ACI Route-Profile Usage

Contents

Route-Profile Overview

-2.3(1) Apic SW was used for all testing

-Export route-control enforcement is assumed.

Route-profiles are used in ACI to apply some sort of policy to routes. It consists of a match rule defining the routes that the policy should be applied to and a set rule, which defines how the route attributes should be changed. For example, a route-profile would be used to match a specific prefix and change the OSPF metric-type to 1. The available criteria to match and set on is based upon what is supported in each ACI version.

Route-profiles can be applied at several different levels depending on what your goal is. These include:
-The Bridge Domain L3 Configuration
-The Bridge Domain Subnet Configuration
-The default-import and default-export policies configured under the l3out
-The L3out EPG (network) in the import or export direction. Furthermore the route-profile can be applied to specific L3out EPG subnets rather than the whole EPG.
-The Interleak Policy configured at the l3out level

Note that route-profiles can be configured in the import direction but the configuration won't take effect unless "Import" Route Control Enforcement is selected at the L3out Level

Configurating a Route-Profile

A route-profile can be configured underneath a specific l3out or under 'External Routed Networks'. If the route-profile is being used for an Interleak Policy then it should be applied under 'External Routed Networks'. For all other uses the route-profile should be configured under the l3out where the policy will be applied.

When configuring the Route-Profile you will see the below window:

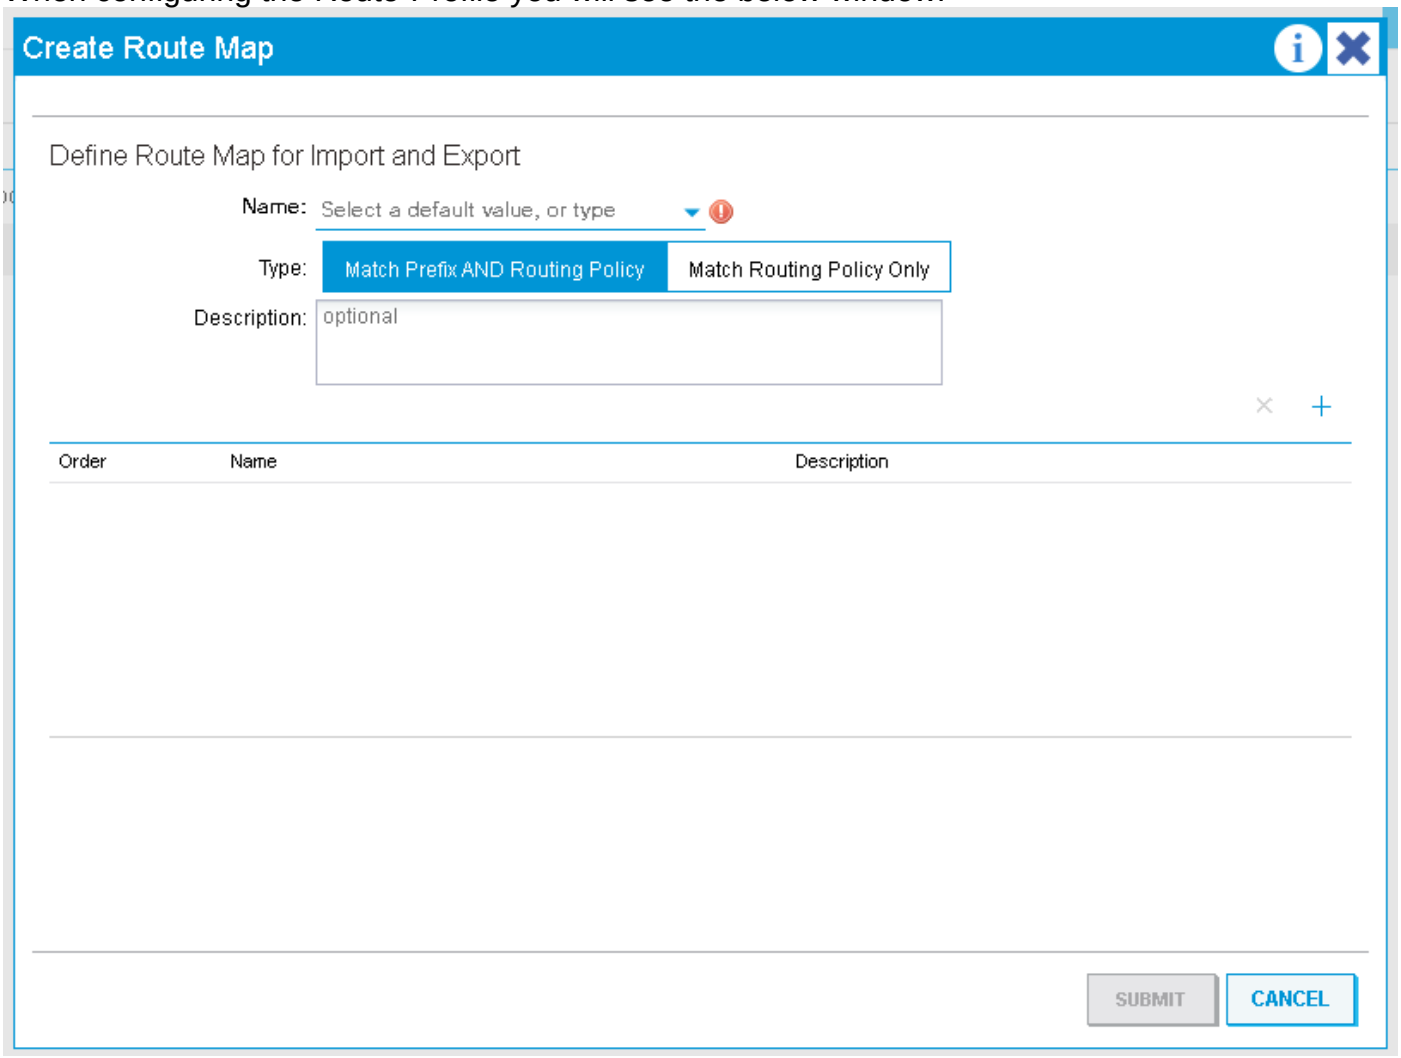

You will have the option to choose between **"Match Prefix and Routing Policy**" and "**Match Routing Policy Only**". These options will take affect depending on what level the route-profile is applied to. Generally speaking though "Match Prefix and Routing Policy" defines the profile as 'combinable'. This means that every match rule that is defined will implicitly include the BD subnets that are set to 'advertise externally' and anything else that is explicitly matched by the match rule. "Match Routing Policy Only" makes the route-profile 'non-combinable'. This means that the profile will only match what is explicitly matched by the match rules. BD subnets are not implicity included. When applied at the External EPG level 'combinable' means that "export route-control subnets" are implicitly matched in each rule rather than BD subnets.

A route-profile requires contexts:

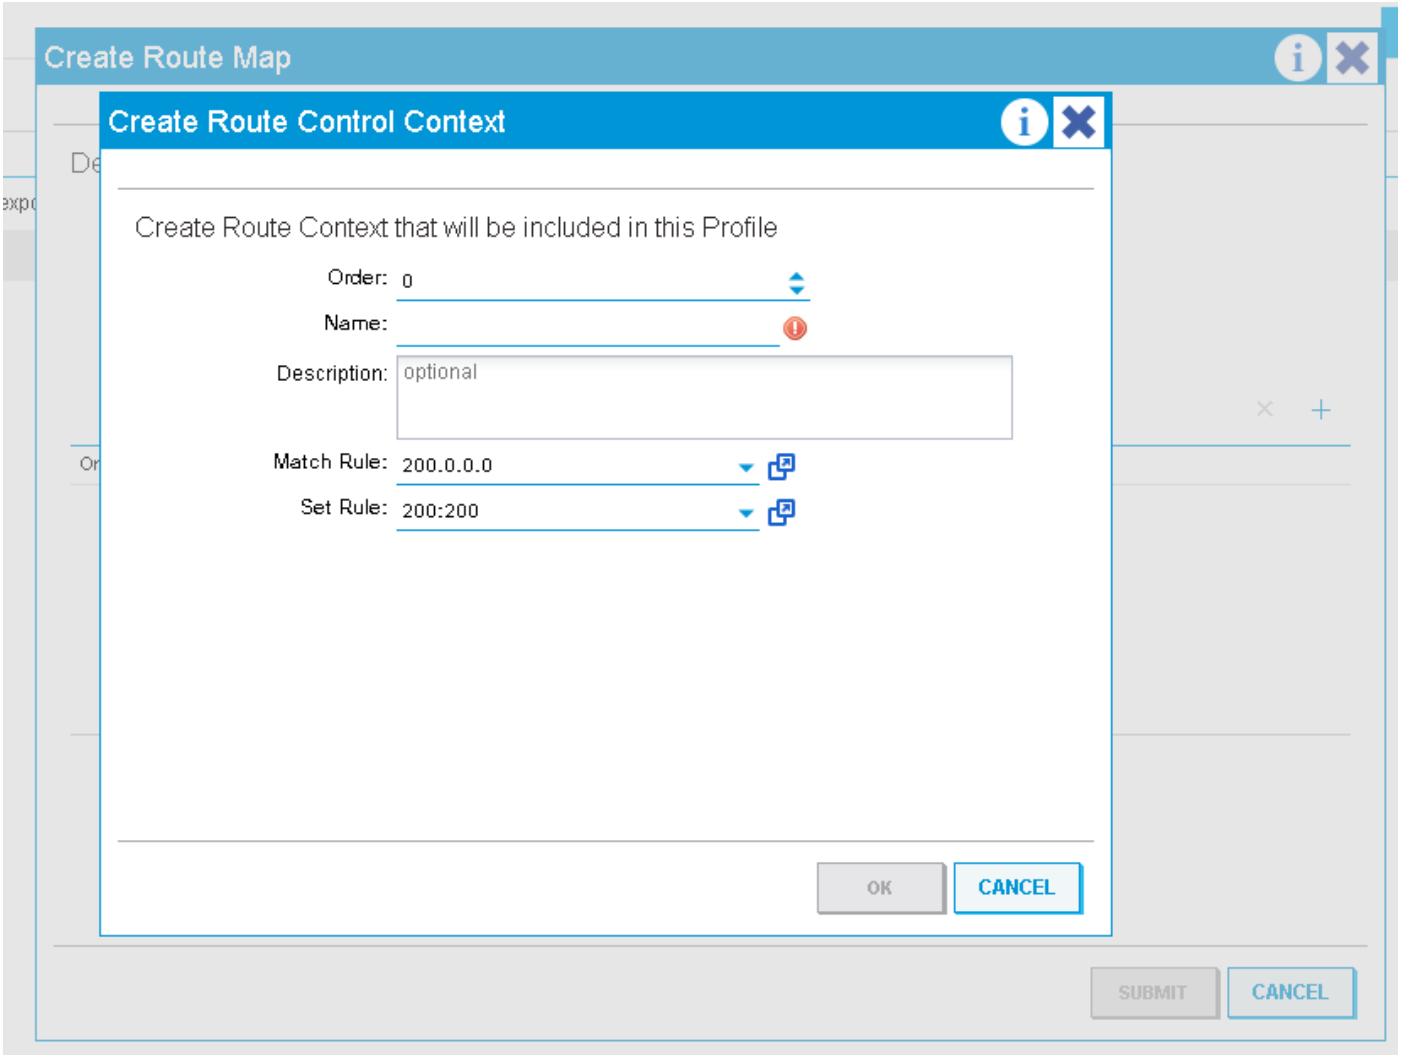

A context is an object that contains a Match rule and a Set rule. Each context has an order (0-9) which defines the order in which the contexts should be evaluated if there are more than one. Once a route-profile is created with at least one context it can be applied.

Applying a Route-Profile at the Bridge Domain Level

A route-profile at the Bridge Domain level is typically used to apply a policy to all subnets defined under a specific BD. To configure this go to 'L3 Configurations' under the Bridge Domain, select the L3out that will apply the policy when advertising the Subnet, and then select the route-profile that is configured under that l3out.

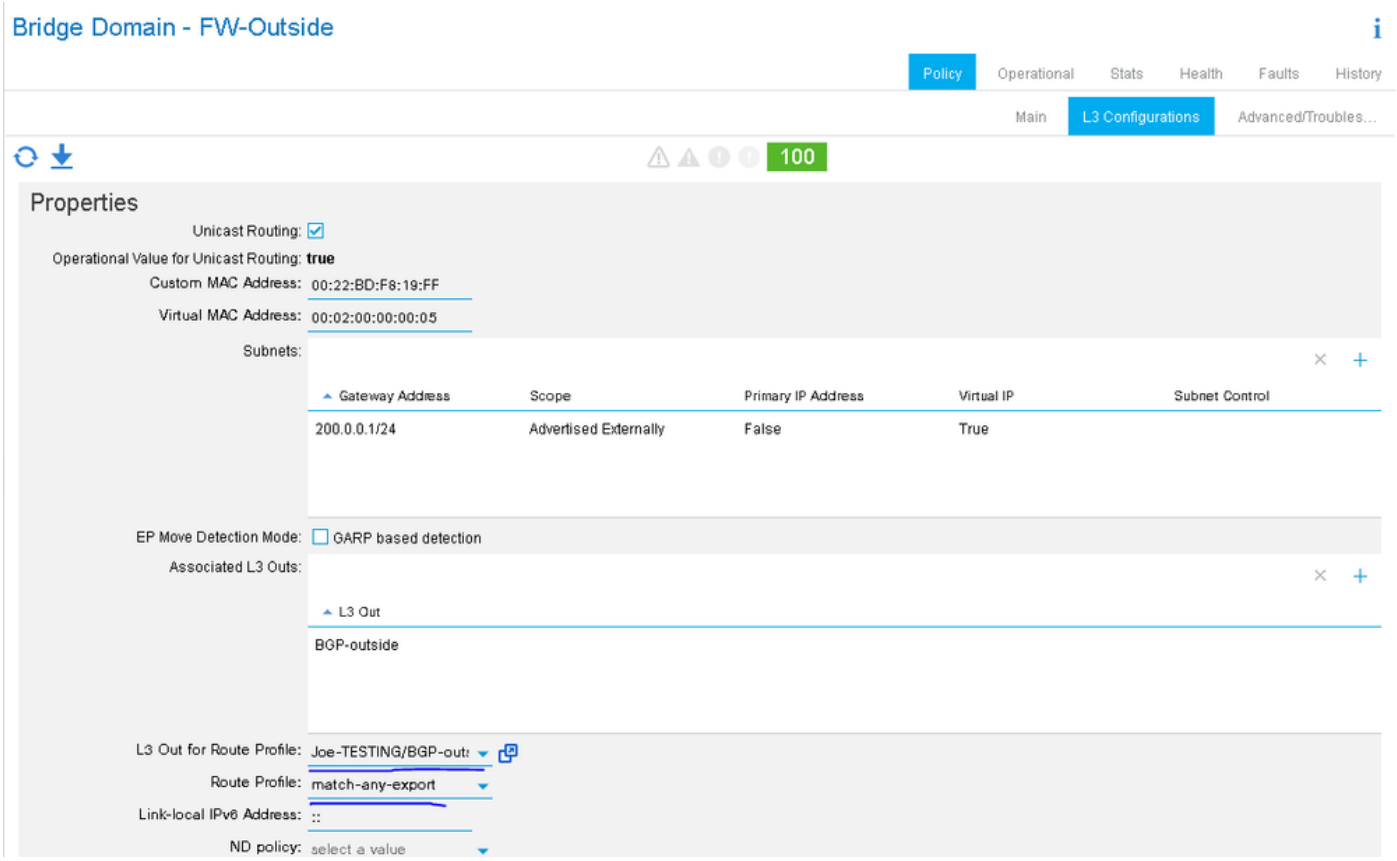

In this example the BD subnet is 200.0.0.0/24 and the route-profile has one match rule that matches 210.0.0.0/24 and sets the community to 200:200. Because the route-profile is set to combinable "Match Prefix AND Routing Policy" the rule will explicitly match 210.0.0.0/24 and implicitly match 200.0.0.0/24 (BD Subnet).

Depending on the external protocol that is being used the route-profile will be applied as an outbound route-map to the neighbor (BGP) or at the protocol level when redistributing the static BD subnet into the external protocol (OSPF).

To verify this configuration when BGP is the l3out protocol...

-Find the neighbor address:

```
leaf6# show bgp ipv4 unicast summary vrf Joe-TESTING:Joe-VRF
BGP summary information for VRF Joe-TESTING:Joe-VRF, address family IPv4 Unicast
BGP router identifier 106.106.106.106, local AS number 100
BGP table version is 97, IPv4 Unicast config peers 1, capable peers 1
7 network entries and 7 paths using 1204 bytes of memory
BGP attribute entries [4/576], BGP AS path entries [1/6]
BGP community entries [0/0], BGP clusterlist entries [6/24]
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
2.2.2.2 4 12345 5833 5924 97 0 0 4d01h 3
```

-Find the outbound route-map used for that neighbor:

```
leaf6# show bgp ipv4 un neighbor 2.2.2.2 vrf Joe-TESTING:Joe-VRF | grep map
Inbound route-map configured is permit-all, handle obtained
Outbound route-map configured is exp-l3out-BGP-outside-peer-3080194, handle obtained
```

-Look at the contents of the route-map:

```
leaf6# show route-map exp-l3out-BGP-outside-peer-3080194
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 4001
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-int-out-match-any-export2any0210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
community 200:200 additive
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 7801
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-int-inferred-export-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8000
Match clauses:
route-type: static
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8001
Match clauses:
route-type: direct
Set clauses:
```

```
leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-int-out-match-any-export2any0210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-int-out-match-any-export2any0210.0.0.0-dst: 2 entries
seq 1 permit 210.0.0.0/24 << Match rule  seq 2 permit 200.0.0.1/24 << Implicit match because route-profile is combinable.
```

In the above example sequence 7801 will match BD subnets so the BD subnet will be matched implicity in both sequence 4001 and 7801. If the route-profile were set to "Match Routing Policy Only" then the match rule would only include 210.0.0.0/24 and not the BD subnet. The BD subnet would still be matched implicitly in a later sequence number so it would be permitted (not sure if this is the same behavior for earlier software releases).

Applying a Route-Profile at the Bridge Domain Subnet Level

The route-profile can directly be associated to the BD Subnet. One of the only use-cases for doing this would be when there is more than one subnet configured under the BD and policy should be applied to these as they are advertised out more than one l3out. (currently only one l3out for route-profile can be associated at BD level)

The configuration can be seen below:

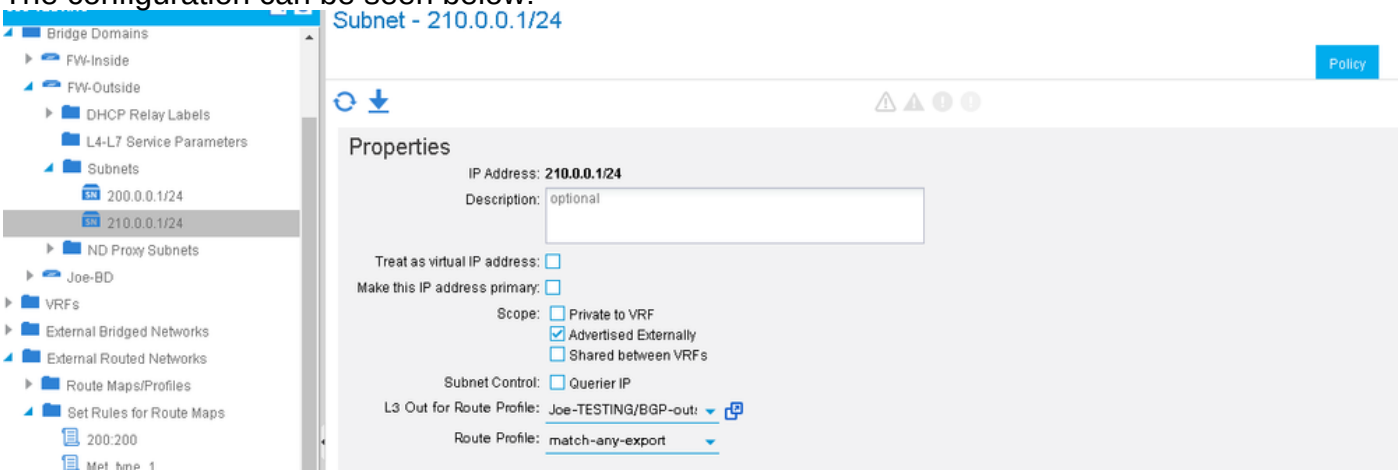

The only difference between applying the route-profile at the BD level vs. the BD subnet level is

that when "Match Prefix AND Routing Policy" is selected only the associated BD subnet will be implicitly included in each match rule. So if there were more than one BD subnet in the same BD only the subnet that the route-profile is tied to would implicitly be matched. This can be verified in the same way as applying the route-profile at the BD level. This example will use OSPF.

A BD is configured with **200.0.0.0/24** and **210.0.0.0/24** subnets. A route-profile is configured under the OSPF l3out and associated to the 210.0.0.0/24 BD subnet. The route-profile is set to '**combinable**' so it should match 210.0.0.0/24 (explicit match), 210.0.0.1/24 (implicit match), and not 200.0.0.0/24 (other bd subnet). 200.0.0.0/24 will implicity be matched at the end of the route-profile and permitted. The route-map will set the ospf metric-type to 1.

-Get the route-map that is used for static to ospf redistribution:

```
leaf6# show ip ospf vrf Joe-TESTING:Joe-VRF | grep -A 4 Redistributing
Redistributing External Routes from
static route-map exp-ctx-st-3080194
direct route-map exp-ctx-st-3080194
bgp route-map exp-ctx-proto-3080194
eigrp route-map exp-ctx-proto-3080194
leaf6# show route-map exp-ctx-st-3080194
route-map exp-ctx-st-3080194, permit, sequence 2001
Match clauses:
ip address prefix-lists: IPv4-st10934-3080194-exc-int-out-non-default-export100210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
metric-type type-1
route-map exp-ctx-st-3080194, permit, sequence 7801
Match clauses:
ip address prefix-lists: IPv4-st10934-3080194-exc-int-inferred-export-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:

leaf6# show ip prefix-list IPv4-st10934-3080194-exc-int-out-non-default-export100210.0.0.0-dst
show ip pip prefix-list IPv4-st10934-3080194-exc-int-out-non-default-export100210.0.0.0-dst: 2
entries
seq 1 permit 210.0.0.1/24
seq 2 permit 210.0.0.0/24

leaf6# show ip prefix-list IPv4-st10934-3080194-exc-int-inferred-export-dst
ip prefix-list IPv4-st10934-3080194-exc-int-inferred-export-dst: 2 entries
seq 1 permit 210.0.0.1/24
seq 2 permit 200.0.0.1/24
```

***Due to CSCvd68302 if a route-profile is associated at the BD subnet level and then removed the route-map may not be removed. The workaround is to make some change in the route-profile (ex: toggle a set rule) to trigger a cleanup. This will be fixed in a future SW release.

Applying a Route-Profile at the 'Default' Level

There are two different default route-profiles that can be configured at the l3out level. These are the 'default-import' and 'default-export' route-profiles. These do not have to be applied anywhere. As long as they exist they will affect matched routes that are being advertised out that l3out. The configuration is identical to any other route-profile creation except that the name must be specified as 'default-export' or 'default-import'. If the software version is late enough then these two names

will appear as options in a drop-down.

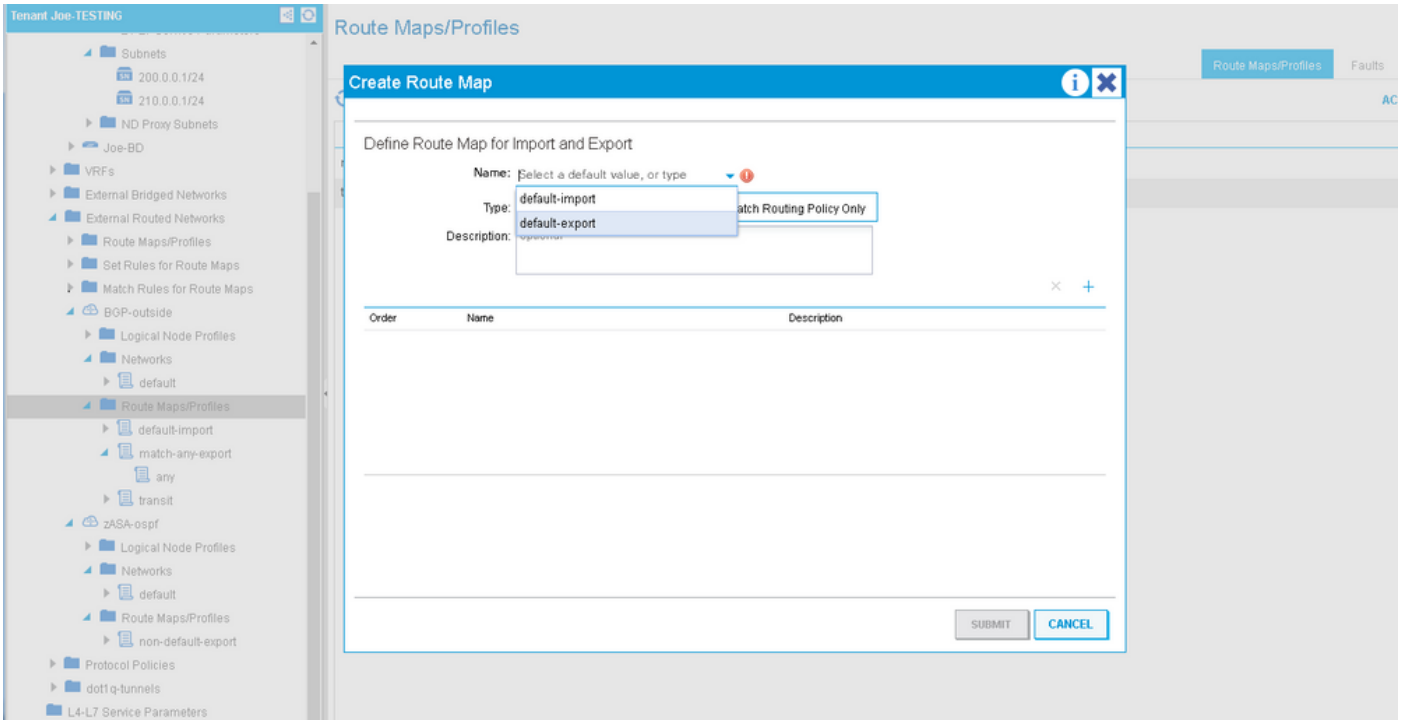

The default-export route-map creates match entries that apply to two different types of routes:

1. External routes that are being advertised out (transit prefixes). The associated route-map entry will match whatever is matched in the default-export match rule(s), perform the set rule specified in the context, and implicitly set the route-tag to the vrf tag. The implicit tag set is done any time transit routing is done in ACI. The border leafs will never install a route in the routing table that has this tag set so setting it on transit prefixes ensures that the prefixes are never looped back into ACI and installed in the routing table in the same VRF.

2. Internal routes that are being advertised out (BD prefixes). This associated route-map entry will match whatever is matched in the default-export match rule(s) and perform the associated set action. If the route-profile is set to 'combinable' (Match Prefix AND Routing Policy) then this entry(ies) in the route-map will include all BD subnets implicitly. If it is not set to combinable it will only match whatever is matched in the match rule.

*******IMPORTANT, setting the default-export to 'Match Routing Policy Only' (non-combinable) will cause BD subnets to stop being advertised if they aren't explicitly matched in the route-profile.***

In the following example the BD subnets are 200.0.0.0/24 and 210.0.0.0/24. The route-profile has one context which matches 210.0.0.0/24 and sets the community to 200:200. The default-export is applied and set to non-combinable.

```
leaf6# show route-map exp-l3out-BGP-outside-peer-3080194
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 4001
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-ext-out-default-export200210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
tag 4294967295
community 200:200 additive
```

```
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 4002
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-int-out-default-export200210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
community 200:200 additive
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8000
Match clauses:
route-type: static
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8001
Match clauses:
route-type: direct
Set clauses:

leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-ext-out-default-export200210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-ext-out-default-export200210.0.0.0-dst: 1 entries
seq 1 permit 210.0.0.0/24

leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-int-out-default-export200210.0.0.0-dst ip
prefix-list IPv4-peer10932-3080194-exc-int-out-default-export200210.0.0.0-dst: 1 entries seq 1
permit 210.0.0.0/24
```

The route-map entry with prefix list "ext-out" is for transit prefixes. It only matches what is matched in the match rule and sets the tag to the vrf default tag. The second route-map entry with prefix-list "int-out" is for internal prefixes (BD subnets) being advertised out. Since the route-profile is not set to cominable it only matches 210.0.0.0/24 since that is what the match rule specified. The other BD subnet 200.0.0.0/24 is not matched and traffic to this subnet could be blackholed.

After changing the route-profile to combinable:

```
leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-ext-out-default-export200210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-ext-out-default-export200210.0.0.0-dst: 1 entries
seq 1 permit 210.0.0.0/24

leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-int-out-default-export200210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-int-out-default-export200210.0.0.0-dst: 3 entries
seq 1 permit 210.0.0.0/24 seq 2 permit 210.0.0.1/24 seq 3 permit 200.0.0.1/24
```

The route-map entry for transit prefixes remains the same but the entry for internal prefixes now includes all BD prefixes as well as what is specified in the match rule.

Applying a Route-Profile at the External EPG and External EPG Subnet levels

A Route-Profile can also be applied directly to an external epg level or the subnet level within an external epg. This is intended for applying policy to transit prefixes but can also be used to apply policy to internal prefixes. The only caveat being that the internal prefixes (if matched) will receive the default vrf tag. If those subnets are supposed to be advertised back into ACI in a different VRF then make sure to change the default tag for that vrf so that the prefixes are accepted and installed in the routing-table.

If the route-profile is set to 'non-combinable' then there is no difference between applying the route-profile at the Ext EPG level vs. the Ext EPG subnet level. The route-map entries will only match what is explicitly matched in the match rule. If the route-profile is set to combinable and the route-profile is applied at the Ext EPG level then each match entry will match what is explicitly

specified and any subnets that are defined as 'export route-control subnet'. If the route-profile is set to combinable and applied at the Ext EPG subnet level then the route-profile will match what is explicitly specified and implicitly match the EPG subnet it is applied to IF that subnet is set to "export route-control subnet".

In this example the BD subnets are 200.0.0.0/24 and 210.0.0.0/24. 89.89.89.89/32 and 90.90.90.90/32 are specified as L3out networks with "export route control subnet" set. The route-map profile has a context that matches 210.0.0.0/24 and sets the community to 200:200. The route-profile is applied at the Ext EPG level and is non-combinable.

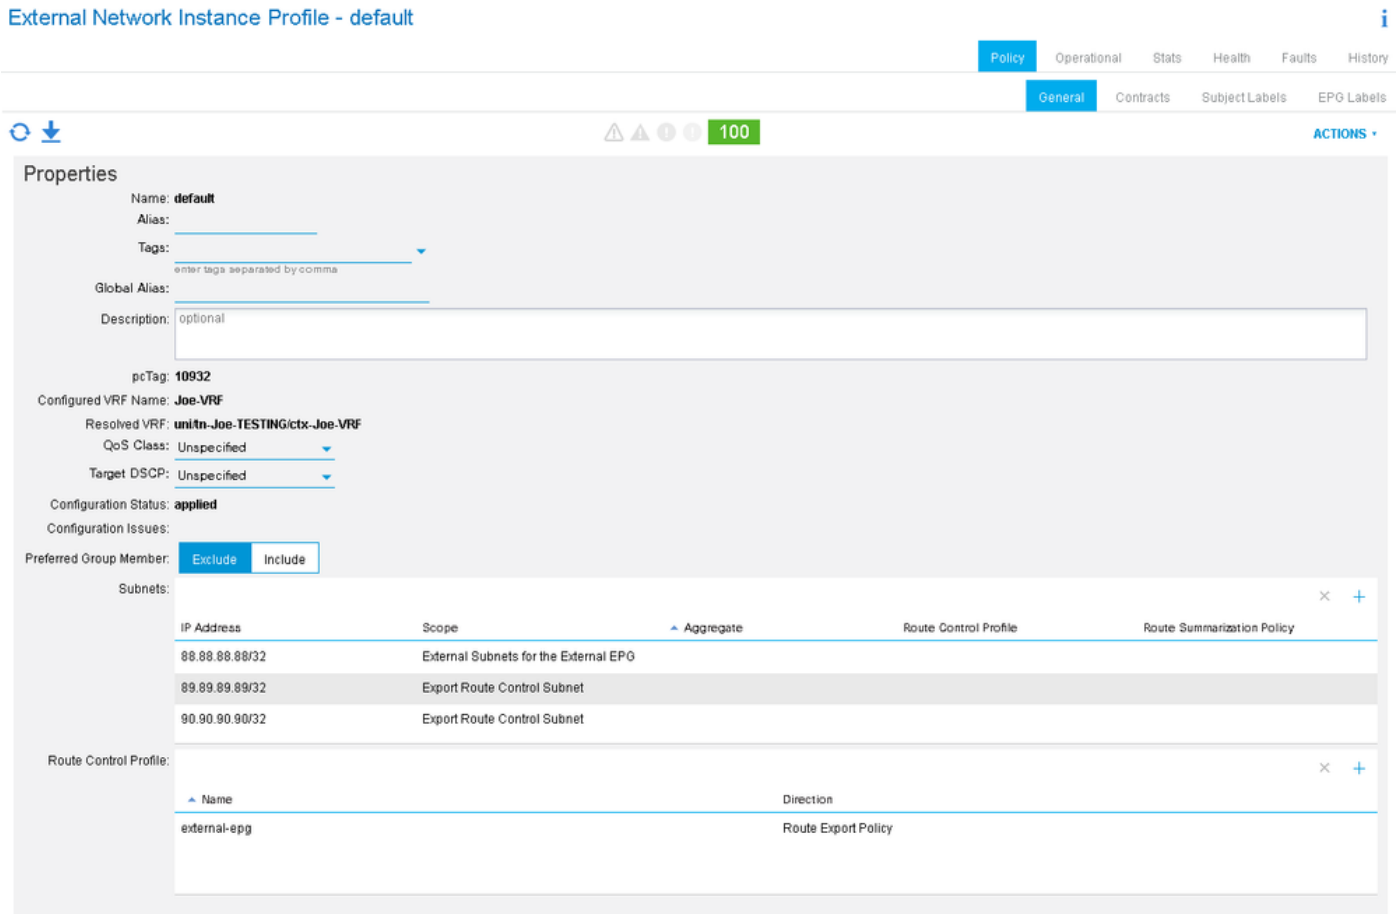


```
leaf6# show bgp ipv4 un neighbors 2.2.2.2 vrf Joe-TESTING:Joe-VRF | grep map
Inbound route-map configured is permit-all, handle obtained
Outbound route-map configured is exp-l3out-BGP-outside-peer-3080194, handle obtained

leaf6# show route-map exp-l3out-BGP-outside-peer-3080194
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 4001
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-ext-out-external-epg200210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
tag 4294967295
community 200:200 additive
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 7801
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-int-inferred-export-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8000
Match clauses:
route-type: static
Set clauses:
```

```
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8001
Match clauses:
route-type: direct
Set clauses:

leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-ext-out-external-epg200210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-ext-out-external-epg200210.0.0.0-dst: 1 entries
seq 1 permit 210.0.0.0/24
leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-int-inferred-export-dst
ip prefix-list IPv4-peer10932-3080194-exc-int-inferred-export-dst: 2 entries
seq 1 permit 210.0.0.1/24
seq 2 permit 200.0.0.1/24
```

Notice that the route-map entry only matches what is specified in the match rule even though subnets are defined with "export route-control subnet". There is still an entry in the route-map that permits all BD subnets that are set to "advertise externally" and are associated to this L3out.

If the route-profile is changed to combinable:

```
leaf6# show route-map exp-l3out-BGP-outside-peer-3080194
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 4001
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-ext-out-external-epg200210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
tag 4294967295
community 200:200 additive
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 7801
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-int-inferred-export-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8000
Match clauses:
route-type: static
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8001
Match clauses:
route-type: direct
Set clauses:
leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-ext-out-external-epg200210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-ext-out-external-epg200210.0.0.0-dst: 3 entries
seq 1 permit 210.0.0.0/24
seq 2 permit 89.89.89.89/32
seq 3 permit 90.90.90.90/32

leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-int-inferred-export-dst
ip prefix-list IPv4-peer10932-3080194-exc-int-inferred-export-dst: 2 entries
seq 1 permit 210.0.0.1/24
seq 2 permit 200.0.0.1/24
```

Notice now that the entry that applies the policy matches all subnets that are set to "export route-control subet".

If the route-profile is combinable and applied directly to one of the subnets that is set to "export route-control subnet":

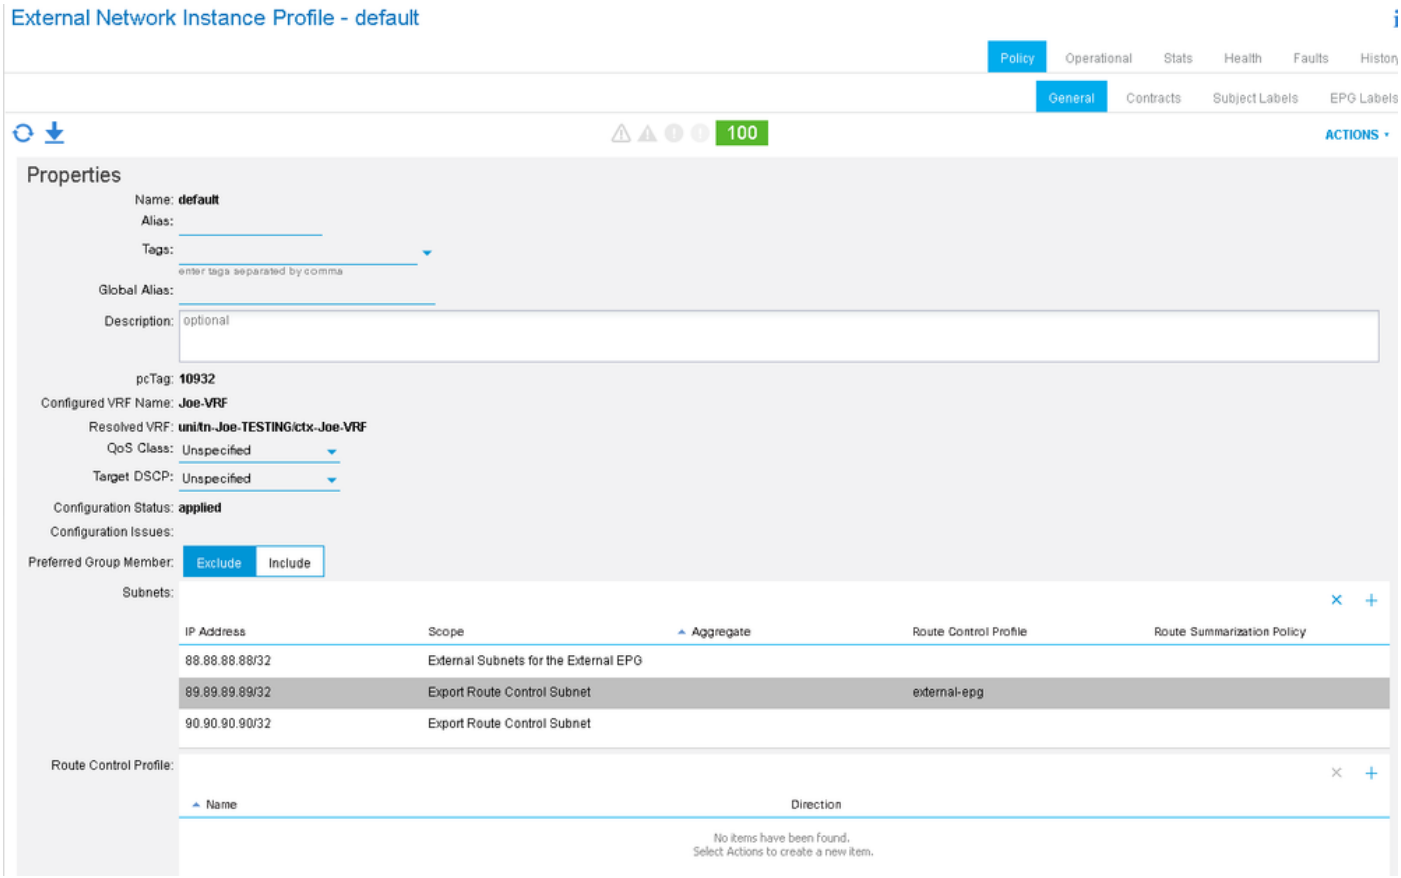


```
leaf6# show route-map exp-l3out-BGP-outside-peer-3080194
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 2001
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-ext-out-external-epg100210.0.0.0-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
tag 4294967295
community 200:200 additive
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 7801
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-int-inferred-export-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, permit, sequence 7802
Match clauses:
ip address prefix-lists: IPv4-peer10932-3080194-exc-ext-inferred-export-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
tag 4294967295
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8000
Match clauses:
route-type: static
Set clauses:
route-map exp-l3out-BGP-outside-peer-3080194, deny, sequence 8001
Match clauses:
route-type: direct
Set clauses:
leaf6# show ip prefix-list IPv4-peer10932-3080194-exc-ext-out-external-epg100210.0.0.0-dst
ip prefix-list IPv4-peer10932-3080194-exc-ext-out-external-epg100210.0.0.0-dst: 2 entries
seq 1 permit 210.0.0.0/24
seq 2 permit 89.89.89.89/32
```

Notice that the route-map entry that applies the policy includes what is matched in the route-profile context and the subnet that it is applied to since "export route-control subnet" is selected. The other subnet that has "export route-control subnet" is not included in the route-map entry that

applies the policy though it is matched in an implicit rule that simply permits it and sets the transit tag.

Applying a Route-Profile at the L3out Level as an Interleak Policy:

The "Route Profile for Interleak" is intended specifically to set policy when redisributing prefixes from some external protocol into BGP. This is the only case where the route-profile should be configured under "External Routed Networks" rather than under the l3out. The route-profile is then applied on the source External protocol (non-bgp) as a "Route Profile for Interleak" policy. This is useful for setting BGP attributes when a prefix is redistributed into the internal fabric bgp process or it can also be used to set bgp attributes when advertising transit prefixes from a non-bgp l3out to a bgp l3out.

In this example 89.89.89.89/32 is being received from OSPF. An interleak route-profile is being applied to the OSPF l3out that matches 89.89.89.89/32 and sets the BGP community to 200:200. The policy is applied as the OSPF route is redistributed into BGP. To verify this you would look at the route-map that gets set in the BGP process.

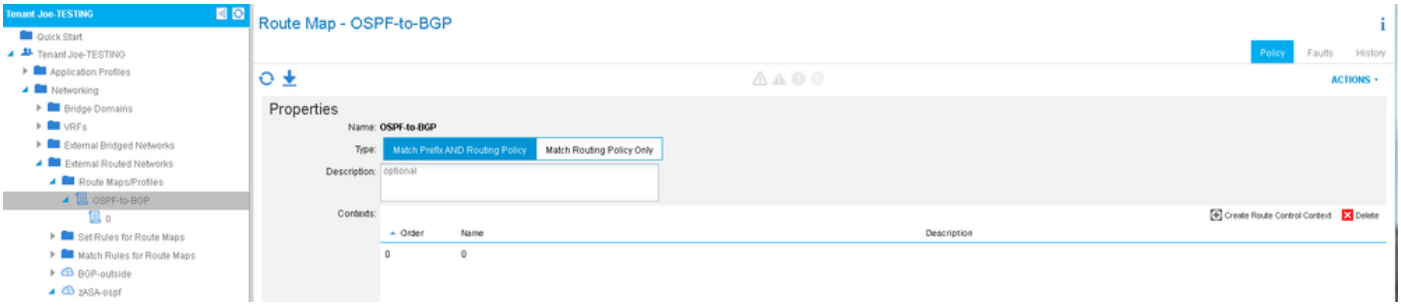

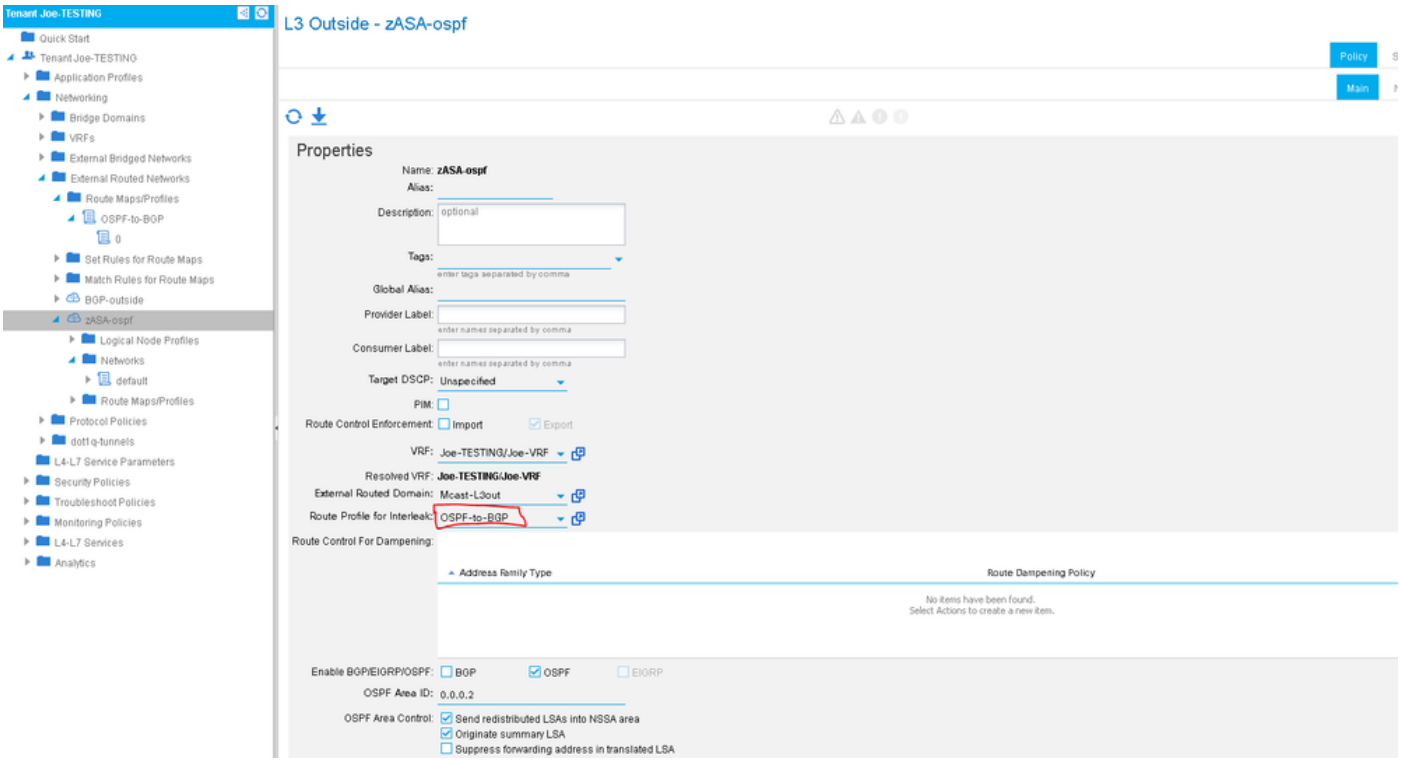

Use "**show bgp process**" to verify the route-map that is being used for redistribution from OSPF to BGP.

```
leaf6# show bgp process vrf Joe-TESTING:Joe-VRF | grep -A 4 Redistri
Redistribution
direct, route-map permit-all
static, route-map imp-ctx-bgp-st-interleak-3080194
ospf, route-map imp-ctx-proto-interleak-3080194
route-map imp-ctx-proto-interleak-3080194, permit, sequence 1
Match clauses:
ip address prefix-lists: IPv4-st10934-3080194-ext-in-OSPF-to-BGP00089.89.89.89-dst
ipv6 address prefix-lists: IPv6-deny-all
Set clauses:
community 200:200 additive

leaf6# show ip prefix-list IPv4-st10934-3080194-ext-in-OSPF-to-BGP00089.89.89.89-dst
ip prefix-list IPv4-st10934-3080194-ext-in-OSPF-to-BGP00089.89.89.89-dst: 1 entries
seq 1 permit 89.89.89.89/32
```

Note that the OSPF epg also includes the "0.0.0.0" subnet but the only thing being redistributed into BGP from OSPF is 89.89.89.89. Setting the route-profile to "combinable" vs "non-combinable" has no affect on interleak policies.

Its important to know that nothing is implicitly allowed into BGP when an interleak policy is set. If there is no interleak policy set (default) then everything is permitted; if a route-profile for interleak is set then nothing is allowed except what is explicitly matched. Misunderstanding of this could lead to outages when configuring interleak policies.

Deny Rules

The ability to deny specific prefixes was added in 2.3(1) software. Previously only permit rules could be matched so there was no ability to deny specific prefixes using route-profiles. The deny action is set in the route-profile context:

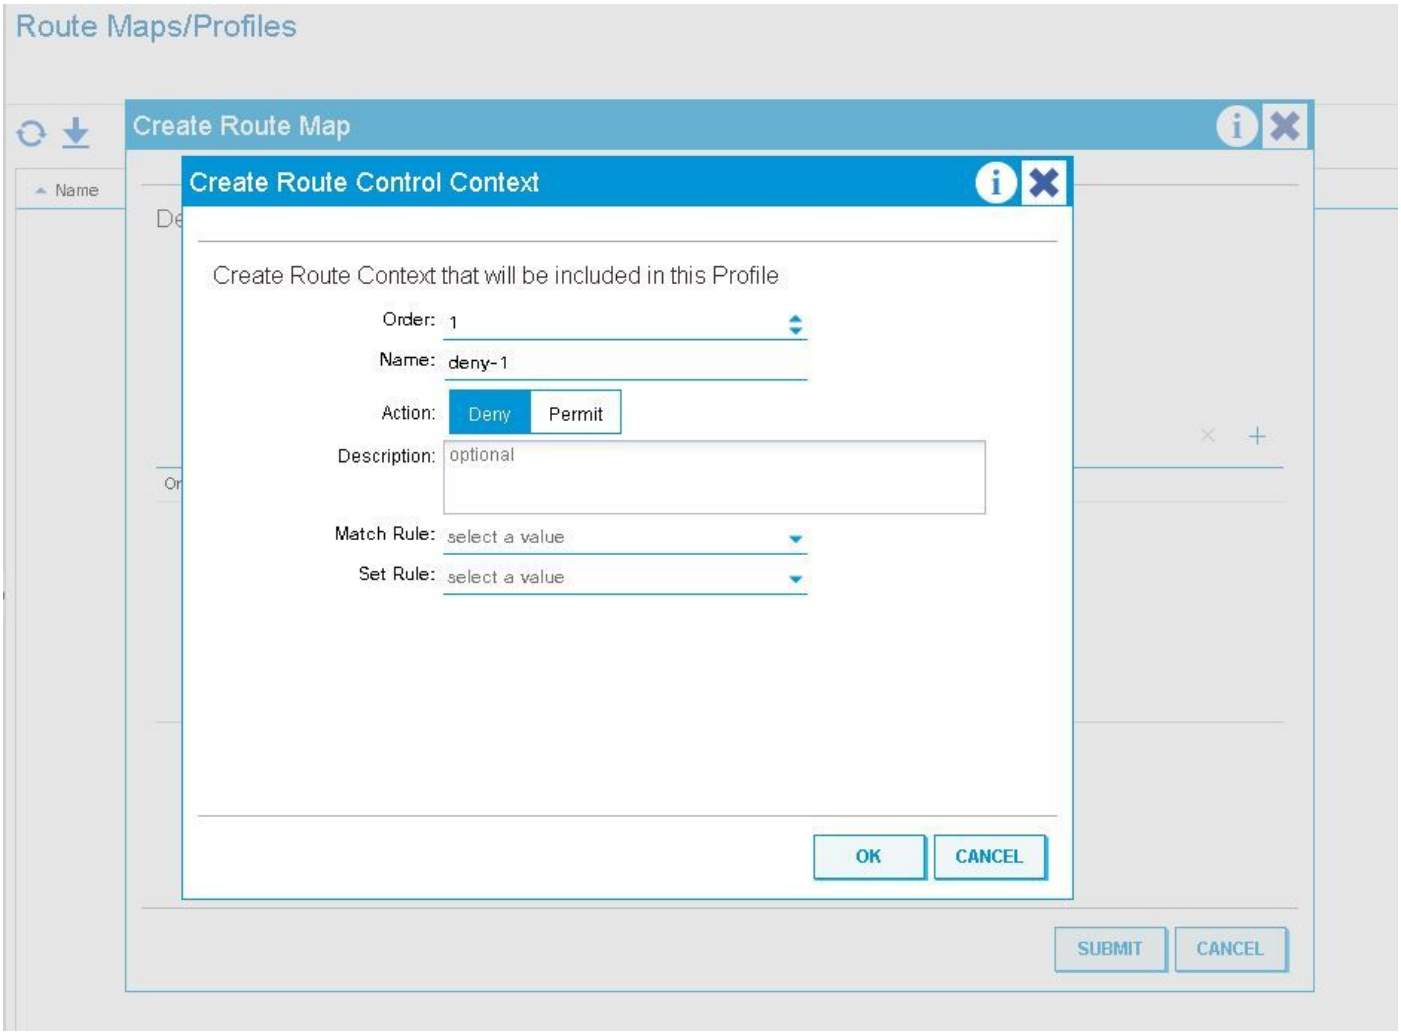

Special care should be used when using deny rules with a Route-Profile that is set to 'combinable' (Match Prefix AND Routing Policy).

The following lists the behavior of deny rules when the route-profile is set to combinable v. non-combinable

Deny Rule behavior with route-profile applied at the Bridge Domain subnet level

Combinable - Deny rules will match whatever is specified in the match rule as well as the BD subnet that the route-profile is applied to.
Non-combinable - Deny rules will only match what is specified in the match rule.

Deny Rule behavior with route-profile applied at the Bridge Domain level

Combinable - Deny rules will match whatever is specified in the match rule as well as all subnets that are configured within that BD.
Non-combinable - Deny rules will match only what is specified in the match rule.

Deny Rule behavior with route-profile applied at the Default-Export level

Combinable - Deny rules will implicitly match ALL BD subnets that are set to be advertised externally as well as what is matched in the rule
Non-Combinable - Deny rules will only match what is specified in the match rule.

Deny Rule behavior with export route-profile applied at the L3out Network Instance level

Combinable - Deny rules will implicitly match all networks with "export route control subnet" set as well as what is matched in the match rule.
Non-combinable - Deny rules will only match what is match in the match rule.

Deny Rule behavior with export route-profile applied at the L3out Network Subnet level

Combinable - If the network that the export route-profile is applied to has "export route control subnet" selected it will be matched as well as what is match in the match rule.
Non-combinable - Deny rules will only match what is match in the match rule.

Deny Rule behavior with export route-profile applied at the "Route Profile for Interleak" level

-Deny rules are not intended for usage here. Regardless of whether 'deny' is set, the resolved route-map on the leaf will have a match rule. Denying prefixes inbound should be done with import security or route-filtering on the external device.

Other Notes

The RPM process is used internally for configuring route-map's from route-profiles. Most useful commands to see RPM information can be seen with "show system internal rpm ...". One way to verify that a route-map is actually getting applied, removed, or changed when a configuration is changed is by looking at the RPM event-history on the Leaf swtich:

```
show system internal rpm event-history events
```