuration will allow any traffic pertaining to any incoming VLANs on the interface to be matched and service chained.

Creating Bucket and Load Balancing

ePBR computes the number of traffic buckets based on the service that has maximum number of service-end-points in the chain. If you configure the load balance buckets, your configuration will have the precedence. ePBR supports load balancing methods of source IP and destination IP but does not support L4-based source or destination load balancing methods.

ePBR Object Tracking, Health Monitoring, and Fail-Action

Layer-2 ePBR performs link state monitoring of the service end-points by default. The user may additionally enable CTP (Configuration Testing Protocol) if supported by the service.

You can configure the ePBR probe options for a service or for each of the forward or reverse end points. You can also configure frequency, timeout, and retry up and down counts. The same track objects is re-used for all policies using the same ePBR service.

If no probe method is defined at the end point level, the probe method configured for the service level will be used.

ePBR supports the following fail-action mechanisms for its service chain sequences:

- Bypass
- Drop on Fail

- Forward

Bypass of a service sequence indicates that the traffic must be redirected to the next service sequence when there is a failure of the current sequence.

Drop on fail of a service sequence indicates that the traffic must be dropped when all the service-end-points of the service become unreachable.

Forward is the default option and indicates that upon failure of the current service, traffic should forward to the egress interface. This is the default fail-action mechanism.



Note Symmetry is maintained when fail-action bypass is configured for all the services in the service chain. In other fail-action scenarios, when there are one or more failed services, symmetry is not maintained in the forward and the reverse flow.

ePBR Session-based Configuration

ePBR sessions allow addition, deletion or modification of the following aspects of in-service services or policies. The in-service refers to a service that is associated with a policy that has been applied to an active interface or a policy that is being modified and currently configured on an active interface.

- Service endpoints with their interfaces and probes
- Reverse endpoints and probes
- Matches under policies
- Load-balance methods for matches
- Match sequences and fail-action



Note In ePBR Sessions, you cannot move interfaces from one service to another service in the same session. To move interfaces from one service to another service, perform the following steps:

1. Use a session operation to first remove it from the existing service.
2. Use a second session operation to add it to the existing service.

ACL Refresh

ePBR session ACL refresh allows you to update the policy generated ACLs, when the user-provided ACL gets modified or added or deleted with ACEs. On the refresh trigger, ePBR will identify the policies that are impacted by this change and create or delete or modify the buckets' generated ACLs for those policies.

For ePBR scale values, see [Cisco Nexus 9000 Series NX-OS Verified Scalability Guide](#).

Guidelines and Limitations for ePBR L2

ePBR has the following guidelines and limitations:

- When fail-action is specified in any match statement, probe is mandatory in the configuration.
- To disable MAC learning on the switch, use the command **mac-learn disable**.
- Do not share the same user defined ACL across multiple match statements in the ePBR configuration.
- Symmetry in traffic is maintained only when fail-action bypass is configured for ePBR Service. For the other fail-actions such as forward/drop in the service chain, symmetry is not maintained for the forward and reverse flow of traffic.
- Feature ePBR and feature ITD cannot co-exist with the same ingress interface.
- With scaled ePBR configuration, it is recommended to remove the policies before you use the **no feature epbr** command.
- ePBRv6 over VXLAN is not supported on Cisco Nexus 9500 series switches.
- If you want to remove the ePBR service endpoint which is configured to a port-channel that is removed from the system, perform the following steps:
 1. Delete the existing ePBR policy.
 2. Delete the existing ePBR service.
 3. Reconfigure the ePBR service endpoint to the required port-channel.
- Please do not modify the dynamically created access-list entries of ePBR that begin with the name “epbr_”. These access-lists are reserved for ePBR internal use.



Note Modifying these prefix strings can cause the ePBR to not function properly and would impact ISSU.

- All redirection rules are programmed in ACL TCAM using ing-ifacl region. This region needs to be carved and allocated prior to the application of ePBR L2 policies.



Note For steps on how to carve TCAM region, refer to the **Configuring IP ACLs** section of *Cisco Nexus 9000 Series NX-OS Security Configuration Guide*.

- ePBR policies require at least one match with redirect action.
- ePBR L2 requires a VLAN range to be reserved for VLAN translation and Q-in-Q. It is recommended that this range does not overlap with the VLANs used for traffic match configuration.
- The ePBR 'infra' VLANs should be reserved prior to the application of the ePBR Layer-2 policies.
- For production interfaces configured as trunk ports, enable VLAN trunking only for the VLANs specified in the ePBR 'infra vlan' range.

- ePBR L2 expects the service appliance to be configured to forward the packet as is without modifying or stripping the VLAN headers.
- ePBR L2 policies can only be applied on a single interface in the forward and single interface in the reverse direction. The policy will need to be replicated in order to similarly service chain on a different pair of interfaces.
- Each match in an ePBR L2 policy needs to have a unique match VLAN or unique VLAN range when applied on trunk interfaces. Only a single match with 'vlan all' can exist in a policy that is applied on trunk interfaces.
- ePBR L2 policy definition can be applied to a maximum of 32 interfaces of supported interface types across forward and reverse directions.
- In order to load-balance between multiple service devices and uniquely detect failure of these devices via CTP health-checks, each service device should be defined as a unique endpoint in the ePBR service.
- Bucket-based load-balance is not supported for layer-2 matches in the ePBR policy.
- In order to service-chain or redirect IPv6 traffic such as Neighbor discovery, ICMPv6 aces with protocol types of ND-NA and ND-NS should be explicitly defined in the user-defined match access-list.
- In order to service-chain or redirect Layer-2 traffic for protocols such as ARP (0x806), VN-tag (0x8926), FCOE (0x8906), MPLS Unicast (0x8847), MPLS Multicast (0x8848), the protocol information should be explicitly added to the ACEs inside the user-defined match access-list.
- Layer-2 ePBR does not support service-chaining or redirection, enforcement of exclude and deny actions for control traffic (IP, IPv6 and Layer-2) that is typically copied or redirected to the supervisor on Cisco NX-OS 9000 series switches.
- Defaulting ePBR production and/or service interfaces while they are in use should be avoided to prevent any unintended behavior.
- Configuration rollback and configuration replace are supported only when the ePBR policy is not associated with any interfaces and the ePBR service definitions are not used in any active ePBR policy in both the source and target configurations. However, configuration rollback and configuration replace do not support policy to interface association and disassociation.

The following guidelines and limitations apply to the match ACL feature:

- Only ACEs with the permit method are supported in the ACL. ACEs with any other method (such as deny or remark) are ignored.
- A maximum of 256 permit ACEs are supported in one ACL.
- Layer-4 ACE rules with port operations other than port equality operations are not supported.

The following guidelines and limitations apply to inter-VRF service chaining:

- Beginning with Cisco NX-OS Release 10.2(3)F, to minimize traffic disruptions during session operations of endpoint additions, service sequence additions, deletions, and modifications, it is recommended to have load-balance buckets configured ahead and avoid modification to the load-balance configuration. Ensure that the configured buckets for load-balance are greater than the number of endpoints configured in services for every sequence in the chain.

Configuring ePBR Service, Policy, and Associating to an Interface

The following section provides information about configuring the ePBR Service, ePBR Policy, and associating the policy on to an interface.

SUMMARY STEPS

1. **configure terminal**
2. **[no] epbr infra vlans [vlan range]**
3. **epbr service *service-name* type l2**
4. **mode [full duplex | half duplex]**
5. **probe {ctp} [frequency seconds] [timeout seconds] [retry-down-count *count*] retry-up-count *count*]**
6. **service-endpoint [interface *interface-name* *interface-number*]**
7. **reverse interface *interface-name* *interface-number***
8. **exit**
9. **epbr policy *policy-name***
10. **match { [ip address *ipv4 acl-name*] | [ipv6 address *ipv6 acl-name*] | [l2 address *l2 acl-name*] } {drop | exclude | redirect | vlan{vlan | vlan range | all}}**
11. **[no] load-balance [method { src-ip | dst-ip }] [buckets *count*]**
12. **sequence-number set service *service-name* [fail-action { bypass | drop | forward }]**
13. **interface *interface-name* *interface-number***
14. **epbr {l2} policy *policy-name* egress-interface *interface-name* [reverse]**
15. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config) #</pre>	Enters configuration mode.
Step 2	[no] epbr infra vlans [vlan range]	VLAN range is using to indicate the VLANs reserved for selective dot1q translation while redirecting to the service devices.
Step 3	epbr service <i>service-name</i> type l2 Example: <pre>switch(config)# epbr service firewall type l2</pre>	Creates a new ePBR L2 service.
Step 4	mode [full duplex half duplex]	Configures the service to be in half-duplex or full-duplex mode.
Step 5	probe {ctp} [frequency seconds] [timeout seconds] [retry-down-count <i>count</i>] retry-up-count <i>count</i>	Configures the probe for the ePBR service.

	Command or Action	Purpose
	<p>Example:</p> <pre>switch(config) # probe icmp</pre>	<p>The options are as follows:</p> <ul style="list-style-type: none"> frequency—Specifies the frequency of the probe in seconds. The range is from 1 to 604800. retry-down-count—Specifies the number of recounts undertaken by the probe when the node goes down. The range is from 1 to 5. retry-up-count —Specifies the number of recounts undertaken by the probe when the node comes back up. The range is from 1 to 5. timeout —Specifies the length of the timeout period in seconds. The range is from 1 to 604800.
Step 6	<p>service-endpoint [interface <i>interface-name interface-number</i>]</p> <p>Example:</p> <pre>switch(config-epbr-svc) # service-end-point interface Ethernet1/3</pre>	<p>Configures service endpoint for the ePBR service.</p> <p>You can repeat steps 2 to 5 to configure another ePBR service.</p>
Step 7	<p>reverse interface <i>interface-name interface-number</i></p> <p>Example:</p> <pre>switch(config-epbr-fwd-svc) # reverse interface Ethernet1/4</pre>	Defines the reverse interface where the traffic policies are applied.
Step 8	<p>exit</p> <p>Example:</p> <pre>switch(config-epbr-reverse-svc) # exit switch(config-epbr-fwd-svc) # exit switch(config-epbr-svc) # exit switch(config) #</pre>	Exits ePBR service configuration mode and enters global configuration mode.
Step 9	<p>epbr policy <i>policy-name</i></p> <p>Example:</p> <pre>switch(config) # epbr policy Tenant_A-Redirect</pre>	Configures the ePBR policy.
Step 10	<p>match { [ip address <i>ipv4 acl-name</i>] [ipv6 address <i>ipv6 acl-name</i>] [l2 address <i>l2 acl-name</i>] } {drop exclude redirect vlan{vlan vlan range all}}</p> <p>Example:</p> <pre>switch (config) # match ip address WEB vlan 10</pre>	<p>Matches an IPv4, or IPv6 address, or a mac address against an IP, or IPv6, or MAC ACLs. Redirect is the default action for a match traffic. Drop is used when the traffic needs to be dropped on the incoming interface. Exclude option is used to exclude certain traffic from service-chaining on the incoming interface.</p> <p>You can repeat this step to match multiple ACLs based on the requirement.</p>
Step 11	<p>[no] load-balance [method { src-ip dst-ip }] [buckets <i>count</i>]</p> <p>Example:</p>	Computes the load balance method and the number of buckets to be used by the ePBR service.

Modifying a Service Using ePBR Session

	Command or Action	Purpose
	switch(config) # load-balance method src-ip mask-position 3	
Step 12	sequence-number set service service-name [fail-action { bypass drop forward}] Example: switch(config) # set service firewall fail-action drop	Configures the fail-action mechanism.
Step 13	interface interface-name interface-number Example: switch(config) # interface Ethernet1/1	Enters into interface configuration mode.
Step 14	epbr {l2} policy policy-name egress-interface interface-name [reverse] Example: epbr l2 policy Tenant_A_Redirect egress-interface Ethernet1/2	An interface may be associated at any time with one forward policy and one reverse policy of the following: <ul style="list-style-type: none">• an IPV4 policy in the forward direction• an IPv4 policy in the reverse direction• an IPv6 policy in the forward direction• an IPv6 policy in the reverse direction• a l2 policy in the forward direction• a l2 policy in the reverse direction
Step 15	exit Example: switch(config-if) # end	Exits policy configuration mode and returns to global mode.

Modifying a Service Using ePBR Session

The following steps explain how to modify a service using ePBR session.

SUMMARY STEPS

1. **epbr session**
2. **epbr service service-name type l2**
3. **[no] service-endpoint [interface interface-name]**
4. **service-endpoint [interface interface-name]**
5. **reverse [interface interface-name]**
6. **commit**
7. **abort**

DETAILED STEPS

	Command or Action	Purpose
Step 1	epbr session Example: switch(config)# epbr session	Enters ePBR session mode.
Step 2	epbr service <i>service-name type l2</i> Example: switch(config-epbr-sess)# epbr service TCP_OPTIMIZER	Specifies the configured ePBR service in the ePBR session mode.
Step 3	[no] service-endpoint [interface <i>interface-name</i>] Example: switch(config-epbr-sess-svc)# no service-end-point interface ethernet 1/3	Disables the configured service endpoint for the ePBR service.
Step 4	service-endpoint [interface <i>interface-name</i>] Example: switch(config-epbr-sess-svc)# service-end-point interface ethernet 1/15	Add a service endpoint to the service.
Step 5	reverse [interface <i>interface-name</i>] Example: switch(config-epbr-sess-fwd-svc)# reverse interface ethernet 1/4	Defines the reverse interfaces where the traffic policies are applied.
Step 6	commit Example: switch(config-epbr-sess)#commit	Completes the modification of the ePBR service using the ePBR session. Note Restart the ePBR session after you complete this step.
Step 7	abort Example: switch(config-epbr-sess)# abort	Aborts the session and clears or resets the current configuration under the session. Use this command to abandon the current session configuration in case of errors or unsupported configuration identified during commits. Note Restart a new ePBR session after this with the rectified configuration.

Modifying a Policy Using ePBR Session

The following steps explain how to modify a policy using ePBR Session.

SUMMARY STEPS

1. **epbr session**

Modifying a Policy Using ePBR Session

2. **epbr policy *policy-name***
3. [no] **match { [ip address *ipv4 acl-name*] | [ipv6 address *ipv6 acl-name*] | [l2 address *mac acl-name*] } vlan {all | *vlan-id* | *vlan-id-range*}**
4. **match { [ip address *ipv4 acl-name*] | [ipv6 address *ipv6 acl-name*] | [l2 address *mac acl-name*] } vlan {all | *vlan-id* | *vlan-id-range*}**
5. **sequence-number set service *service-name* [fail-action { bypass | drop | forward }]**
6. [no] **load-balance [method { src-ip | dst-ip }] [buckets *count*]**
7. **commit**
8. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	epbr session	
Step 2	epbr policy <i>policy-name</i> Example: <pre>switch(config-epbr-sess)# epbr policy Tenant_A-Redirect</pre>	Specifies the configured ePBR policy in the ePBR session mode.
Step 3	[no] match { [ip address <i>ipv4 acl-name</i>] [ipv6 address <i>ipv6 acl-name</i>] [l2 address <i>mac acl-name</i>] } vlan {all <i>vlan-id</i> <i>vlan-id-range</i>} Example: <pre>switch(config-epbr-sess-pol)# no match ip address WEB</pre>	Disables the match against IP, IPv6, or L2 ACLs.
Step 4	match { [ip address <i>ipv4 acl-name</i>] [ipv6 address <i>ipv6 acl-name</i>] [l2 address <i>mac acl-name</i>] } vlan {all <i>vlan-id</i> <i>vlan-id-range</i>} Example: <pre>switch(config-epbr-sess-pol)# match ip address HR</pre>	Modifies the match against the IP, IPv6 or L2 ACLs.
Step 5	sequence-number set service <i>service-name</i> [fail-action { bypass drop forward }] Example: <pre>switch(config-epbr-sess-pol-match)# set service firewall fail-action drop</pre>	Configures the fail-action mechanism.
Step 6	[no] load-balance [method { src-ip dst-ip }] [buckets <i>count</i>] Example: <pre>switch(config)# load-balance method src-ip mask-position 3</pre>	Configures the load-balance method and buckets for the match. Note On omitting this configuration in the session context while modifying the service-chain for an existing match, the load-balance configuration for the match will be reset to default.

	Command or Action	Purpose
Step 7	commit Example: switch(config-epbr-sess)#commit	Completes the modification of the ePBR policy using the ePBR session.
Step 8	end Example: switch(config-epbr-sess)#end	Exits the ePBR session mode.

Updating the Access-list Used by ePBR Policies

The following steps explain how to update the access-list used by ePBR policies:

SUMMARY STEPS

1. **epbr session access-list *acl-name* refresh**
2. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	epbr session access-list <i>acl-name</i> refresh Example: switch(config)# epbr session access-list WEB refresh	Updates or refreshes the policy generated ACLs.
Step 2	end Example: switch(config)# end	Exits the global configuration mode.

ePBR Show Commands

The following list provides the show commands associated with ePBR.

SUMMARY STEPS

1. **show epbr policy *policy-name* [reverse]**
2. **show epbr statistics *policy-name* [reverse]**
3. **show tech-support epbr**
4. **show running-config epbr**
5. **show startup-config epbr**

DETAILED STEPS

	Command or Action	Purpose
Step 1	show epbr policy <i>policy-name</i> [reverse] Example: switch# show epbr policy Tenant_A-Redirect	Displays information on the ePBR policy applied in forward or reverse direction.
Step 2	show epbr statistics <i>policy-name</i> [reverse] Example: switch# show ePBR statistics policy pol2	Displays the ePBR policy statistics.
Step 3	show tech-support epbr Example: switch# show tech-support epbr	Displays the technical support information for ePBR.
Step 4	show running-config epbr Example: switch# show running-config epbr	Displays the running configuration for ePBR.
Step 5	show startup-config epbr Example: switch# show startup-config epbr	Displays the startup configuration for ePBR

Verifying ePBR Configuration

To verify the ePBR configuration, use the following commands:

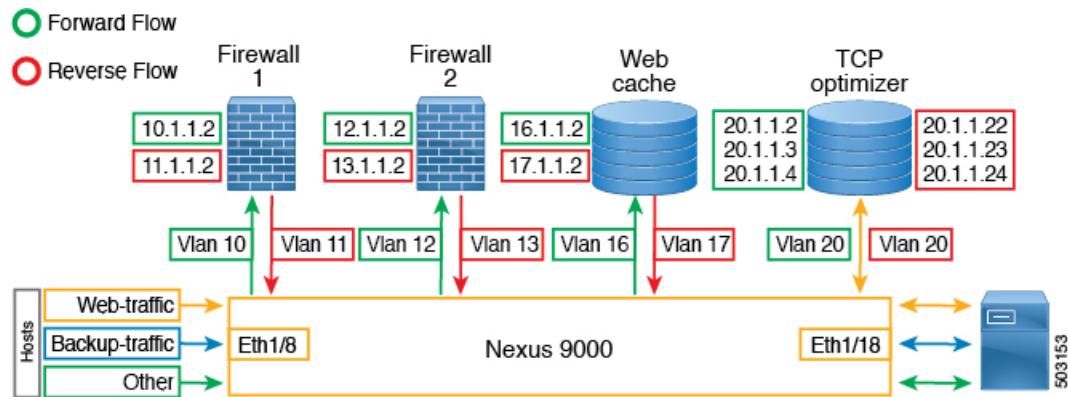
Command	Purpose
show ip access-list <access-list name> dynamic	Displays the traffic match criteria for a bucket access-list.
show ip sla configuration dynamic	Displays the IP SLA configuration generated by ePBR, for the service-end-points in the chain, when probes are enabled.
show track dynamic	Displays the tracks generated by ePBR, for the service-end-points in the chain, when probes are enabled.
show ip access-list summary	Displays the summary of the traffic match criteria for a bucket access-list.
show [ip ipv6 mac] access-lists dynamic	Displays the dynamic entries of match criteria.

Configuration Examples for ePBR

Example: ePBR NX-OS Configuration

The following topology illustrates ePBR NX-OS configuration:

Figure 3: ePBR NX-OS Configuration



Example: Service Configuration for Access and Trunk Ports

The following configuration example shows how to perform service configuration for access and trunk ports:

```
epbr infra vlans 100-200

epbr service app_1 type 12
    service-end-point interface Ethernet1/3
    reverse    interface Ethernet1/4

epbr service app_2 type 12
    probe ctp frequency 2 retry-down-count 1 retry-up-count 1 timeout 1
    service-end-point interface port-channel10
    reverse    interface port-channel11

epbr service app_3 type 12
    probe ctp frequency 2 retry-down-count 1 retry-up-count 1 timeout 1
    service-end-point interface Ethernet1/9
    reverse    interface Ethernet1/10

epbr service app_4 type 12
    probe ctp frequency 2 retry-down-count 1 retry-up-count 1 timeout 1
    service-end-point interface port-channel12
    reverse    interface port-channel13
```

Example: Configuring Access Ports

The following example shows how to configure access ports:

```
epbr policy p1
    statistics
    match ipv6 address flow2 vlan 10
        load-balance buckets 2
        10 set service app_1
        20 set service app_3
        25 set service app_4
        30 set service app_2
    match 12 address flow3 vlan 10
```

Configuration Examples for ePBR

```

20 set service app_2
25 set service app_4
50 set service app_3
match ip address flow1 vlan 10
10 set service app_1
15 set service app_3
20 set service app_2

interface Ethernet1/1
switchport
switchport access vlan 10
no shutdown
epbr 12 policy p1 egress-interface Ethernet1/2

interface Ethernet1/2
switchport
switchport access vlan 10
no shutdown
epbr 12 policy p1 egress-interface Ethernet1/1 reverse

```

Example: Configuring Trunk Ports

The following configuration example shows how to configure trunk ports:

```

epbr policy p3
statistics
match ip address flow1 vlan 10
load-balance buckets 2
10 set service app_1
20 set service app_2
match ipv6 address flow2 vlan 20
load-balance buckets 2
10 set service app_3
20 set service app_4
match 12 address flow3 vlan 30
10 set service app_1
20 set service app_2

interface Ethernet1/27
switchport
switchport mode trunk
no shutdown
epbr 12 policy p3 egress-interface Ethernet1/28

interface Ethernet1/28
switchport
switchport mode trunk
no shutdown
epbr 12 policy p3 egress-interface Ethernet1/27 reverse

Collecting statistics

```

Collecting statistics:

```

itd-san-2# show epbr statistics policy p1

Policy-map p1, match flow2

Bucket count: 2

traffic match : bucket 1
    app_1 : 8986 (Redirect)
    app_3 : 8679 (Redirect)

```

```
    app_4 : 8710 (Redirect)
    app_2 : 8725 (Redirect)
traffic match : bucket 2
    app_1 : 8696 (Redirect)
    app_3 : 8680 (Redirect)
    app_4 : 8711 (Redirect)
    app_2 : 8725 (Redirect)

Policy-map p1, match flow3

Bucket count: 1

traffic match : bucket 1
    app_2 : 17401 (Redirect)
    app_4 : 17489 (Redirect)
    app_3 : 17461 (Redirect)

Policy-map p1, match flow1
```

```
Bucket count: 1

traffic match : bucket 1
    app_1 : 17382 (Redirect)
    app_3 : 17348 (Redirect)
    app_2 : 17411 (Redirect)
```

Example: Viewing ePBR Policy

The following example shows how to view an ePBR policy:

```
show epbr policy p3

Policy-map : p3
Match clause:
ip address (access-lists): flow1
action:Redirect
service app_1, sequence 10, fail-action No fail-action
Ethernet1/3 track 4 [UP]
service app_2, sequence 20, fail-action No fail-action
port-channel10 track 10 [UP]
Match clause:
ipv6 address (access-lists): flow2
action:Redirect
service app_3, sequence 10, fail-action No fail-action
Ethernet1/9 track 13 [UP]
service app_4, sequence 20, fail-action No fail-action
port-channel12 track 3 [UP]
Match clause:
layer-2 address (access-lists): flow3
action:Redirect
service app_1, sequence 10, fail-action No fail-action
Ethernet1/3 track 4 [UP]
service app_2, sequence 20, fail-action No fail-action
port-channel10 track 10 [UP]
Policy Interfaces:
egress-interface Eth1/28
```

