
cisco Live!
Let's go

The bridge to possible

Advanced SD-WAN Policies Troubleshooting

And well-known issues with centralized policies

Eugene Khabarov, SD-WAN Escalation Engineer, BU

cisco Live!

BRKENT-3797

Baseline and Objectives

- Cisco SD-WAN at least basic level knowledge is a must
- This is advanced level session, very technical
- The session main objectives:
  - demonstrate useful policies troubleshooting tools
  - share experience about some typical failures seen in the field to help you avoid them in your network
- Not a complete guide, there are always more issues...
- Consider this session as a ”cookbook” for SD-WAN policies failures, but not a “Tour de Force”
- The session is mainly oriented on centralized policies, but we will briefly discuss localized policies as well
- Main topics touched:
  - Policies troubleshooting workflow
  - Internal components of IOS-XE responsible for policies programming and execution
  - Troubleshooting toolset
  - Common pitfalls and underwater stones
- Heavily CLI based, old-school classic ☺
- Recommended prerequisite session: Advanced SD-WAN Routing Troubleshooting (BRKENT-3793)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- Part 1: SD-WAN Polices Troubleshooting Basics
  - 1.1 SD-WAN Policies Quick Overview
  - 1.2 Troubleshooting SD-WAN policies from vManage perspective
  - 1.3 Centralized Control Policies troubleshooting workflow
  - 1.4 Centralized Data and AAR Policies troubleshooting workflow
- Part 2: Issues seen in the field
  - 2.1 Issues with control policies in disjoined underlays
  - 2.2 Not-so-well-known failures with centralized control policies
  - 2.3 Interesting cases with centralized data and AAR policies

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Part 1. SD-WAN Policies Troubleshooting Basics

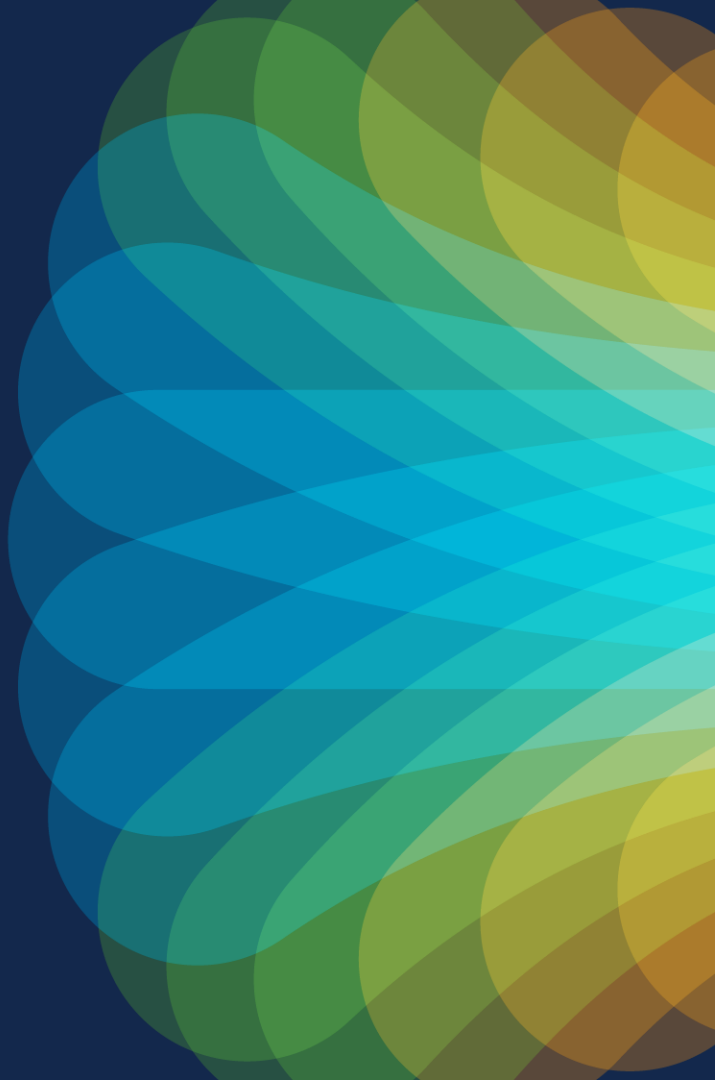

SD-WAN Policies Quick Overview

Cisco Catalyst SD-WAN Overlay

Fabric Components Quick Recap

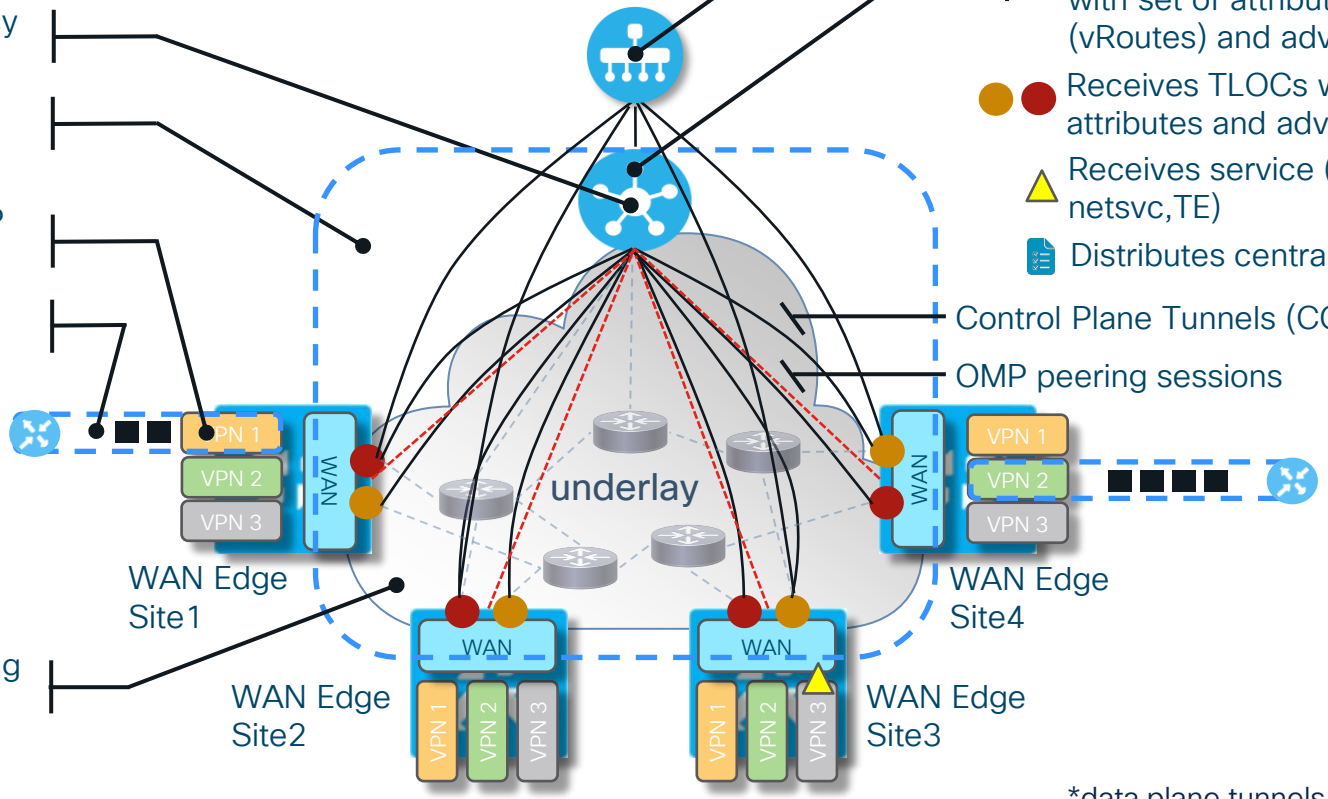

*data plane tunnels are not shown

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

New Cisco Catalyst SD-WAN components naming

- vManage (NMS) == Catalyst SD-WAN Manager
- vBond (orchestrator) == Catalyst SD-WAN Validator
- vSmart == Catalyst SD-WAN Controller

We will stick to legacy names vManage/vBond/vSmart in our slides

Why? Because we like them and to avoid confusion. They are historically called so and in all CLI outputs and all codebase their names will remain the same (vmanage/vbond/vsmart), no plans to change it.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco SD-WAN Policy Architecture

Policy Categories

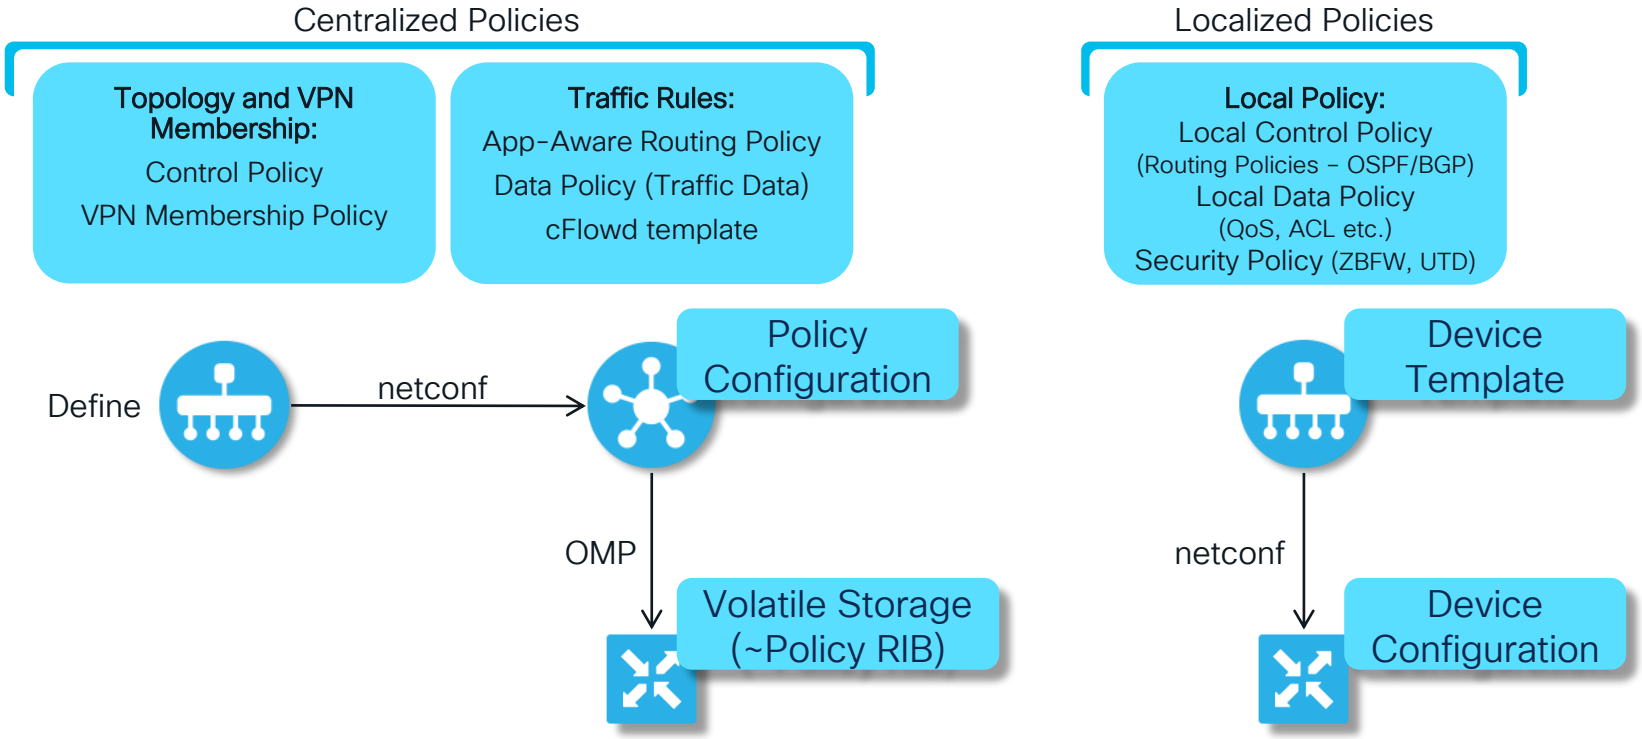

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Building Blocks of Centralized Policies

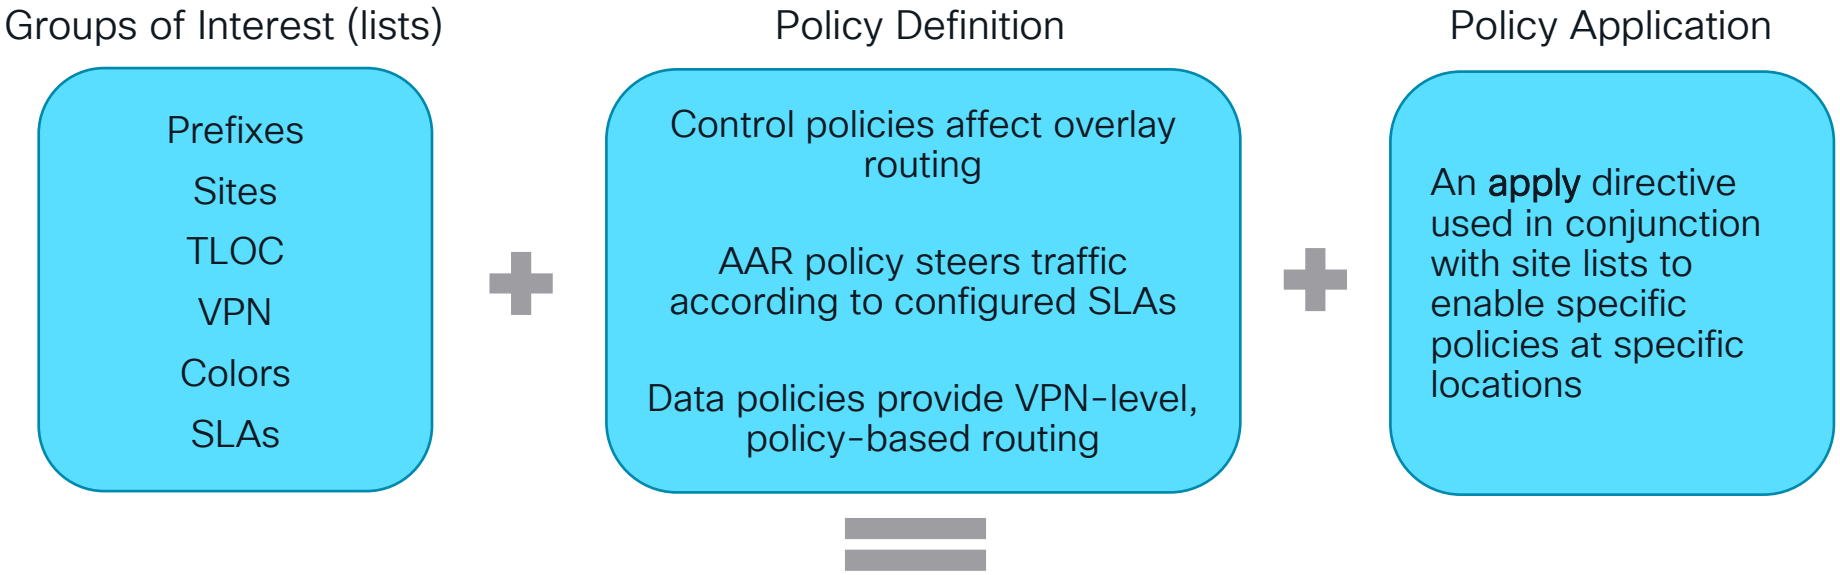

Centralized policy definition is configured on vManage and enforced across the network (on a device or vSmart controller depending on type)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Generalized Policy Troubeshooting Workflow

For Your Reference

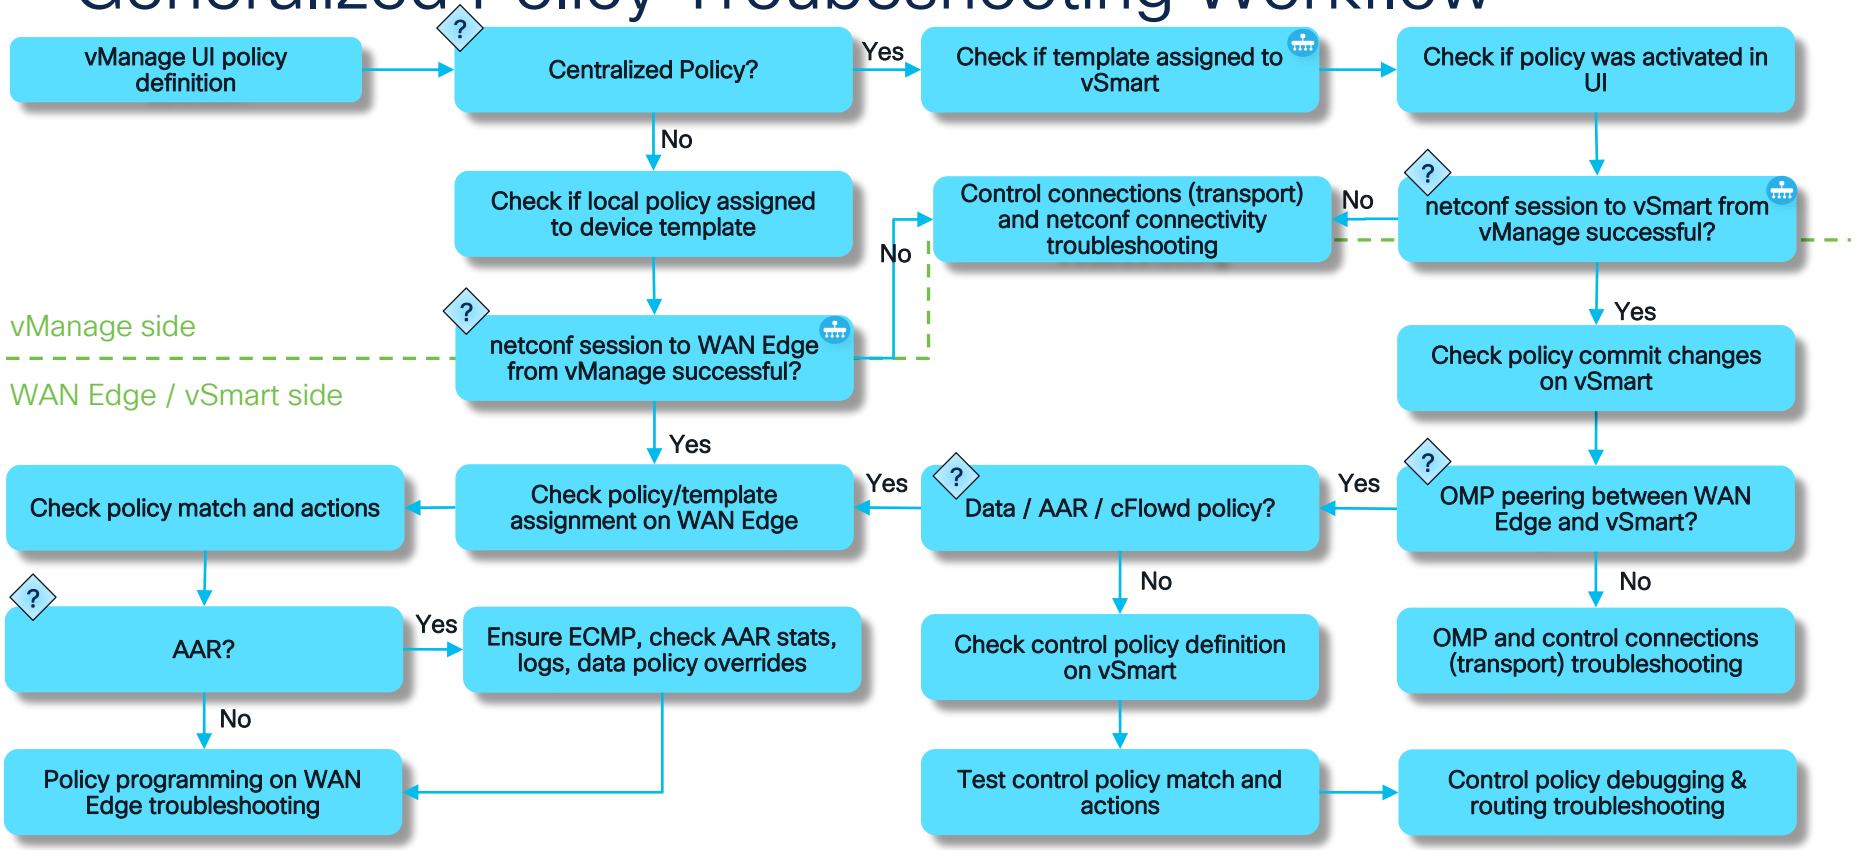

* vManage automatically verify this and warns us in case of problems

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Troubleshooting SD-WAN policies from vManage perspective

Centralized Policy: Check if template assigned to vSmart

Configuration -> Devices -> Controllers

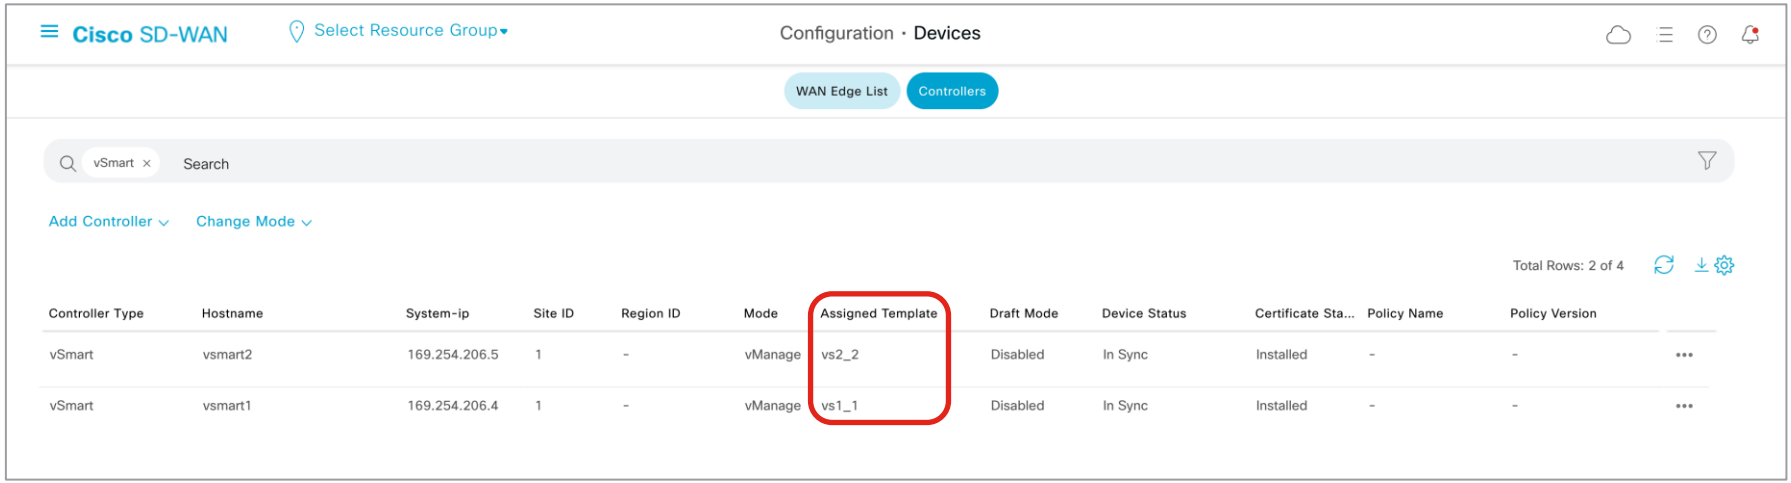

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Policy: Check if policy was activated in UI

Configuration -> Policices -> Centralized Policy

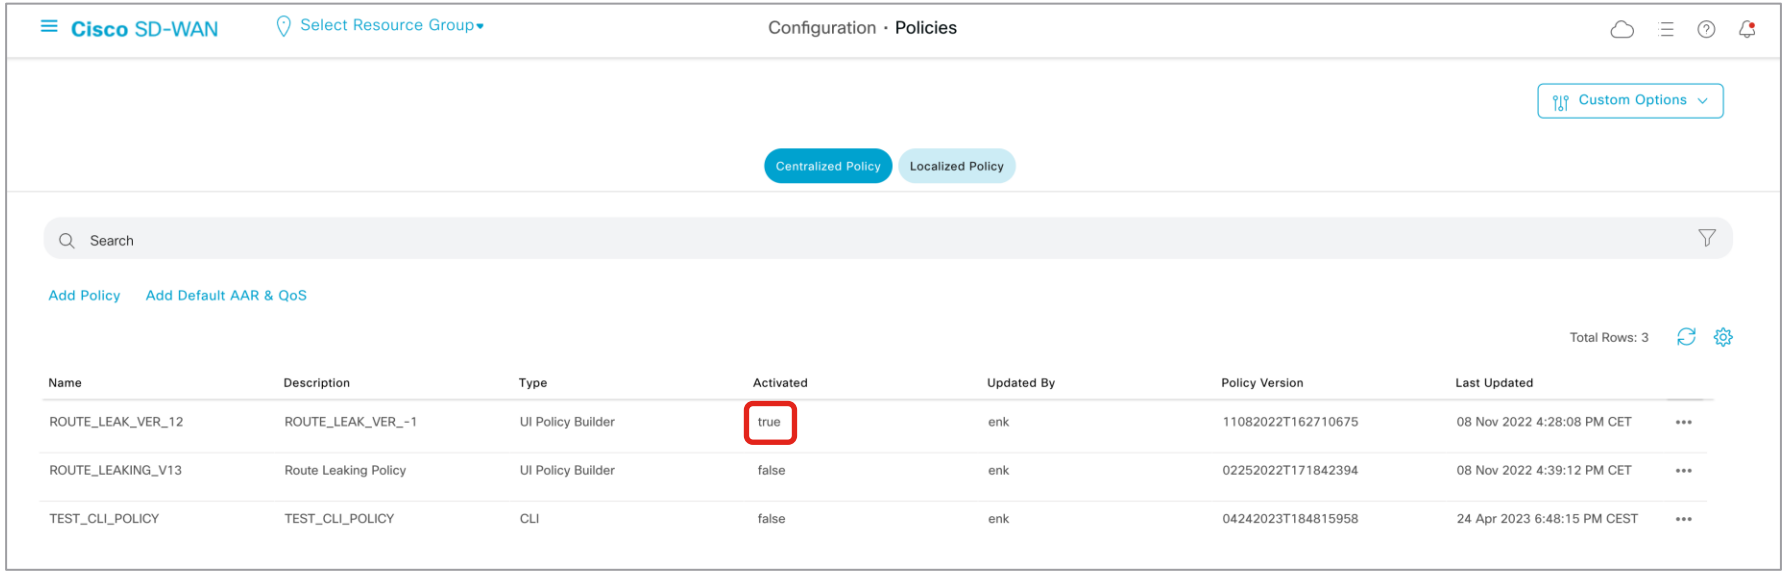

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Policy: Check if template assigned to vSmart

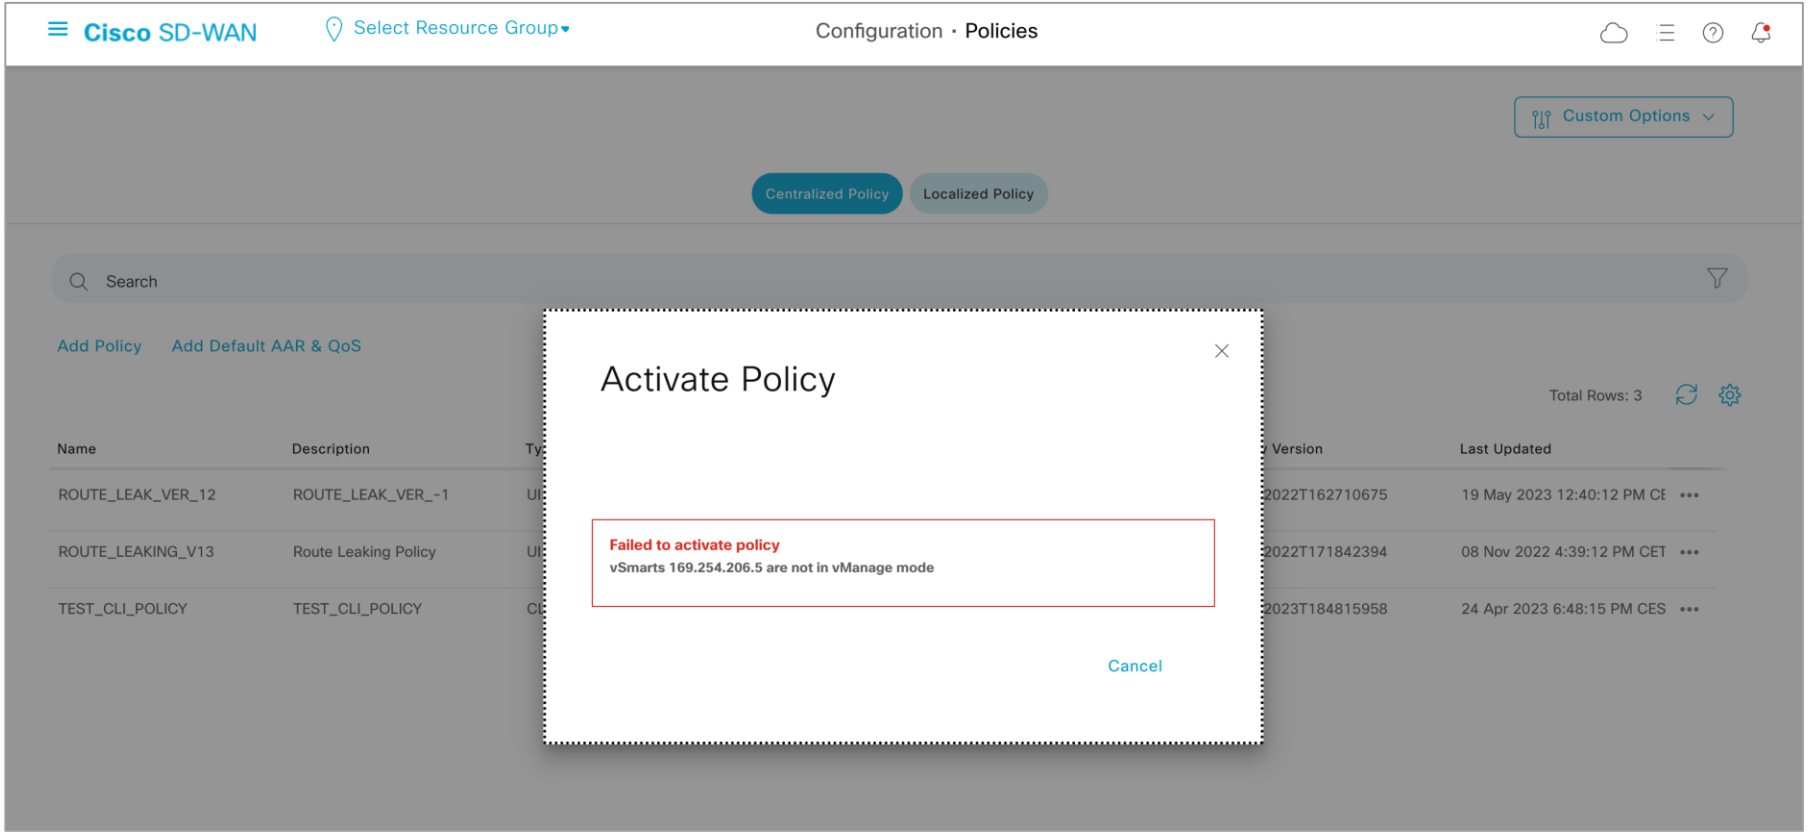

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Policy: Policy Activation Issues

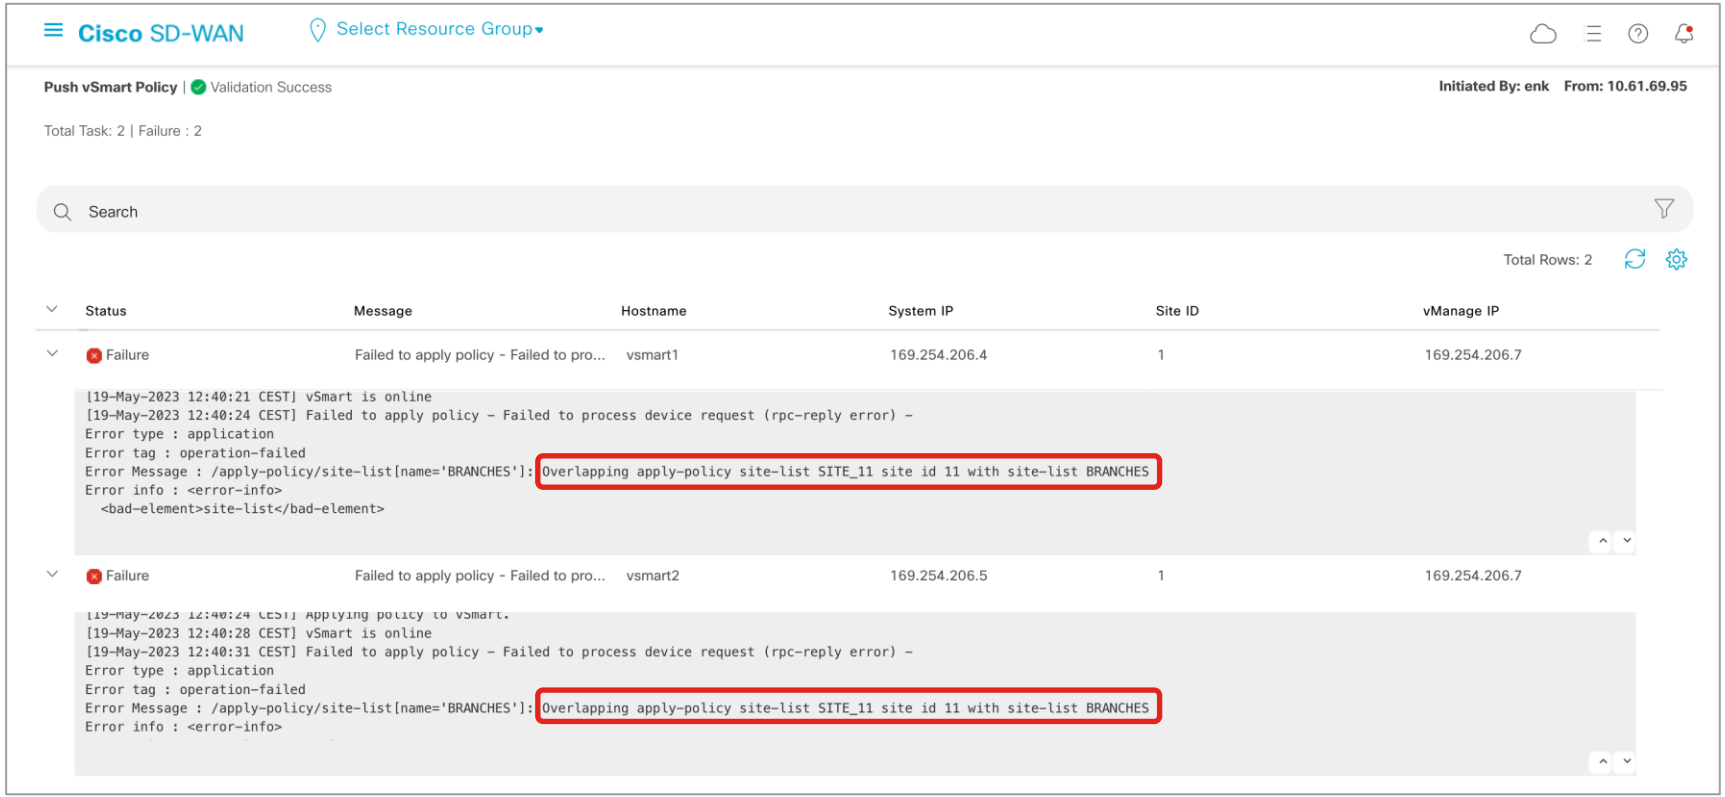

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Localized Policy: Check if policy assigned to device template

Configuration -> Templates -> Device Template -> Additional Templates section

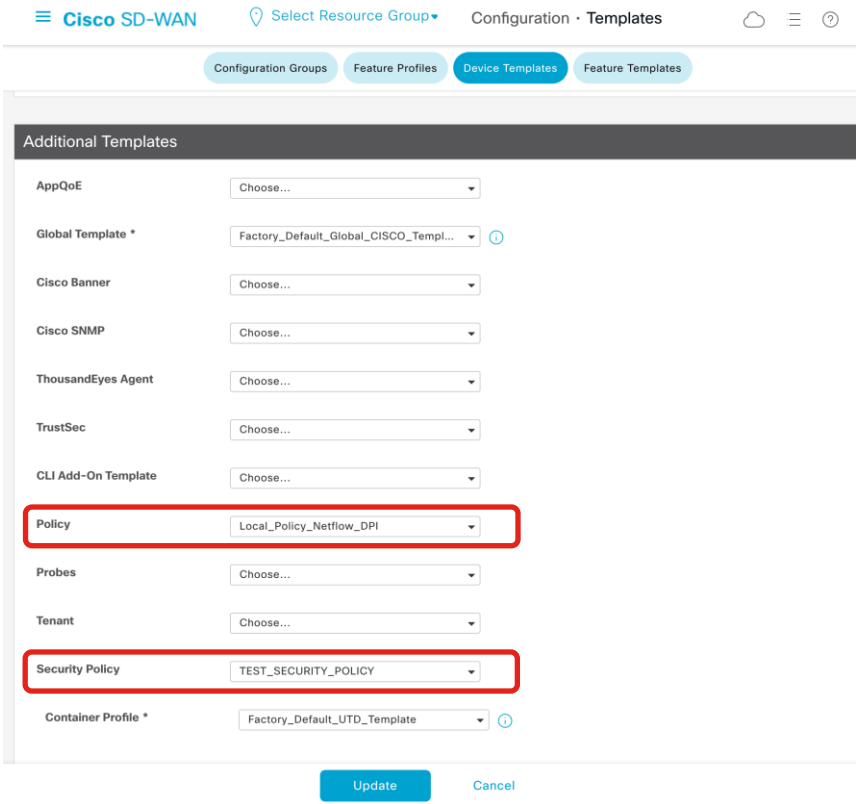

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Localized Policy: Device Template Assignment Issues

Configuration -> Templates -> Device Template -> ... -> Attach Devices

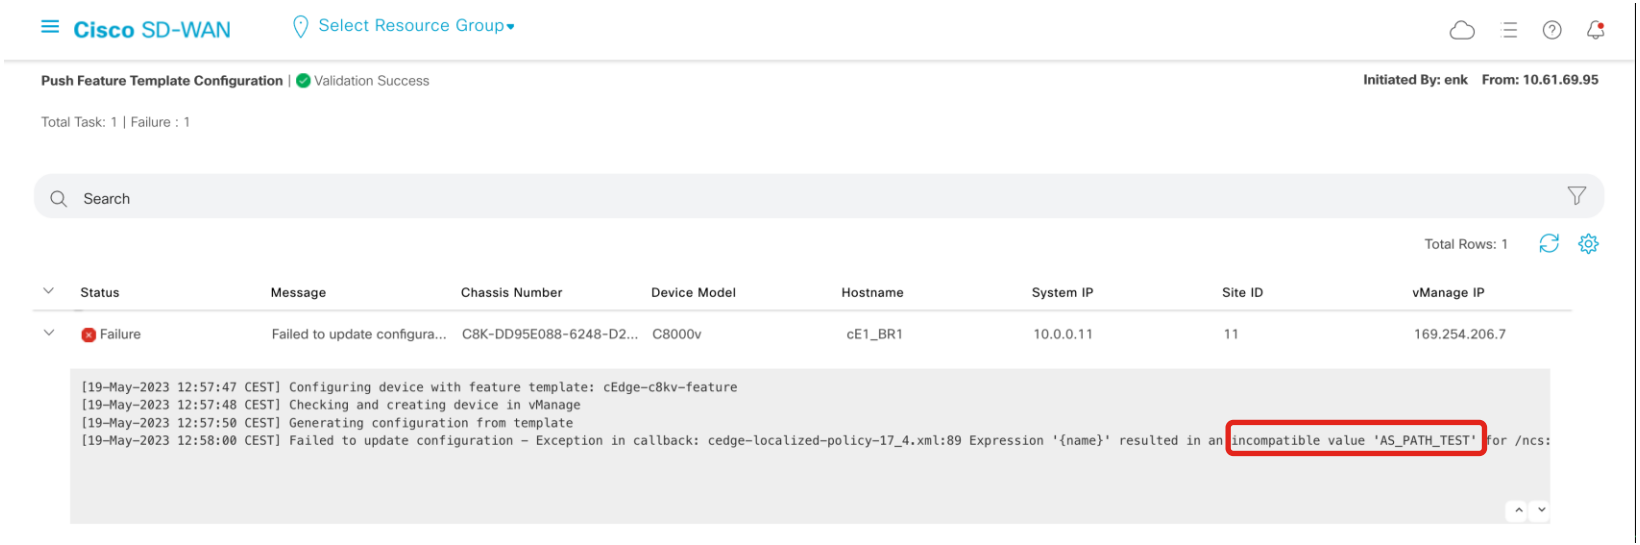

* Here is the reason that AS_PATH_TEST contains typo "^^*$"

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Troubleshooting SD-WAN policies from vSmart and WAN Edge perspective

Centralized Control Policy Troubleshooting 101

cisco Live!

Centralized Control Policy Application

Most important concept to rememer for control policies troubleshooting

Policy applied in the inbound direction

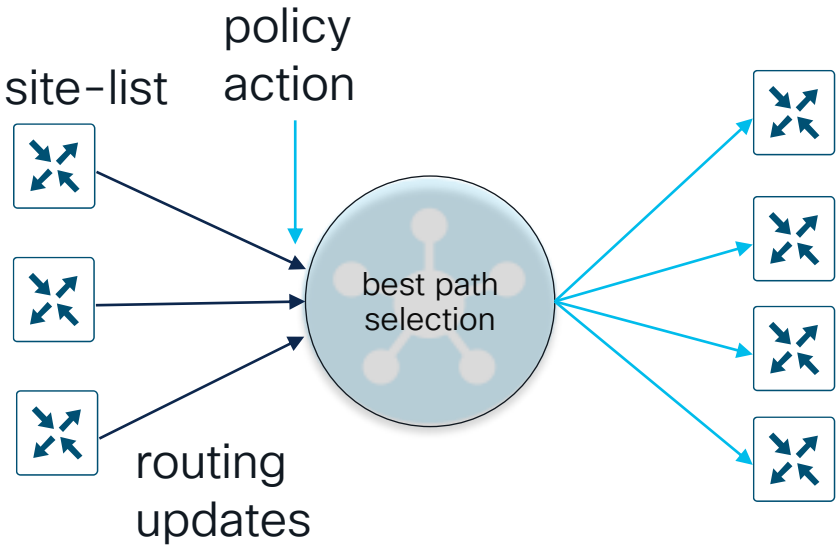

Policy applied in the outbound direction

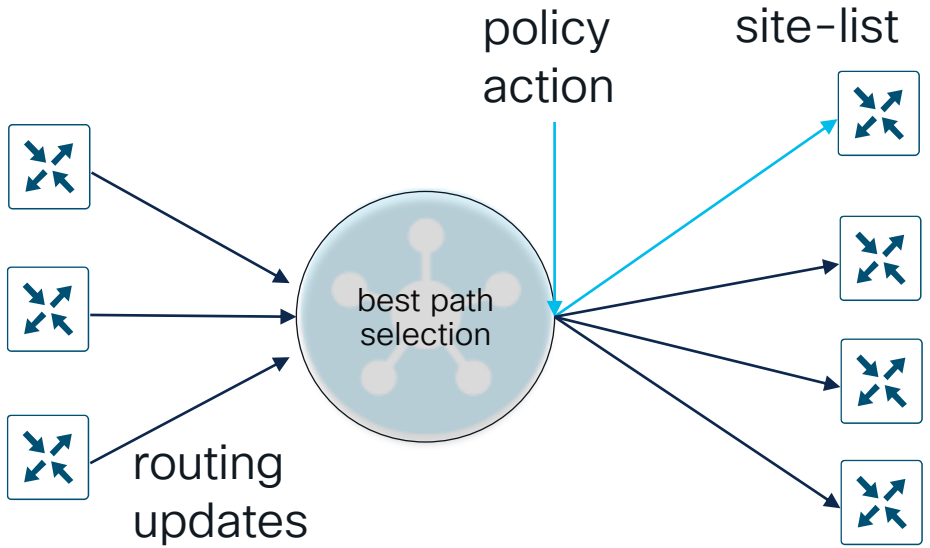

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow (1)

1. Check policy commit changes:

```
show configuration commit changes <number>
```

2. Check OMP peering between WAN Edge and vSmart to ensure policy can be applied on routing updates to/from WAN Edge:

```
show omp peers <system-ip>
```

3. Check control policy assignment and direction

```
show support omp peer peer-ip <system-ip> | include –pol
```

4. Check policy definition (vManage UI polciy definition was sucessfully translated into CLI syntax on vSmart):

```
show configuration commit changes
show run apply-policy site-list <name> control-policy <name>
show run policy list <name>
show run policy control-policy <name>
```

5. Check control policy match and actions

```
test policy match control-policy <name> <conditions>
```

6. Proceed with overlay routing troubleshooting, more in BRKENT-3793

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Recap: overlay routing troubleshooting: missing route(s) algorithm

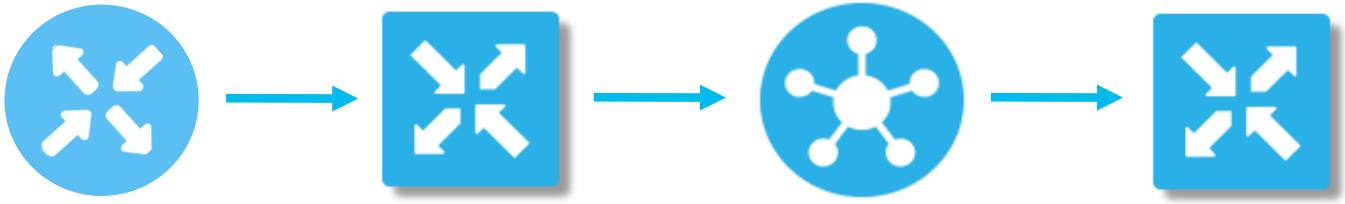

Check on WAN Edge:

1. RIB/FIB (**`show ip route/show [sdwan] ip fib`**)
2. OMP table if route is not in RIB (**`show [sdwan] omp route`**)
3. TLOC information presented (**`show [sdwan] omp tloc`**)
4. BFD session with remote TLOC (**`show [sdwan] bfd sessions`**) -> troubleshoot data plane tunnels
5. Local policy filtering on redistribution to/from OMP table (**`show sdwan run "policy"/show run policy/show run route-map`**)

Check on vSmart:

- OMP route and TLOC tables on vSmart (**`show omp route, show omp tloc`**)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow (2)

- Last resort: start debugging on vSmart:
  - **`debug omp policy [level <high|low> peer-address <system-ip> prefix <IP prefix/length> direction <both|received|sent> vpn <number>]`**
  - Before 20.12 logs stored in **`/var/log/tmplog/vdebug`**
  - 20.12+ logs stored in **`/var/log/vdebug`**
    - Ensure to enable disk logging for debug messages:

      `vSmart1(config)# system logging disk enable priority debug`
  - To view them:
    - enter **`vshell`** and use **`tail -f <filename>`**
    - Or simply **`show log <filename> tail -f`**
    - Or **`monitor start <filename>`** and logs will be printed into your terminal

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow:

Commands usage examples

cisco Live!

Centralized Control Policies Troubleshooting Workflow

1. Check policy commit changes **`show configuration commit changes <number>`**:

```
vsmart1# show configuration commit changes 0
!
! Created by: vmanage-admin
! Date: 2023-04-24 19:22:02
! Client: netconf
!
policy
 lists
  site-list BRANCHES
   site-id 11-12
  !
  site-list SITE-30
   site-id 30
  !
  site-list SITE-40
   site-id 40
  !
  prefix-list DEFAULT
   ip-prefix 0.0.0.0/0
  !
 !
 control-policy MY-CONTROL-POLICY-v1
  sequence 1
   match tloc
    site-list SITE-30
   !
   action accept
   !
  !
  sequence 11
   match tloc
    site-list SITE-40
   !
   action accept
   !
  !
  sequence 21
   match route
    prefix-list DEFAULT
    site-list   SITE-30
   !
   action accept
    set
     preference 100
     service netsvc3 vpn 3
    !
   !
  !
  sequence 31
   match route
    prefix-list DEFAULT
    site-list   SITE-40
   !
   action accept
    set
     preference 50
     service IDP vpn 3
    !
   !
  !
  default-action reject
 !
!
apply-policy
 site-list BRANCHES
  control-policy MY-CONTROL-POLICY-v1 out
 !
!
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow

2. Check OMP peering between WAN Edge and vSmart to ensure policy can be applied on routing updates to/from WAN Edge:

`show omp peers <system-ip> [details]`

```
vsmart1# show omp peers 10.0.0.11
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER            TYPE    ID        ID        ID        STATE    UPTIME           R/I/S
-------------------------------------------------------------------------------------
10.0.0.1        vedge   1         1         30        up       1:18:11:17       18/0/94
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow

3. Check control policy assignment and direction of assignment

`show support omp peer peer-ip <system-ip> | include -pol`

Can be used to find which policies applied to a peer and which site-list it belons:

```
vsmart1# show support omp peer peer-ip 10.0.0.11 | include -pol
        site-pol: BRANCHES route-pol-in: None route-pol-out: MY-CONTROL-POLICY-v1 data-pol-in: None
           data-pol-out: None pfr-pol: None mem-pol: None cflowd:None
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow

4. Check policy definition (In essence, check if vManage UI policy definition was sucessfully translated into CLI syntax on vSmart, **`policy`** section):

```
vsmart1# show running-config policy control-policy MY-CONTROL-POLICY-V1
policy
 control-policy REMOTE-TOPOLOGY-POLICY-PPC-rev1
  sequence 1
   match tloc
    site-list SITE-30
   !
   action accept
   !
  !
  sequence 11
   match tloc
    site-list SITE-40
   !
   action accept
   !
  !
  sequence 21
   match route
    prefix-list DEFAULT
    site-list   SITE-30
   !
   action accept
    set
     preference 100
     service netsvc3 vpn 3
    !
   !
  !
  sequence 31
   match route
    prefix-list DEFAULT
    site-list   SITE-40
   !
   action accept
    set
     preference 50
     service IDP vpn 3
    !
   !
  !
  default-action reject
 !
!
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow

... and **`apply-policy`** section:

```
vsmart1# show running-config apply-policy site-list BRANCHES
apply-policy
 site-list BRANCHES
  control-policy MY-CONTROL-POLICY-v1 out
 !
!

vsmart1# show running-config policy lists site-list BRANCHES
policy
 lists
  site-list BRANCHES
   site-id 11-12
  !
 !
!
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies Troubleshooting Workflow

5. Check control policy match and actions

`test policy match control-policy <name> <conditions>`

- Available starting from 20.8.1. Can be used to find matching sequence in a policy on vSmart

```
vsmart1# test policy match control-policy MY-CONTROL-POLICY-V1 site-id 40 ipv4-prefix DEFAULT
Found: "site-id 40 ipv4-prefix-list DEFAULT" matches policy MY-COJNTROL-POLICY-v1 sequence 31
    sequence: 31
        match route [SITE-LIST PFX-LIST  (0x11) ]
            site-list: SITE-40 (0x7f15b90bfc00)
            IPv4 prefix-list: DEFAULT (0x7f15b90bfc80)
        action: accept
        set: [PREF SERVICE  (0x44) ]
            preference: 50
            service: 3 vpn: 3 tloc: :: : invalid : ipsec [none]
```

6. Examples on routing related policy troubleshooting will folow in Part 2.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Data and AAR Policies Troubleshooting

Order of Operation on WAN Edge

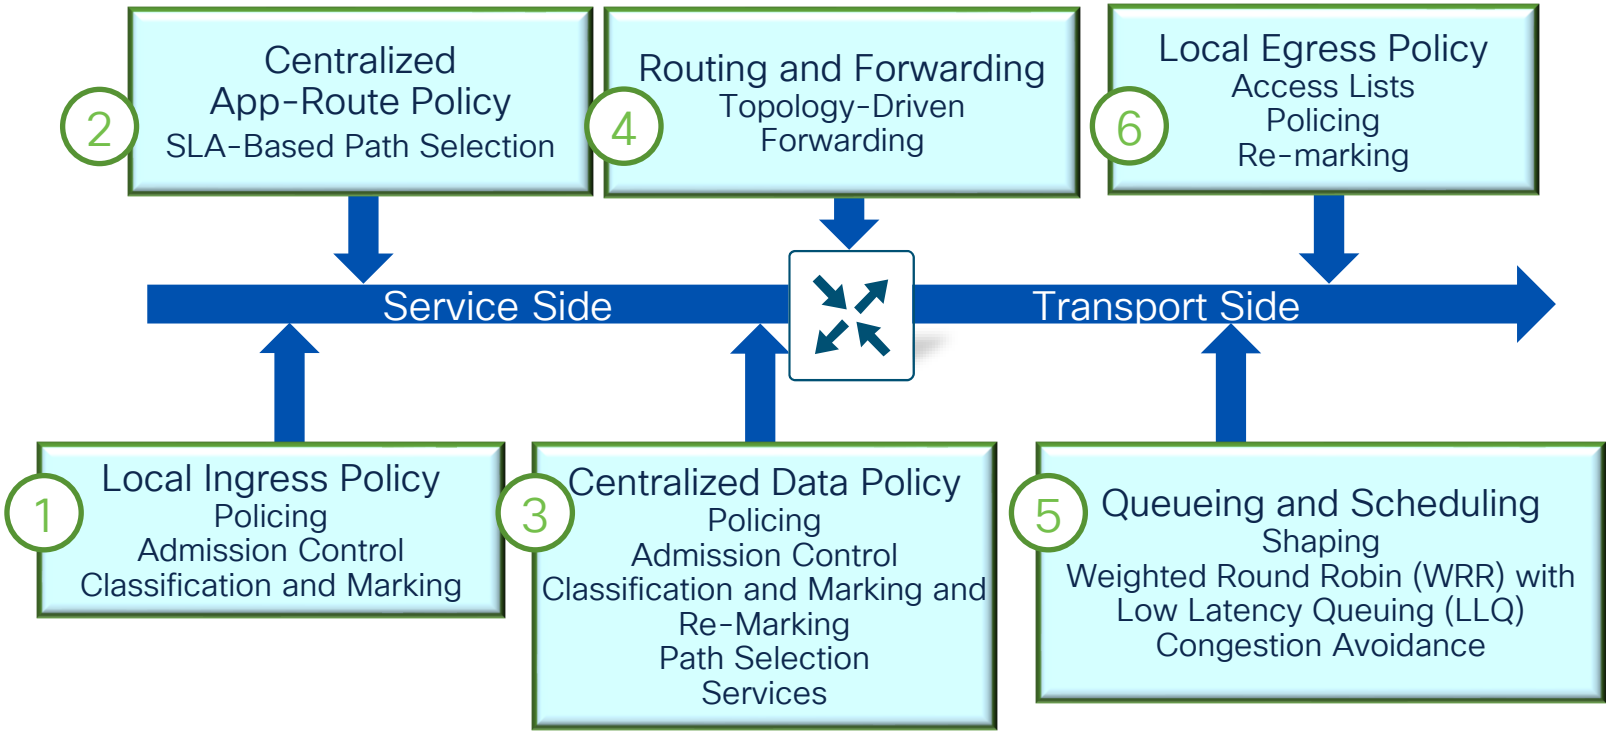

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow

From vSmart perspective, it is similar to control policy workflow:

1. Check policy commit changes:
   `show configuration commit changes <number>`
2. Check OMP peering between WAN Edge and vSmart to ensure policy can be sent:
   `show omp peers <system-ip>`
3. Check AAR/Data policy assignment and direction of assignment
   `show support omp peer peer-ip <system-ip> | include –pol`
4. Check policy definition (vManage UI polciy definition was sucessfully translated into CLI syntax on vSmart):
   `show run policy list <name>`
   `show run policy <data-policy|app-route-policy> <name>`
   `show run apply-policy site-list <data-policy|app-route-policy> <name>`
5. Check policy to XML translation (crafting)*:
   `show support omp peer peer-ip <system-ip>`

The same steps as for control policies

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow (cont.)

From WAN Edge perspective, ensure policy processing:

1. Check policy assignment on WAN Edge

   `show sdwan running-config "policy"`

   ^ for localized policies

   `show sdwan policy from-vsmart`

   ^ for AAR, data policies and cFlow template

2. Ensure correct TLOC/next-hop/color/interface selected accroding to a policy*:

   - For traffic from service-side

     `show sdwan policy service-path vpn <id> interface <name> source-ip <ip-addr> dest-ip <ip-addr> protocol <id> src/dst-port <number> app <name> [all]`

   - For traffic from tunnel-side

     `show sdwan policy tunnel-path vpn <id> interface <name> source-ip <ip-addr> dest-ip <ip-addr> protocol <id> src/dst-port <number> app <name> [all]`

   * Commands can be also used for centralized control policies verification

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow (cont.)

3. Ensure correct policy match occurs from WAN Edge perspective:

- Configure policy counters to ensure traffic match occurs:

```
action [accept|drop]
count <counter name>
```

  - To display counters on the WAN Edge router, depends on type of policy:

```
show sdwan policy <app-route-policy-filter|data-policy-filter|access-list-counters>
```

- Use logging (logs packet header only)

```
action [accept|drop]
log
```

- Use policy troubleshooting tools like packet-trace (CLI) or NWPI (vManage UI)

```
debug platform condition ipv4 <address>/<mask> both
debug platform packet-trace packet <number of packets> [fia-trace]
debug platform condition [start|stop]
show platform packet-trace [summary|packet <number>]
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow (cont.)

4. Other useful Data and AAR policies troubleshooting commands:

- Verify AAR statistics:

```
show sdwan app-route stats
```

- Verify DPI application classification if ”policy app-visibility” enabled (also useful for data policy):

```
show sdwan app-fwd dpi flows
```

- Check traffic flows symmetry and path taken according to neflow data if “policy flow-visibility” or cflowd template configured (also useful control polices):

```
show sdwan app-fwd cflowd flows
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow (cont.)

5. If problems found, proceed to policy programming verification.

Real debuggin just starts here, time to put yourself into my shoes.

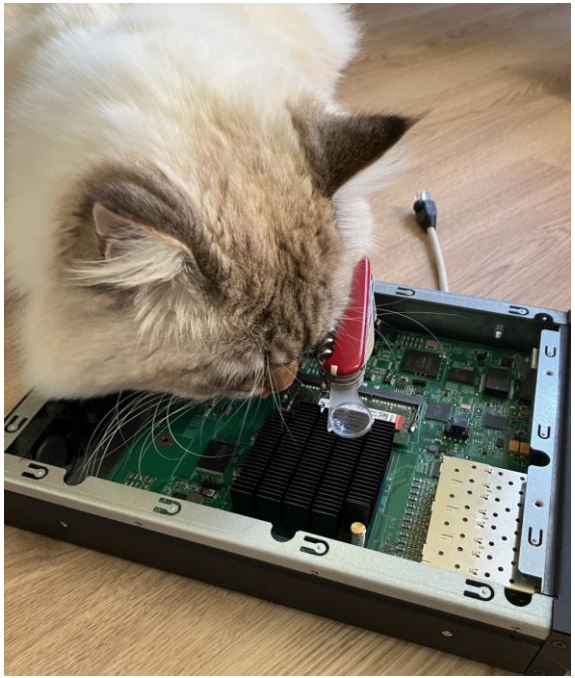

*On the photo my cat trying to put himself my shoes and catch a bug like a mouse ☺

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Side note: how to generate synthetic traffic for testing?

Probem to solve: no user at a site to help with a testing

1. CLI command to trigger synthetic traffic, execution will trigger one probe.
2. Probe result will be reported using log. Use **`show logging`** to see it

CLI request syntax:

`request platform software sdwan synthetic-traffic probe vpn-id 1 url www.cisco.com [dscp <code> [dns <address of nameserver>]`

NOTE: **vpn-id** and **url** is must and **dscp/dns** are optional

Example:

```
cEdge1#request sdwan synthetic-traffic probe vpn-id 1 url www.cisco.com

*Apr 11 02:05:34.302: %Cisco-SDWAN-Site25-cEdge-1-DBGD-6-INFO-1500002: Synthetic test probe
result for app: Def-Test, url: www.cisco.com, src_intf: GigabitEthernet7, latency: 253, loss:
0%, score: 6, count 1
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Data and AAR Policies Troubleshooting:

Commands usage examples

cisco Live!

Data and AAR Policies Troubleshooting Workflow (cont.)

From vSmart perspective the same steps as for control policies (so we won't repeat them here), except additional step 5. Check policy XML translation (crafting):

```
vsmart1# show support omp peer peer-ip 10.0.0.11 | begin "Policy received" | until "Statistics"
        Policy received: Complete
       Forwarding policy len: 632
<data-policy>
  <name>VPN_1_NAT</name>
  <vpn-list>
    <name>VPN_1</name>
    <sequence>
      <seq-value>1</seq-value>
      <match>
        <source-data-prefix-list>LAN</source-data-prefix-list>
      </match>
      <action>
        <action-value>accept</action-value>
        <nat>
          <use-vpn>0</use-vpn>
        </nat>
      </action>
    </sequence>
  </vpn-list>
<direction>from-service</direction></data-policy><lists><vpn-list>
  <name>VPN_1</name>
  <vpn>
    <id>1</id>
  </vpn>
</vpn-list>
<data-prefix-list>
  <name>LAN</name>
  <ip-prefix>
    <ip>10.10.10.0/24</ip>
  </ip-prefix>
</data-prefix-list>
</lists>
        Statistics:
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow

From WAN Edge perspective

1. Check policy assignment on WAN Edge:

```
cE1_BR1#show sdwan policy from-vsmart
from-vsmart data-policy VPN_1_NAT
 direction from-service
 vpn-list VPN_1
  sequence 1
   match
    source-data-prefix-list LAN
   action accept
    nat use-vpn 0
    no nat fallback
  default-action drop
from-vsmart lists vpn-list VPN_1
 vpn 1
from-vsmart lists data-prefix-list LAN
 ip-prefix 10.10.10.0/24
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow (2)

From WAN Edge perspective

2. Ensure correct TLOC/next-hop/color/interface selected as a result of a policy*:

```
cE1_BR1#show sdwan policy service-path vpn 1 interface GigabitEthernet 4 source-ip 10.10.10.10 dest-ip 1.1.1.1
protocol 17 dest-port 53
Next Hop: Remote
  Remote IP: 192.168.10.1, Interface GigabitEthernet3 Index: 9
```

Example of problematic state:

```
cE1_BR1#show sdwan policy service-path vpn 1 interface GigabitEthernet 4 source-ip 10.10.1.10 dest-ip 1.1.1.1
protocol 17 dest-port 53 app dns
Next Hop: Blackhole
```

* can be used also for verification of routing decision as a result of centralized control policy as well

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Data and AAR Policies Troubleshooting Workflow (3)

From WAN Edge perspective

3. Ensure correct policy match occurs from WAN Edge perspective:

```
vsmart1# show configuration commit changes 0
policy
 data-policy VPN_1_NAT
  vpn-list VPN_1
   sequence 1
    action accept
     count NATed_pkts
    !
   !
  !
 !
!
```

```
cE1_BR1#show sdwan policy data-policy-filter
VPN_1_NAT
data-policy-filter VPN_1_NAT
 data-policy-vpnlist VPN_1
  data-policy-counter NATed_pkts
   packets 5
   bytes   500
  data-policy-counter default_action_count
   packets 76652
   bytes   9023632
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-WAN Policy troubleshooting tools
(step 3. ensure correct policy match and actions)

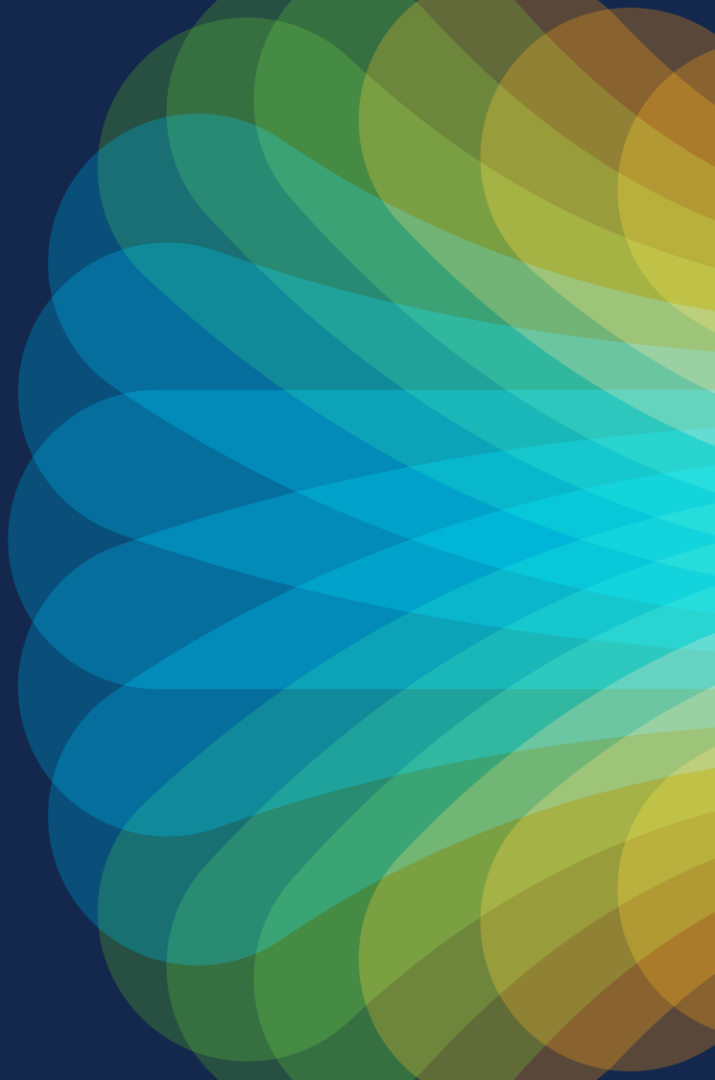

Packet-trace
a.k.a
FIA-trace

Enabling packet-trace

Set debug conditions (match filter) and enable packet-trace:

```
cEdge1#debug platform condition <ipv4|ipv6|mac|mpls> <address/mask | access-list name> both
cEdge1#debug platform packet-trace packet <number of packets> [fia-trace]
cEdge1#debug platform condition start
```

Optionally, dump packet data in a hex format:

```
cEdge1#debug platform packet-trace copy packet both size <…>
```

If you want to trace only internally dropped packets, check QFP statistics first:

```
cE1_BR1#show platform hardware qfp active statistics drop detail
Last clearing of QFP drops statistics : never

-------------------------------------------------------------------------------
   ID  Global Drop Stats                          Packets                Octets
-------------------------------------------------------------------------------
   62  IpTtlExceeded                                 28                  1748
   56  IpsecInput                                    14                  2402
   19  Ipv4NoRoute                                 4909                786205
  483  SdwanDataPolicyDrop                          650                 78230
  479  SdwanImplicitAclDrop                      261280              44905782
```

Then enable trace only for specific drop code ID:

```
cEdge1#debug platform packet-trace drop code <id>
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using packet-trace

You can check overall statistics and number of packets captured:

```
cEdge1# show platform packet-trace statistics
Packets Summary
  Matched  1165
  Traced   1024
Packets Received
  Ingress  1085
  Inject   80
    Count       Code  Cause
    80          3     QFP IPv4/v6 nexthop lookup
Packets Processed
  Forward  928
  Punt     0
  Drop     0
  Consume  237
```

To stop packet-trace and clear all conditions (filters):

```
cEdge1# debug platform condition stop
cEdge1# clear platform condition all
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using packet-trace (2)

Show captured packets summary:

```
cEdge1#show platform packet-trace summary
Pkt    Input                     Output                    State  Reason
0      Gi2                       Gi3                       FWD
1      Tu3                       Gi2                       FWD
2      INJ.3                     Gi2                       FWD
3      internal0/0/recycle:0     Gi2                       FWD
4      Gi2.                      Tu3                       DROP   525 (NoStatsUpdate)
5      internal0/0/recycle:0     Gi2                       FWD
```

BFD return packet dropped, expected

Details of specific packet:

```
cEdge1# show platform packet-trace packet <packet number>
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Packet-trace output

```
cEdge1#show platform packet-trace packet 0
Packet: 0           CBUG ID: 35949496
Summary
  Input     : GigabitEthernet2
  Output    : GigabitEthernet3
  State     : FWD
  Timestamp
    Start   : 1214211941024994 ns (02/24/2020 11:03:14.435466 UTC)
    Stop    : 1214211941530105 ns (02/24/2020 11:03:14.435971 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet2
    Output      : <unknown>
    Source      : 192.168.11.254
    Destination : 192.168.17.254
    Protocol    : 1 (ICMP)
<remove>
  Feature: SDWAN ACL IN
    Interface    : GigabitEthernet2
    CG           : 3
    Seq          : 21
    Policy Flags : 0x100
    Action : SET_FWD_CLASS 3 Prec3
  Feature: SDWAN_ACL_IN
    Entry       : Input - 0x81845740
    Input       : GigabitEthernet2
    Output      : <unknown>
    Lapsed time : 815733 ns
<removed>
```

Local policy in action here

SD-WAN ACL matches flow and assign to QoS class ”Prec3”

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Packet-trace output (2)

```
<skipped>
Feature: NBAR
    Packet number in flow: N/A
    Classification state: Final
    Classification name: ping
<skipped>
Feature: SDWAN App Route Policy
    VRF           : 1
    CG            : 1
    Seq           : 65535
    SLA           : __all_tunnels__ (0)
    Policy Flags : 0x2
    SLA Strict    : No
    Preferred Color : 0x0 none
<removed>
Feature: SDWAN OCE
    Hash Value  : 0xaf6f0c4e
    Encap       : ipsec
    SLA         : 0
    SDWAN VPN   : 1
    SDWAN Proto : IPV4
    Out Label   : 1001
    Local Color : biz-internet
    Remote Color: biz-internet
    FTM Tunnel ID:15
    SDWAN Session Info
      SRC IP      : 172.16.11.254
      SRC Port    : 12346
      DST IP      : 172.16.17.254
      DST Port    : 12346
      Remote System IP : 172.16.255.17
```

NBAR classification is completed
Application is recognized

This flow does not match any app-route policies, so it's load balanced to all available tunnels

Forwarding decision

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Packet-trace output (3)

```
<removed>
 Feature: SDWAN QoS Output
    Fwd Class     : 3
    QoS Queue     : 3
    DSCP Rewrite : No
    CoS  Rewrite : No
<removed>
 Feature: IPSec
    Result     : IPSEC_RESULT_SA
    Action     : ENCRYPT
    SA Handle : 45
    Peer Addr : 172.16.17.254
    Local Addr: 172.16.11.254
<removed>
 Feature: MARMOT_SPA_D_TRANSMIT_PKT
    Entry        : Output - 0x817f19f4
    Input        : GigabitEthernet3
    Output       : GigabitEthernet3
    Lapsed time : 217840 ns
```

QoS queueing in queue3

Encrypting

Packet is transmitted on an interface

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

SD-WAN Datapath: packet-trace to the rescue

Green: Service interface [VPN 1]
Blue: MPLS interface [VPN 0]
Red: biz-internet interface [VPN 0]

```
cEdge1#show platform packet-trace summary
Pkt    Input                Output               State  Reason
0      Gi2                  Gi1                  FWD
1      Gi2                  Gi1                  FWD
2      Tu1                  Gi2                  FWD
3      Gi2                  Gi1                  FWD
4      Gi2                  Gi3                  FWD
5      Tu3                  Gi2                  FWD
6      Tu3                  Gi2                  FWD
```

Why traffic flows to/from Gig1 at the beginning and then switch to Gig3 ?

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

SDWAN Datapath: packet-trace to the rescue (2)

First let's check very first packet in the flow:

```
cEdge1#show platform packet-trace packet 0
Packet: 0           CBUG ID: 423
Summary
  Input     : GigabitEthernet2
  Output    : GigabitEthernet1
  State     : FWD
<removed>
 Feature: NBAR
    Packet number in flow: 1
    Classification state: Not final
<removed>
 Feature: SDWAN App Route Policy
    VRF          : 1
    Seq          : 1
    SLA          : __all_tunnels__ (0)
    Policy Flags : 0x0
    SLA Strict   : No
    Preferred Color : 0x0 none
```

NBAR application recognition does not have enough information to recognize the application yet

SDWAN App route policy will consider all available tunnels

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

SDWAN Datapath: packet-trace to the rescue (3)

Then let's check some later packet from the service side:

```
cEdge1#show platform packet-trace packet 4
Packet: 4           CBUG ID: 423
Summary
  Input     : GigabitEthernet2
  Output    : GigabitEthernet3
  State     : FWD
<removed>
 Feature: NBAR
    Packet number in flow: 5
    Classification state: Final
    Classification name: ssh
    Classification ID: [IANA-L4:22]
<removed>
Feature: SDWAN App Route Policy
    VRF           : 1
    Seq           : 0
    SLA           : TEST1 (1)
    Policy Flags : 0x1
    SLA Strict    : Yes
    Preferred Color : 0x10 biz-internet
    Tunnel Match Reason : MATCHED_SLA_AND_PREF_COLOR
```

NBAR application recognition has discovered the application as SSH

SDWAN App route policy will optimize the flow towards biz-internet tunnel

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Network Wide Path Insight (NWPI)

Enabling NWPI

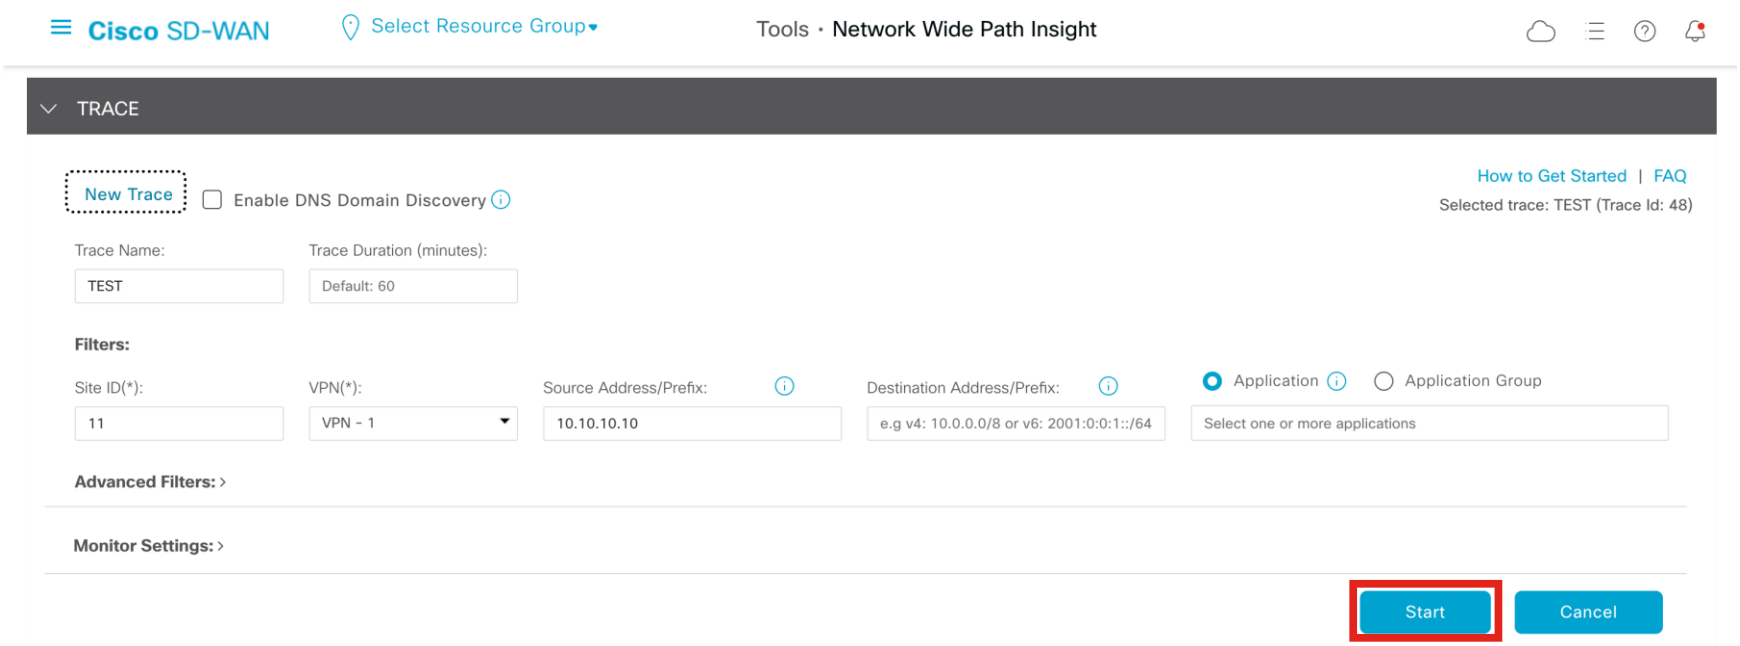

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using NWPI

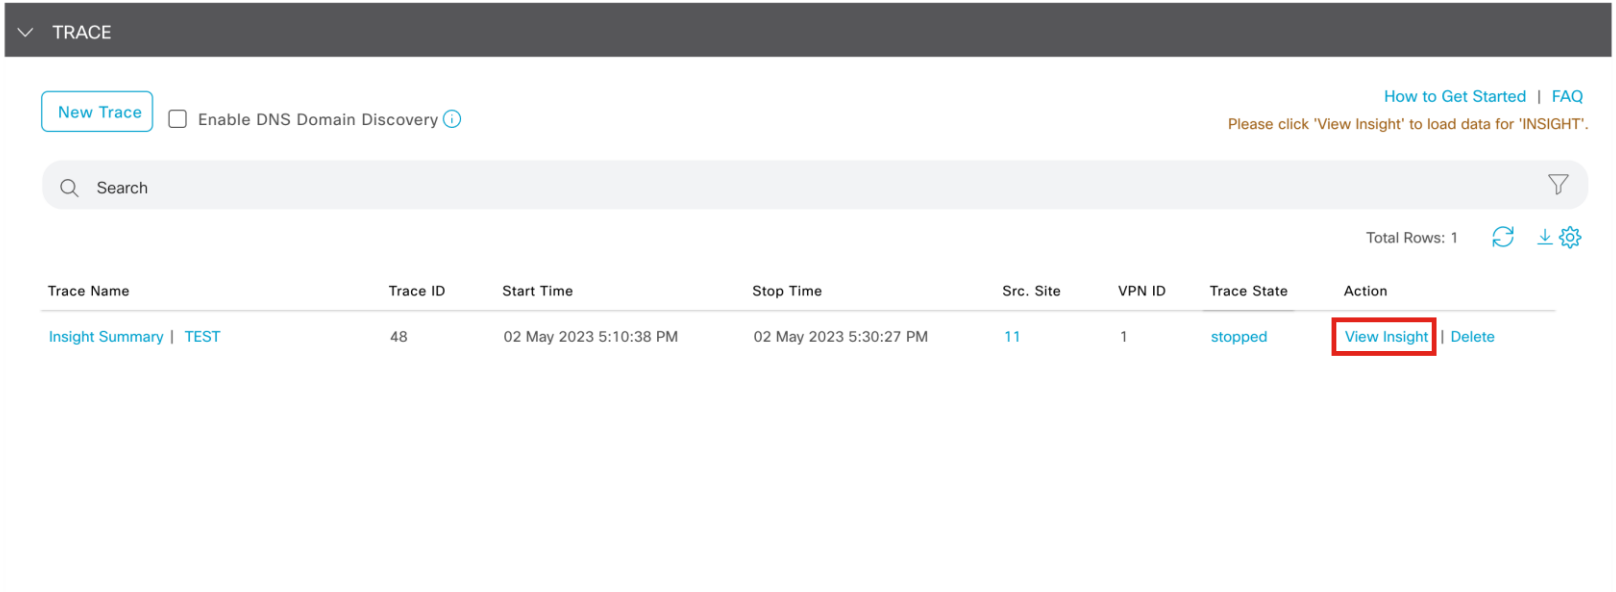

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using NWPI (2)

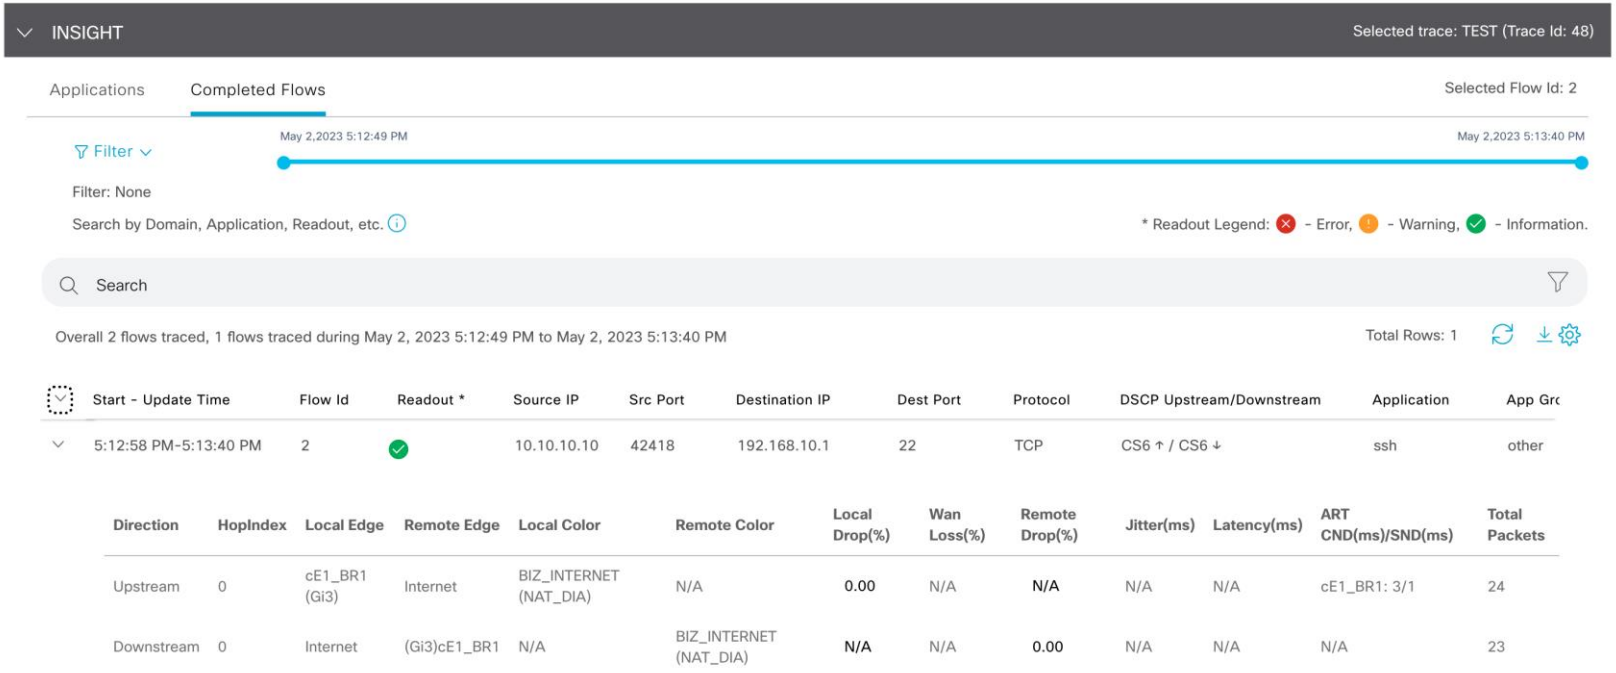

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using NWPI (3)

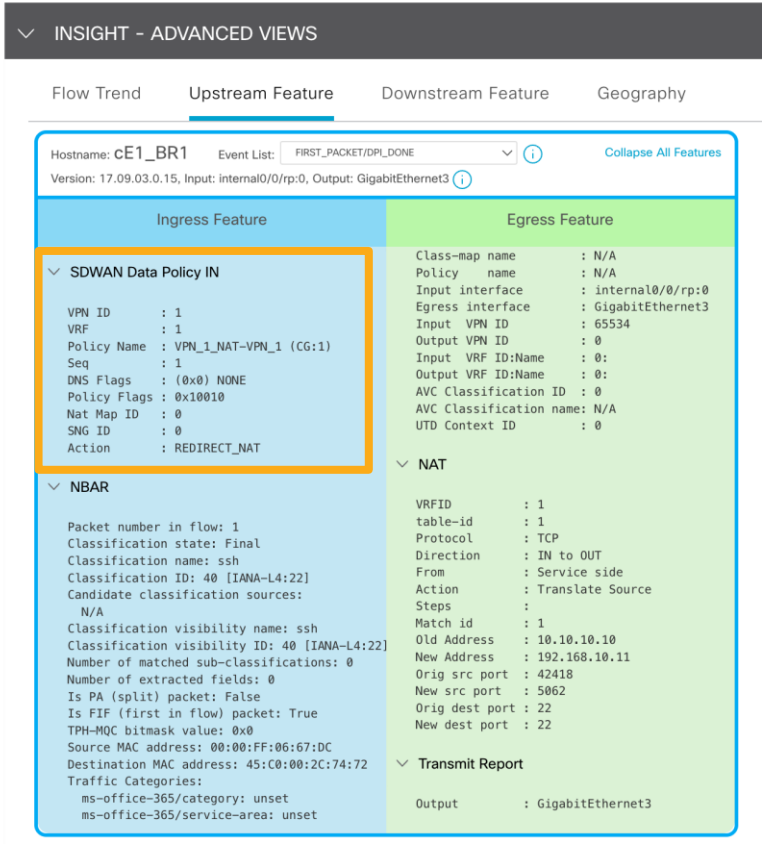

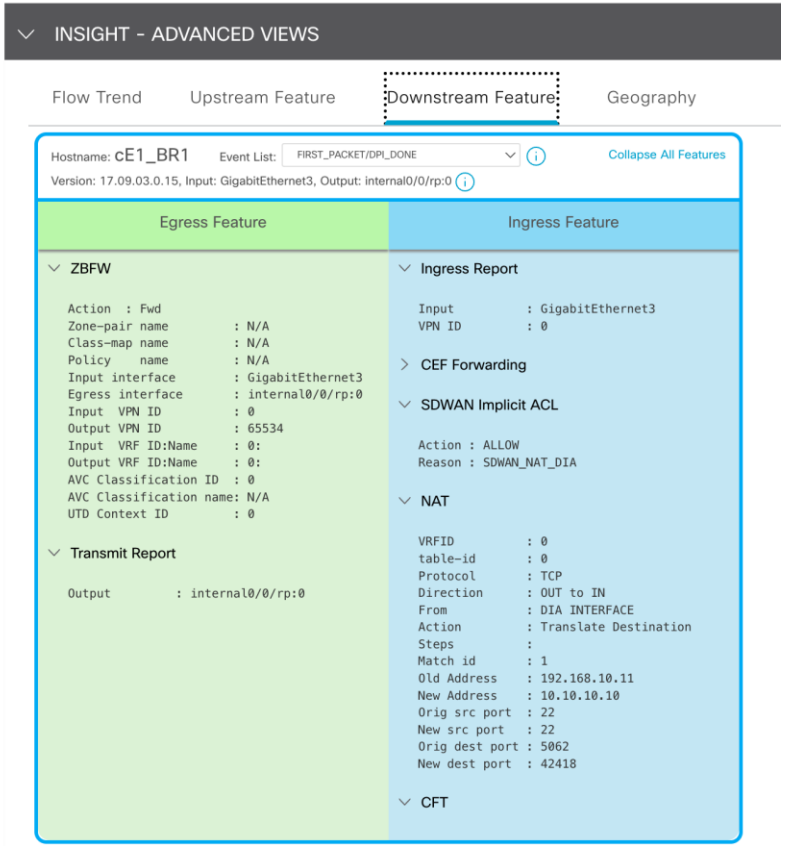

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

What's new in 20.12/17.12: Synthetic Traffic

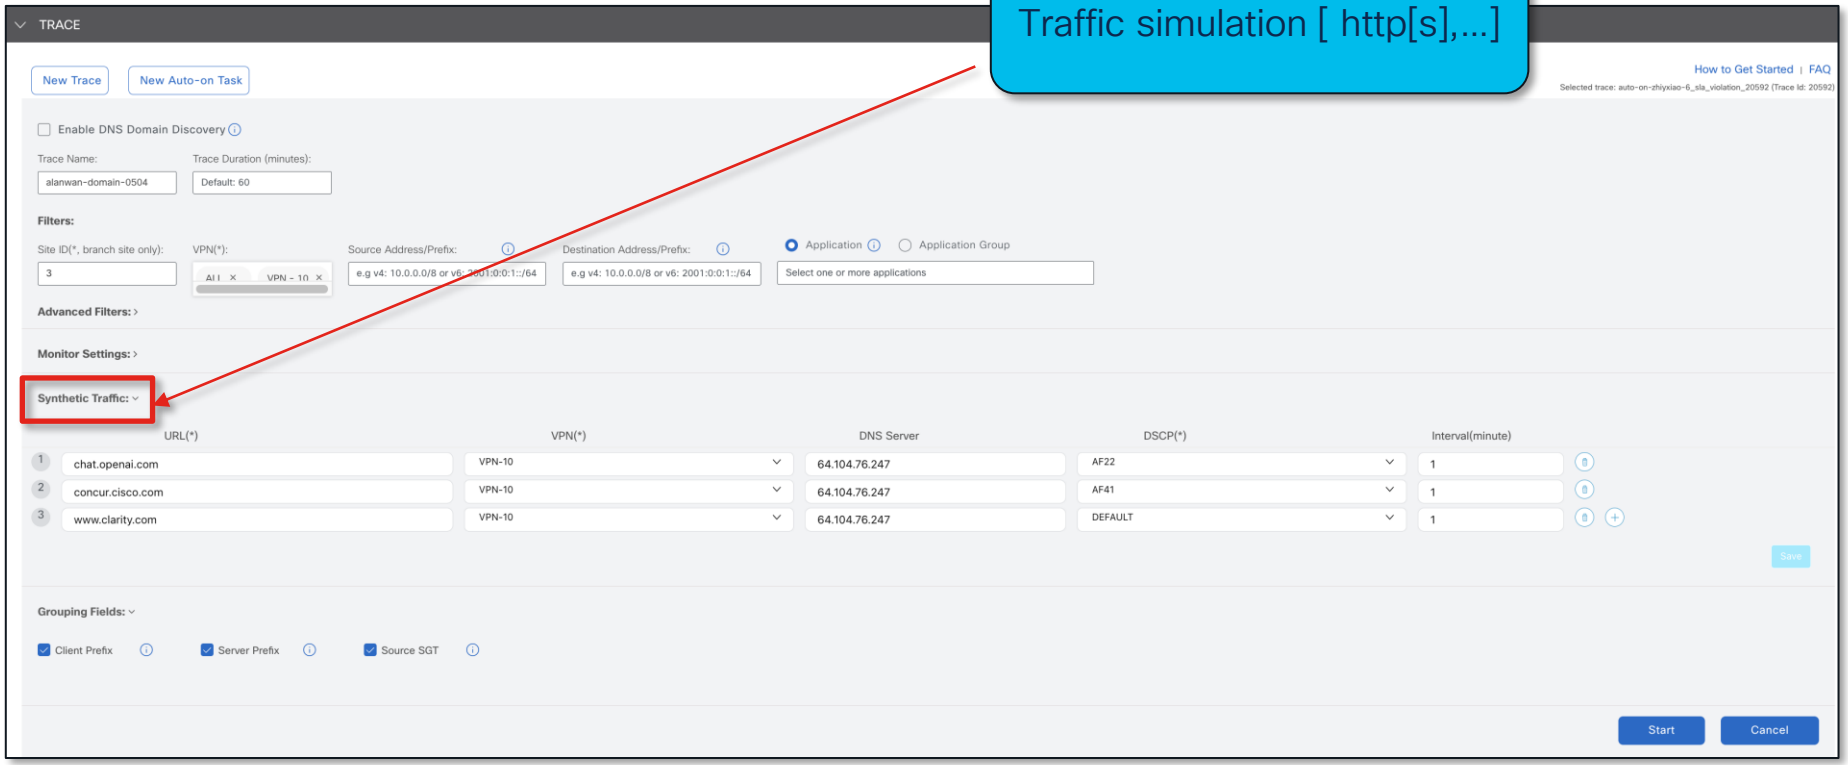

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Other useful commands for AAR and Data Policies troubleshooting (step 4)

Other Useful Commands for AAR troubleshooting (1)

```
cE1_BR1#show sdwan app-route stats remote-color biz-internet remote-system-ip 169.254.206.37 summary
Generating output, this might take time, please wait ...
app-route statistics 192.168.10.11 192.168.10.37 ipsec 12346 12406
 remote-system-ip         169.254.206.37
 local-color              custom2
 remote-color             biz-internet
 sla-class-index          0
 fallback-sla-class-index None
 enhanced-app-route       Disabled
 sla-dampening-index      None
 app-probe-class-list None
  mean-loss    0.000
  mean-latency 0
  mean-jitter  0
      TOTAL                      AVERAGE  AVERAGE  TX DATA       RX DATA       IPV6 TX       IPV6 RX
INDEX  PACKETS       LOSS         LATENCY  JITTER   PKTS          PKTS          DATA PKTS     DATA PKTS
-------------------------------------------------------------------------------------------------------
-
0      661           0            0        0        0             0             0             0
1      664           0            0        0        0             0             0             0
2      663           0            0        0        0             0             0             0
3      662           0            0        0        0             0             0             0
4      665           0            0        0        0             0             0             0
5      664           0            0        0        0             0             0             0
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Other Useful Commands for AAR troubleshooting (2)

```
cE1_BR1#show sdwan app-fwd dpi flows vpn 4
Generating output, this might take time, please wait
app-fwd cflowd flows vpn 4 src-ip 192.168.5.197 dest-ip
192.168.4.196 src-port 36470 dest-port 22 dscp 10 ip-
proto 6
 tcp-cntrl-bits           24
 icmp-opcode              0
 total-pkts               61
 total-bytes              4080
 start-time               "Tue Jan 16 15:26:56 2024"
 egress-intf-name         GigabitEthernet4
 ingress-intf-name        GigabitEthernet3
 application              ssh
 family                   terminal
 drop-cause               "No Drop"
 drop-octets              0
 drop-packets             0
 sla-not-met              0
 color-not-met            0
 queue-id                 2
 initiator                1
 tos                      40
 dscp-output              10
 sampler-id               0
 fec-d-pkts               0
 fec-r-pkts               0
 pkt-dup-d-pkts-orig      0
 pkt-dup-d-pkts-dup       0
 pkt-dup-r-pkts           0
 pkt-cxp-d-pkts           0
 category                 0
 service-area             0
 cxp-path-type            0
 region-id                0
 ssl-read-bytes           0
 ssl-written-bytes        0
 ssl-en-read-bytes        0
 ssl-en-written-bytes     0
 ssl-de-read-bytes        0
 ssl-de-written-bytes     0
 ssl-service-type         0
 ssl-traffic-type         0
 ssl-policy-action        0
 appqoe-action            0
 appqoe-sn-ip             0.0.0.0
 appqoe-pass-reason       0
 appqoe-dre-input-bytes   0
 appqoe-dre-input-packets 0
 appqoe-flags             0
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Other Useful Commands for AAR troubleshooting (3)

```
cE1_BR1#show sdwan app-fwd cflowd flows vpn 4
Generating output, this might take time, please
wait ...
app-fwd cflowd flows vpn 4 src-ip 192.168.5.197
dest-ip 192.168.4.196 src-port 22 dest-port 37748
dscp 4 ip-proto 6
 tcp-cntrl-bits           24
 icmp-opcode              0
 total-pkts               6
 total-bytes              2064
 start-time               "Fri Dec 22 15:35:11
2023"
 egress-intf-name         GigabitEthernet4
 ingress-intf-name        GigabitEthernet3
 application              ssh
 family                   terminal
 drop-cause               "No Drop"
 drop-octets              0
 drop-packets             0
 sla-not-met              0
 color-not-met            0
 queue-id                 2
 initiator                2
 tos                      0
 dscp-output              0
 sampler-id               0
 fec-d-pkts               0
 fec-r-pkts               0
 pkt-dup-d-pkts-orig      0
 pkt-dup-d-pkts-dup       0
 pkt-dup-r-pkts           0
 pkt-cxp-d-pkts           0
 category                 0
 service-area             0
 cxp-path-type            0
 region-id                0
 ssl-read-bytes           0
 ssl-written-bytes        0
 ssl-en-read-bytes        0
 ssl-en-written-bytes     0
 ssl-de-read-bytes        0
 ssl-de-written-bytes     0
 ssl-service-type         0
 ssl-traffic-type         0
 ssl-policy-action        0
 appqoe-action            0
 appqoe-sn-ip             0.0.0.0
 appqoe-pass-reason       0
 appqoe-dre-input-bytes   0
 appqoe-dre-input-packets 0
 appqoe-flags             0
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy programming low-level verification (step 5)

Down the rabbit hole. Are you ready?

*Not a stock photo. Rabbit hole near our Brussles office

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Policy Installation Workflow from IOS-XE Perspective

WAN Edge running IOS-XE

Generalized Software Architecture

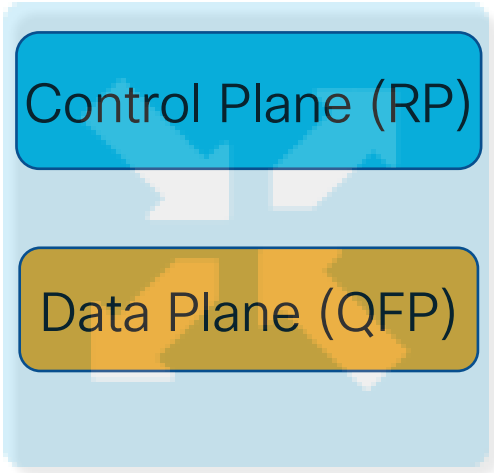

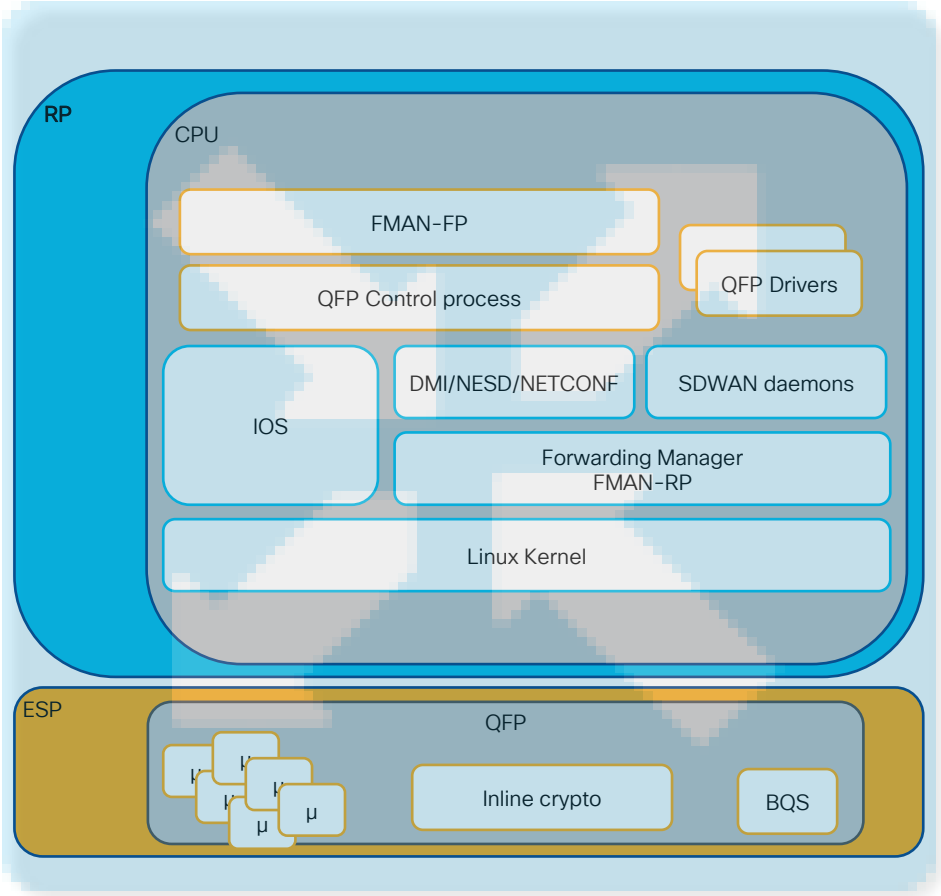

- All platforms share similar architecture
- Key differences:
  - the **data plane (QFP)** is either dedicated CPU/linecard or a Linux software process
  - Crypto implemented either inline or via external crypto accelerator/hardware/ASIC
- Same troubleshooting toolset and approach!

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

DTLS to controllers (1)

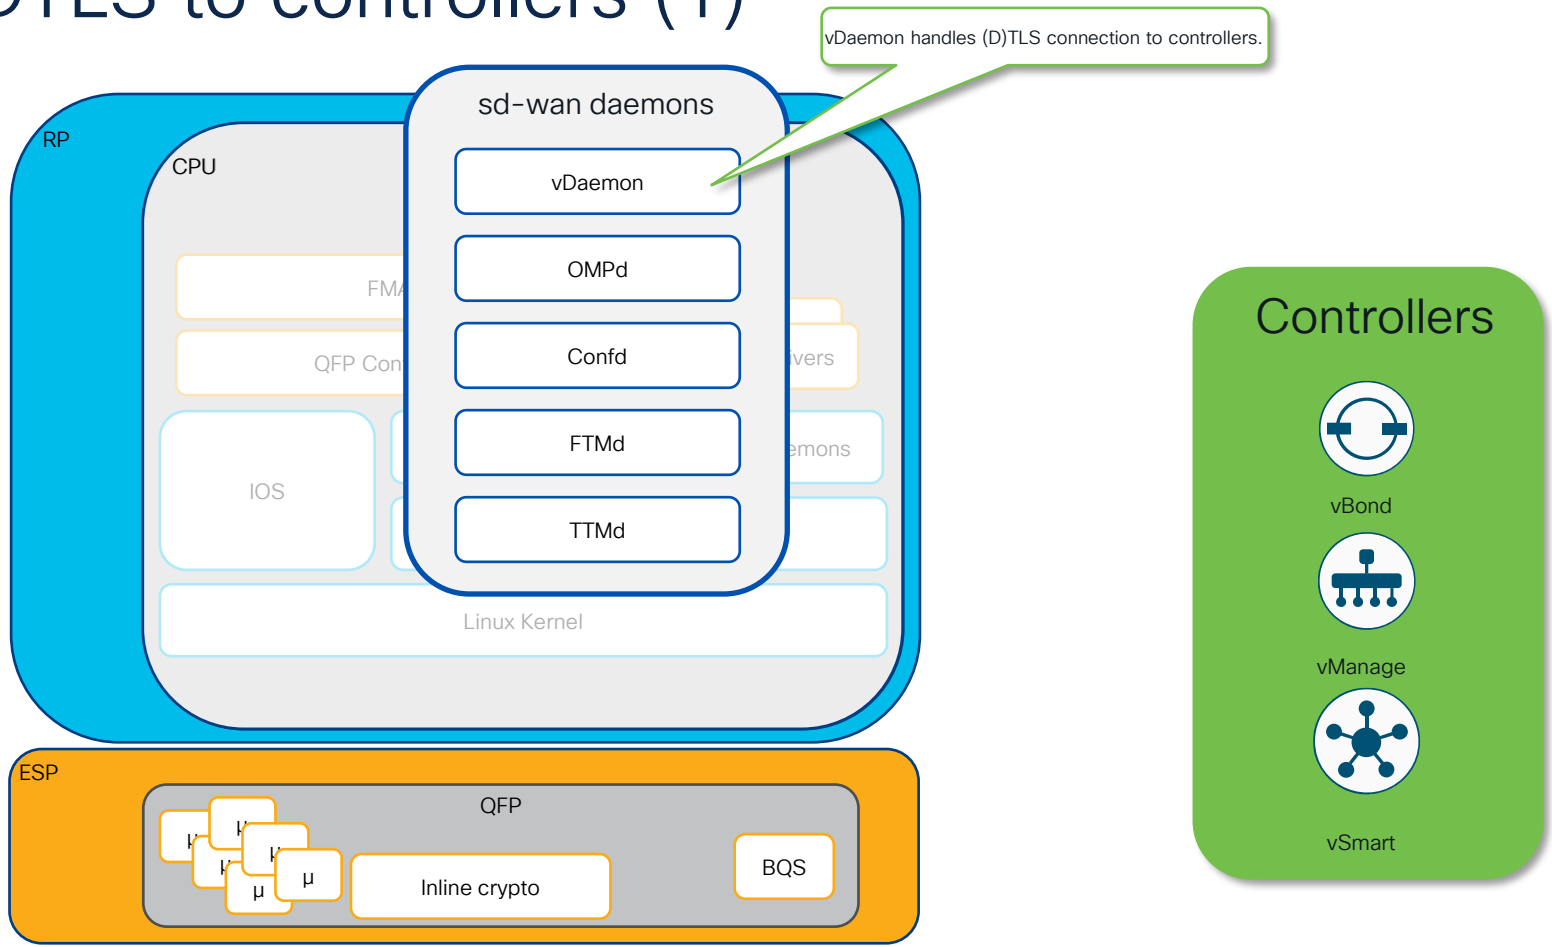

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

DTLS to controllers (2)

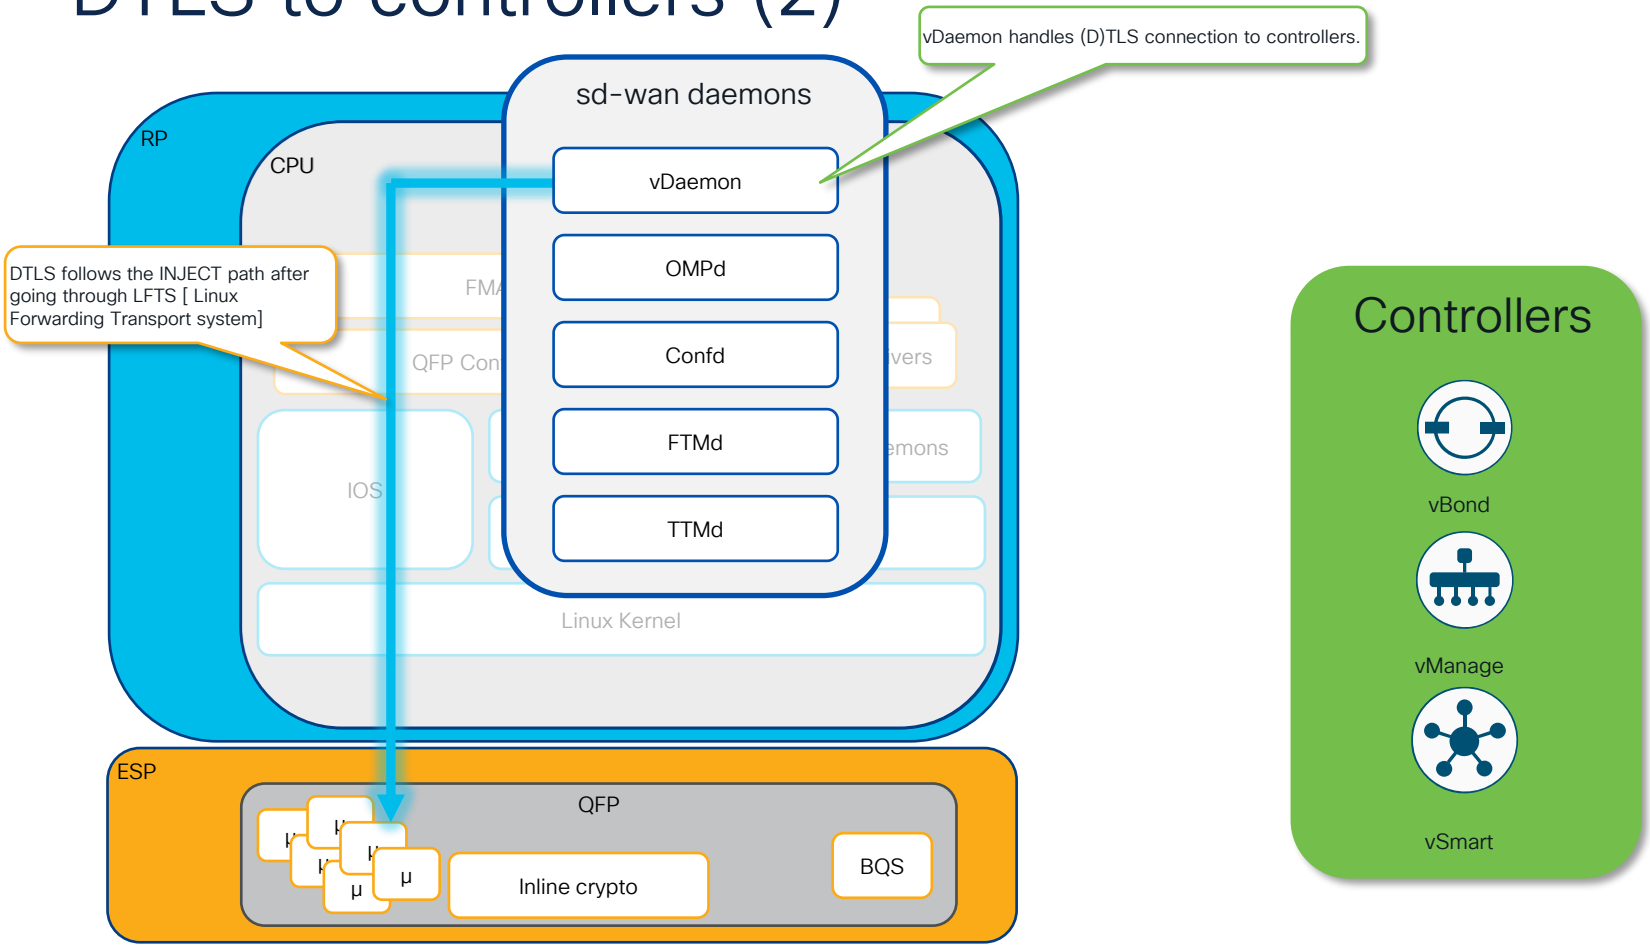

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

DTLS to controllers(3)

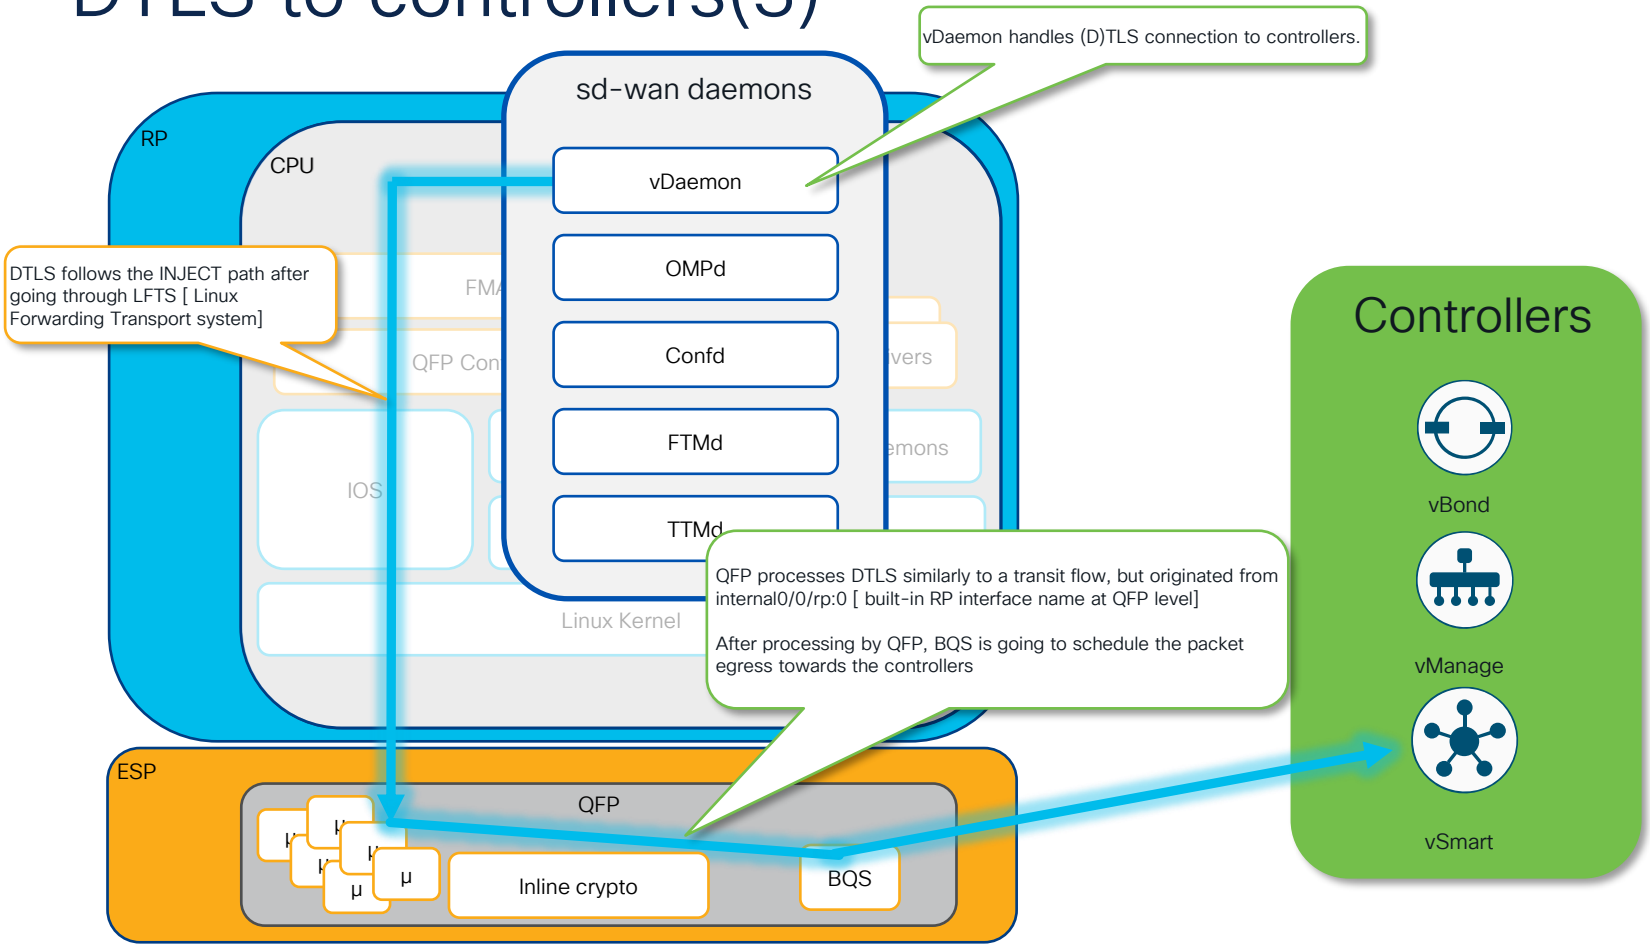

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

DTLS to controllers (4)

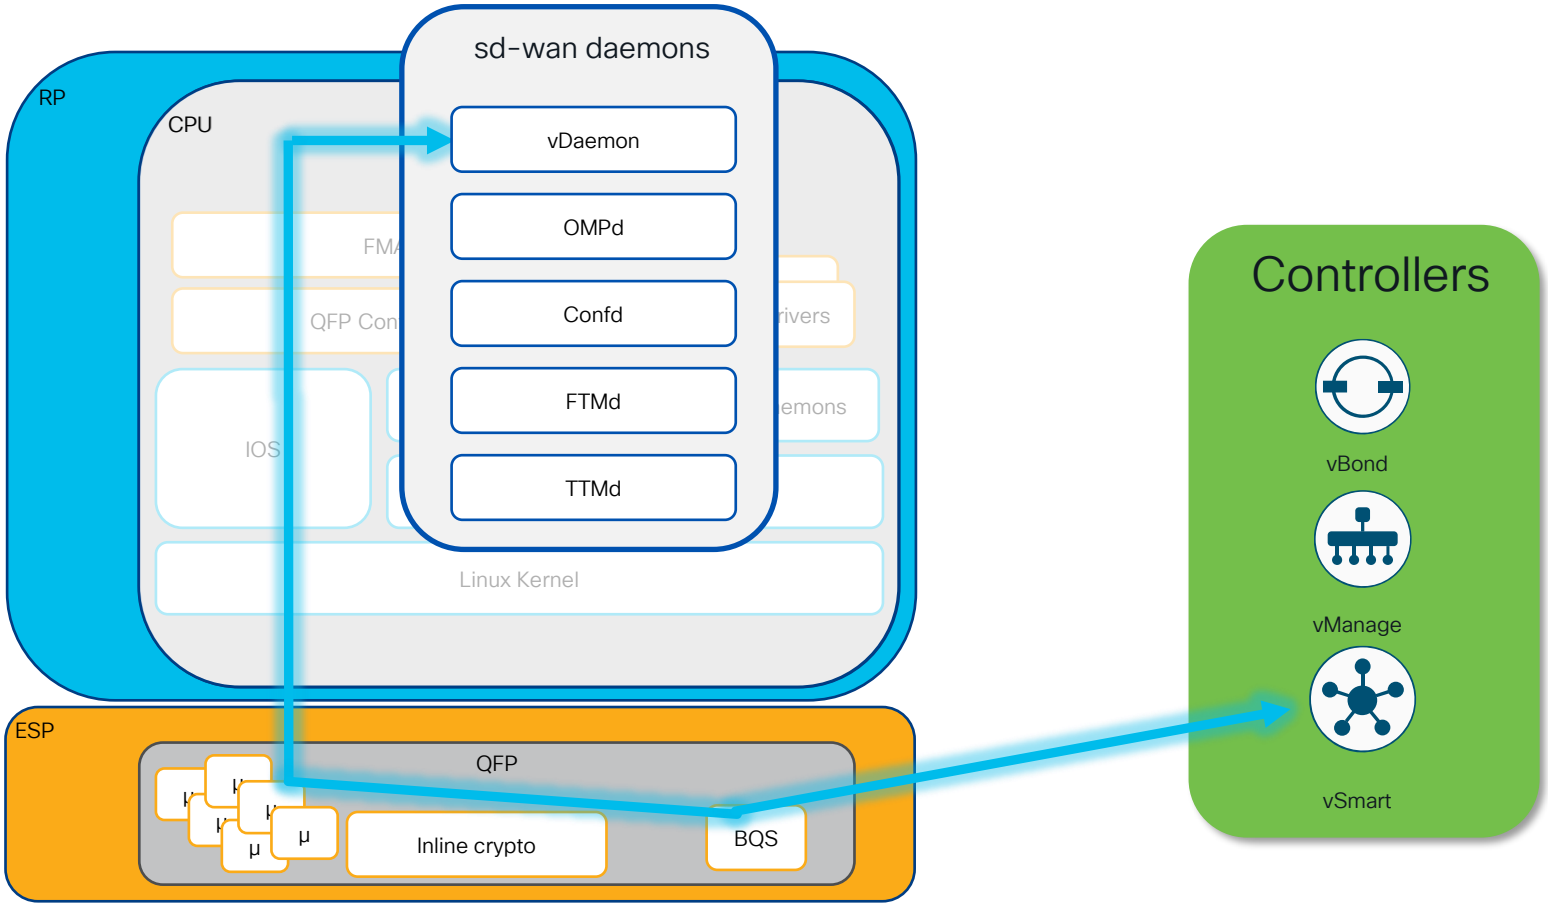

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

OMP centralized policy update (1)

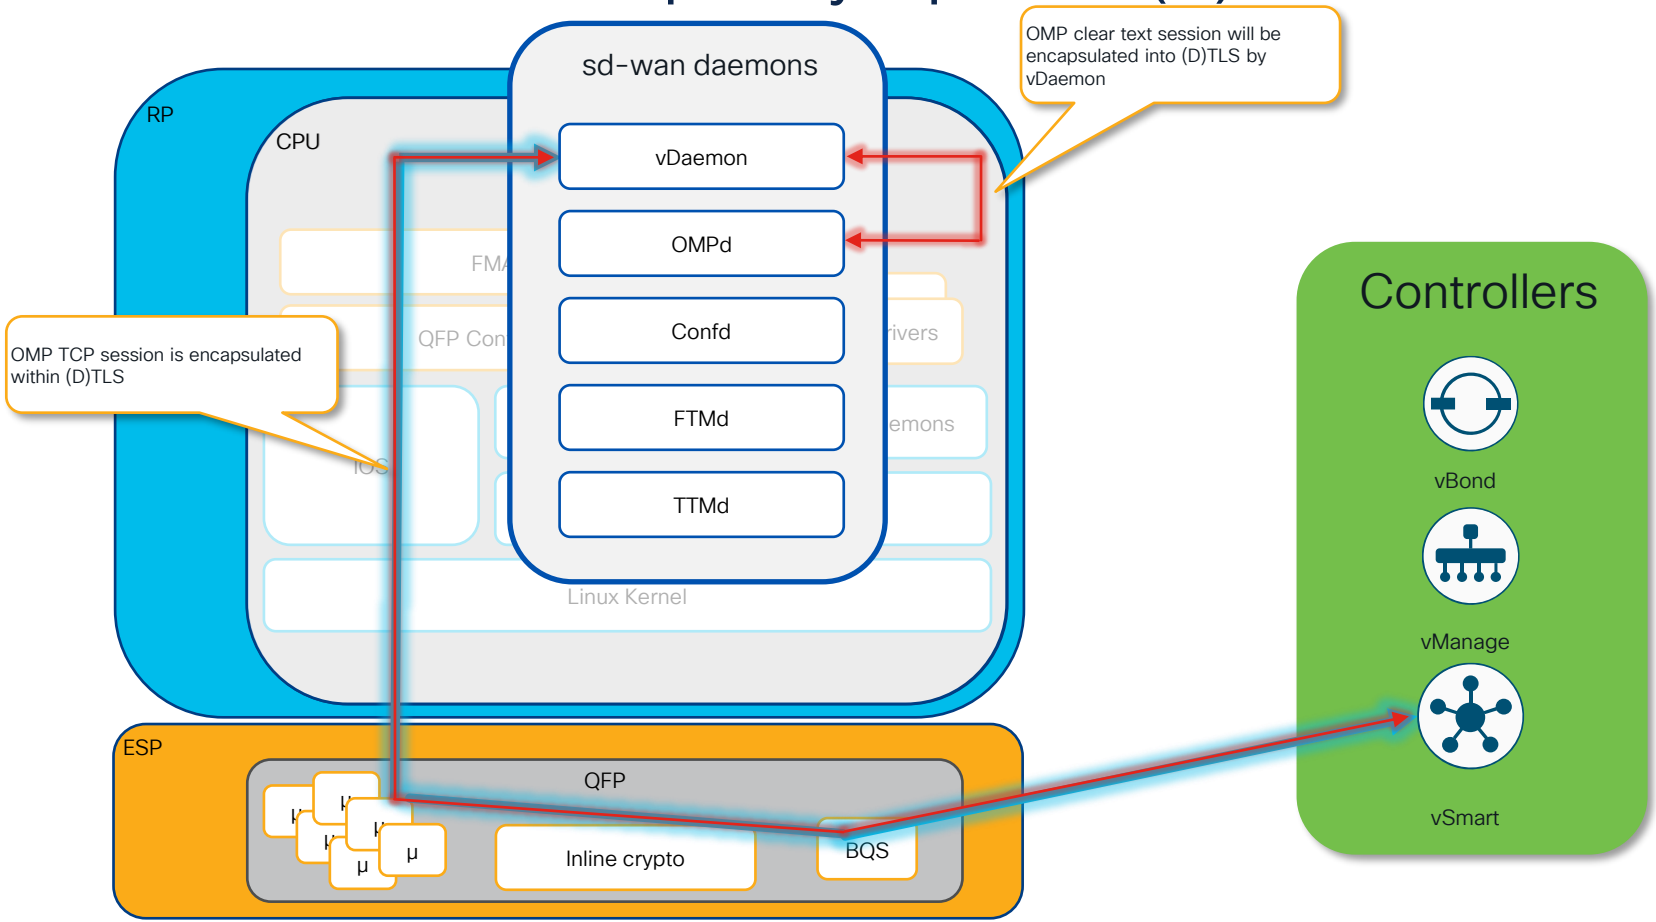

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

OMP centralized policy update (2)

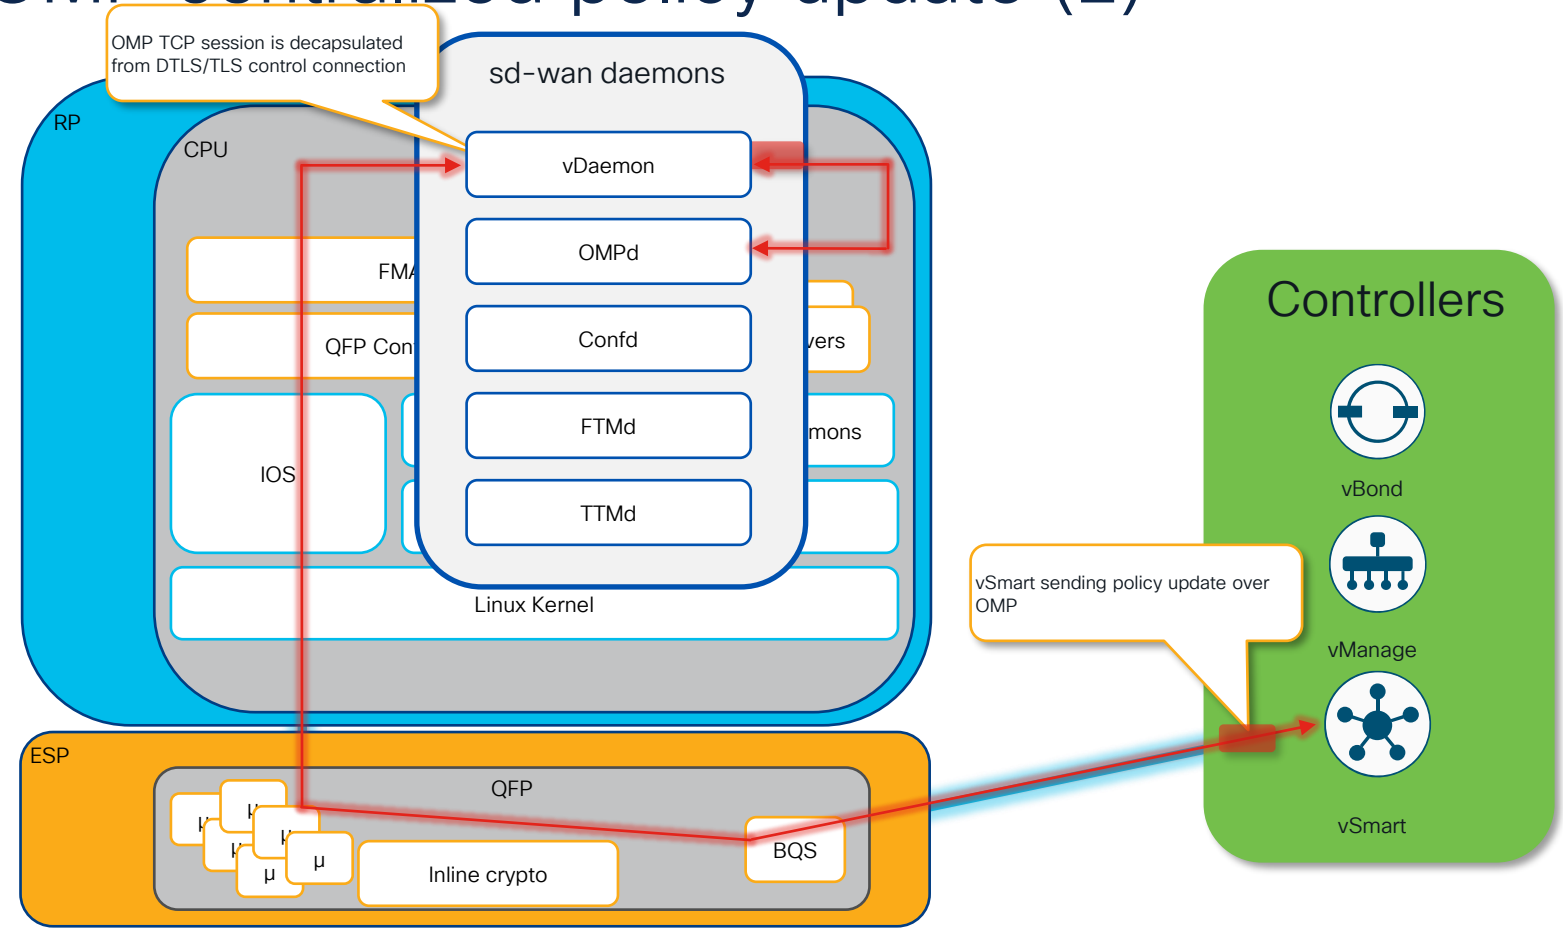

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

OMP centralized policy update (3)

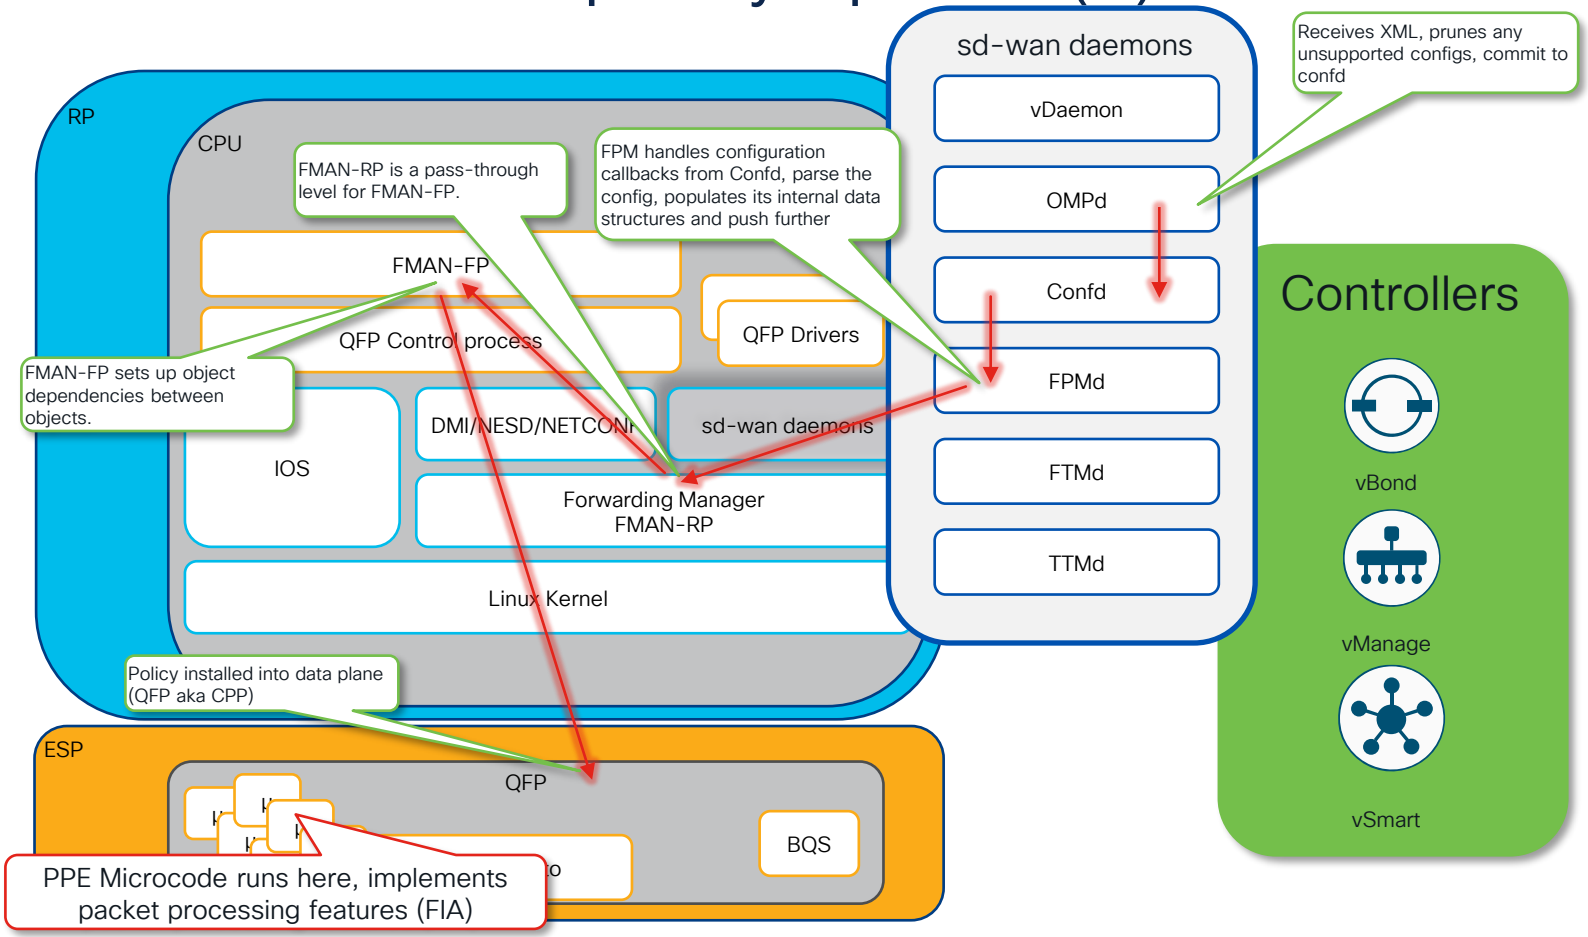

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Key takeaways from the section

- All hardware platfroms running IOS-XE share similar architecture
- Route Processor (RP) runs Linux kernel and various daemons responsible for control plane implementation and inter-process communications
- Data plane aka Quantum Flow Processor (QFP) implements various traffic processing features via Feature Invocation Array (FIA). Data/AAR/ACL/QoS/Security policies (features) executed here.
- Control plane traffic flows through QFP similarly to transit traffic
- Same toolset can be used for troubleshooting on different platforms in the same manner

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Back on track: Policy programming low-level verification (step 5)

“How many times have you been on a WebEx with TAC and thought – ‘Wow – I wish I knew all those cool ninja TAC commands!’?”

One of our customers according to one of my colleagues

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Why? Certainly for fun! Put yourself into shoes of TAC engineer ☺

When? For example, policy does not work and log message like below was seen:

```
%FMFP-3-OBJ_DWNLD_TO_DP_STUCK: AOM download to Data Plane is stuck for more
than 1800 seconds for obj[180464] type[711] pending-issue Req-create Issued-
none 'class class_name NEW-POLICY-seq-10 class_key 12:10’
```

Only few basic things you need to know and remember (simplified):

- AOM is Asynchronous Object Manager, our internal software stuff that allows processes to continue with other tasks without waiting for the operation to finish
- AOM state “Done” == Good and “Pending” == ”Bad” or any other state
- Class-group == Policy, Class == Policy Sequence, just a fancy terms for the known things
- Whole policy should be reflected starting from OMP, via Forwarding Policy Manager Daemon (FPMD) to Forwarding Manager (FMAN) and then to QFP (data plane)
- FMAN-RP is just passthrough level for policy objects, hence nothing to check there, check rather FMAN-FP

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Recap: Data Policy as per FPMD view

I will use simple DIA policy for further demonstration with just single sequence:

```
cE1_BR1#show sdwan policy from-vsmart
from-vsmart data-policy VPN_1_NAT
 direction from-service
 vpn-list VPN_1
  sequence 1
   match
    source-data-prefix-list LAN
   action accept
    nat use-vpn 0
    no nat fallback
  default-action drop
from-vsmart lists vpn-list VPN_1
 vpn 1
from-vsmart lists data-prefix-list LAN
 ip-prefix 10.10.10.0/24
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

OMP policy processing troubleshooting

If FPM related output has some issues already (wrong or incomplete policy, no policy at all), then OMP policy processing to be debugged:

- Set logging marker:
  - **`set logging marker MY_DEBUG`**
- Enable debugs IOS-XE release < 17.10
  - **`debug platform software sdwan omp policy level high`**
- Enable debugs IOS-XE release >= 17.10
  - **`set platform software trace ompd R0 ompd-policy verbose`**
  - **`set platform software trace ompd R0 ompd-event verbose`**
- Reset OMP (similar to BGP hard reset, be careful!)
  - **`clear sdwan omp all`**
- Check logs and look for any errors:
  - **`show logging process ompd internal start marker MY_DEBUG`**

If no failures, config committed successfully to CDB by ConfD and transferred to FPM

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Control Plane (RP)

Check FMAN-FP policy binding to a target VRF:

```
cE1_BR1#show platform software sdwan fp active policy bind summary
Target-id  Target-Type  Dir   AF          CG-Type          Group-id  AOM-id  AOM-Status  CG-Name
------------------------------------------------------------------------------------------------------
1          VRF          IN    V4          DATA             1         116194  Done        VPN_1_NAT-VPN_1
```

Find class-id for corresponding group-id (policy), each class is a policy sequence:

```
cE1_BR1#show platform software sdwan fp active policy class summary
Group-id  Class-id  AOM-id  AOM-status  Class-Name
-------------------------------------------------------------
1         1         116195  Done        VPN_1_NAT-VPN_1-seq-1
1         65535     116192  Done        VPN_1_NAT-VPN_1-def-class
```

*Calss-id 65535 is a default-action of the policy

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Control Plane (RP)

Verify class-group (policy sequence) ID details and programming status in FMAN-FP:

```
cE1_BR1#show platform software sdwan f0 policy cg 1 detail
Policy: VPN_1_NAT-VPN_1, type: DATA, aom_id: 116191, aom_status: Done
  sequence 1
    name: VPN_1_NAT-VPN_1-seq-1, aom_id: 116195, aom_status: Done
    filters:
      match SRC OG IPV4
        value 57345
    actions: fo_aom_id: 116198, aom_status: Done
      action accept
      action nat_dia
  sequence 65535
    name: VPN_1_NAT-VPN_1-def-class, aom_id: 116192, aom_status: Done
    filters:
      match WILDCARD
    actions: fo_aom_id: 116193, aom_status: Done
      action drop
      action count
  target id: 1, dir: IN, af: V4, type: VRF, aom_id: 116194, aom_status: Done
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Control Plane (RP)

Then verify match objects programming in FMAN-FP:

```
cE1_BR1#show platform software common-classification f0 object-group all
Total Number of OGs: 1

og id        og name               og type      lkup in upd  state
-------------------------------------------------------------------------------
57345        LAN_vs                IPV4         0            PD Created

cE1_BR1#show platform software common-classification f0 object-group ipv4 57345
  OG ID: 57345
  OG TYPE: IPV4
  OG Name: LAN_vs
  Pending Entry List Size: 0
  Num LKUPs in hash: 1
  Num LKUPs in Update: 0
  AOM EPOCH: 0
  State: PD Created
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Data Plane (QFP)

Likewise, we need to verify policy in QFP (data plane).

First, ensure that feature was enabled in Features Invocation Array (FIA) for an interface:

```
cE1_BR1#show platform hardware qfp active interface if-name GigabitEthernet 4 | include SDWAN
  SDWAN_POLICY_FIA
```

*If localized data policy was enabled (ACL), you would see also SDWAN_ACL_IN/OUT in the list

If we need to verify ACL (local policy), we will also need QFP interface ID “handle”:

```
cE1_BR1#show platform hardware qfp act interface if-name GigabitEthernet4 | include QFP interface handle
  QFP interface handle: 9
```

And for data-policy or AAR policy, which is applied per-VRF basis, you need to know VRF ID which does not match to VRF name that happen to be a number (1 in this case):

```
cE1_BR1#show ip vrf detail 1  | include Id
VRF 1 (VRF Id = 3); default RD 1:1; default VPNID <not set>
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Data Plane (QFP)

Then find QFP class-group (policy) ID:

```
cE1_BR1#show platform hardware qfp active classification class-group-manager class-group client sdwan all
QFP classification class client all group

  class-group [SDWAN:1] VPN_1_NAT-VPN_1
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Using QFP class ID, dump details of a class-group (policy) match conditions:

Data Plane (QFP)

```
cE1_BR1#show platform hardware qfp active classification class-group-manager class-group client sdwan 1
class-group [sdwan-cg:1] VPN_1_NAT-VPN_1 (classes: 2)
 clients:
 fields: ipv4_og_src:1 any:1  (100000:0:0:200:0:00000000)
  (1) class: logical-expression [1.1] VPN_1_NAT-VPN_1-seq-1 (filters: 1)
         lexp: LOG-EXP: [1]
    (1) filter: generic [1.1.1] (rules: 1)
      (1) rule: generic  [1.1.1.1] (permit)
         match ipv4_og_src 57345
  (65535) class: logical-expression [1.65535] VPN_1_NAT-VPN_1-def-class (filters: 1)
         lexp: LOG-EXP: [1]
    (1) filter: generic [1.65535.1] (rules: 1)
      (1) rule: generic  [1.65535.1.1] (permit)
         match any
```

To decode individual objects like prefix-lists from the policy, use ID of the object or its name:

```
cE1_BR1#show platform hardware qfp active classification class-group-manager object-group all | include 57345
  LAN_vs:57345  Type: IPV4  No. of Entries: 1

cE1_BR1#show platform hardware qfp active classification class-group-manager object-group name LAN_vs
Object-group LAN_vs:57345  Type: IPV4  No. of Entries: 1 AOM Id: 116190
    id: 0xe0010001  10.10.10.0/255.255.255.0
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Data Plane (QFP)

Then check action statements in the whole class-group (policy) or per class (sequence) in QFP:

```
cE1_BR1#show platform hardware qfp active feature sdwan client policy class-group 1 detail
Policy: 1 type: NONE og_lkup: ipv4_src 4 ipv4_dst 0 ipv6_src 0 ipv6_dst 0 app_id 3
  sequence 1
  actions
    accept
    nat_dia
  sequence 65535
  actions
    drop
    count
  target id: 1, dir: IN, af: V4, type: VRF
cE1_BR1#show platform hardware qfp active feature sdwan client policy class-group 1 class 1
QFP sdwan client policy GroupId information


Group id            : 1
Class id            : 1

actions
  accept
  nat_dia
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Data Plane (QFP)

Based on QFP class-group (policy) ID and QFP interface handle (for ACL) or VRF ID (for AAR/Data policy), we can check TCAM programming:

```
cE1_BR1#show platform hardware qfp active classification feature-manager class-group tcam sdwan 1 ?
  acl             sdwan acl feature
  app-route       app route feature
  data-policy     data policy feature
  utd-tls-policy  UTD TLS decryption feature
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

Data Plane (QFP)

```
cE1_BR1#show platform hardware qfp active classification feature-manager class-group tcam sdwan 1 data-policy 3
proto-v4 input detail
QFP classification class group CACE

CACE classification Info::
Total entries: 2  Available entries: 65534  Total RAM used:612 bytes
IPV4 Traffic Classifier: total_entries=2 default_entry_idx=1 num_attr_clusters=2
IPV6 Traffic Classifier: total_entries=1 default_entry_idx=1 num_attr_clusters=2
MPLS Traffic Classifier: total_entries=1 default_entry_idx=1 num_attr_clusters=2
L2   Traffic Classifier: total_entries=1 default_entry_idx=1 num_attr_clusters=2

class-group [sdwan-cg:2] (classes: 2, total number of vmrs: 2)
key name: 320_Viptela_og_02  value size: 0 result size: 16 tcam id: SOFTWARE TCAM
object-group: (ipv4) lkup handle id (source: 4 dest: 0)
              (ipv6) lkup handle id (source: 0 dest: 0)

              (user) lkup handle id (appid: 0)
              (fqdn) is_valid: No version: 0
     internal (ipv6) lkup handle id (source: 0 dest: 0)
              (ext_data1) is_installed: No type: None

   Value: : 01000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00800000 00000000
    Mask: : 01000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00c00000 00000000
  Result: : 10000000 01000000 01000000 00000000

   Value: : 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
    Mask: : 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
  Result: : 02000100 00000000 ffff0000 00000000
```

Sequence 1

Default sequence 65535

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Policy pogramming verification on WAN Edge

What if you finnaly happened to (not) find some (any) problems at such a low level?

KEEP CALM
AND
CALL
CISCO TAC

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Part 2. Issues seen in the field

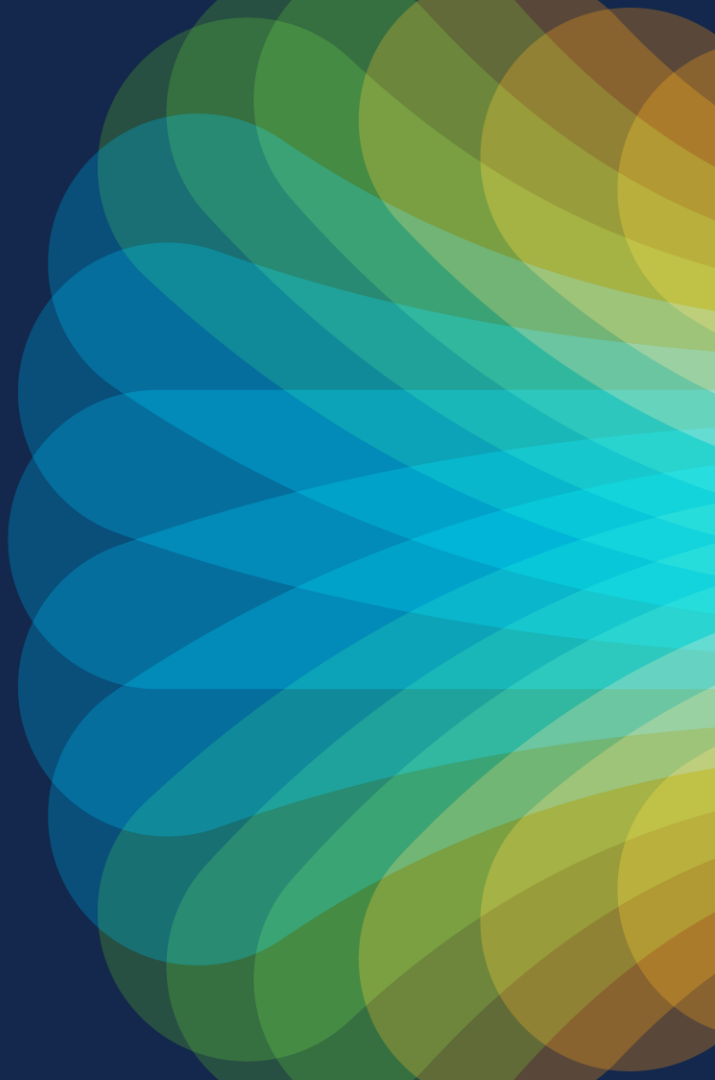

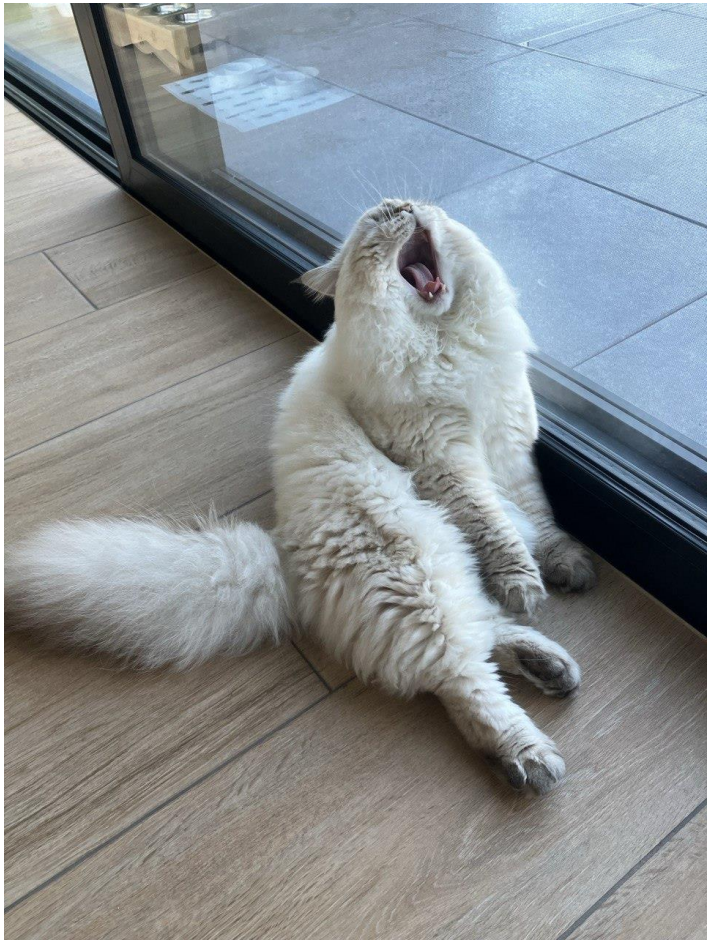

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Centralized Control Policies: Failures in overlays with disjoined underlays

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

What is an overlay with disjoined underlay?

This is an overlay with an underlay connectivity topology where different sites connected to various types of transports and transports have no direct connectivity between them

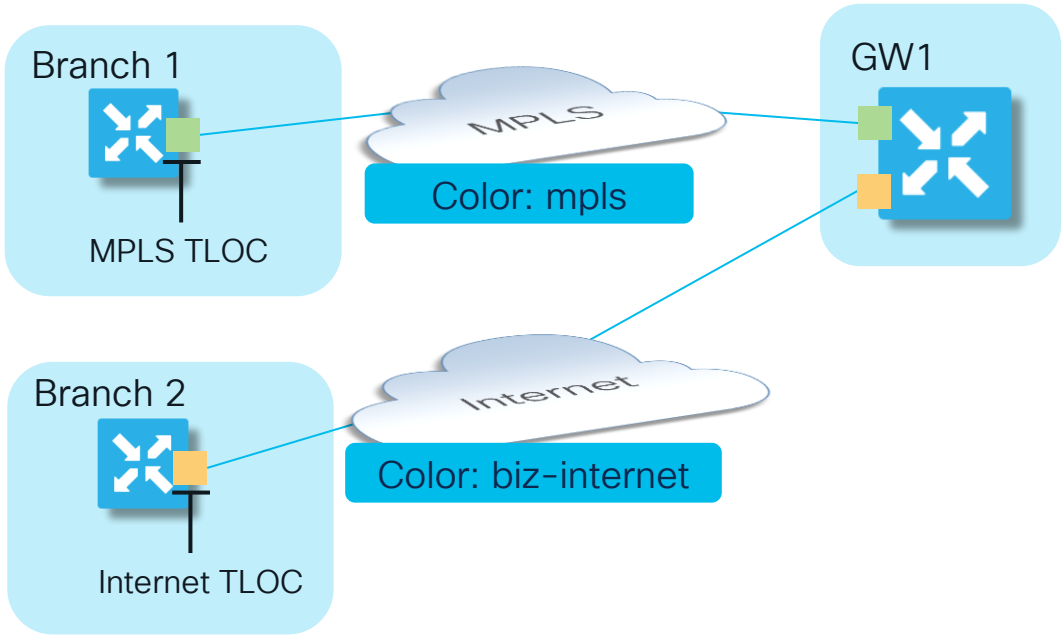

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1: Disjoined underlay without control policy

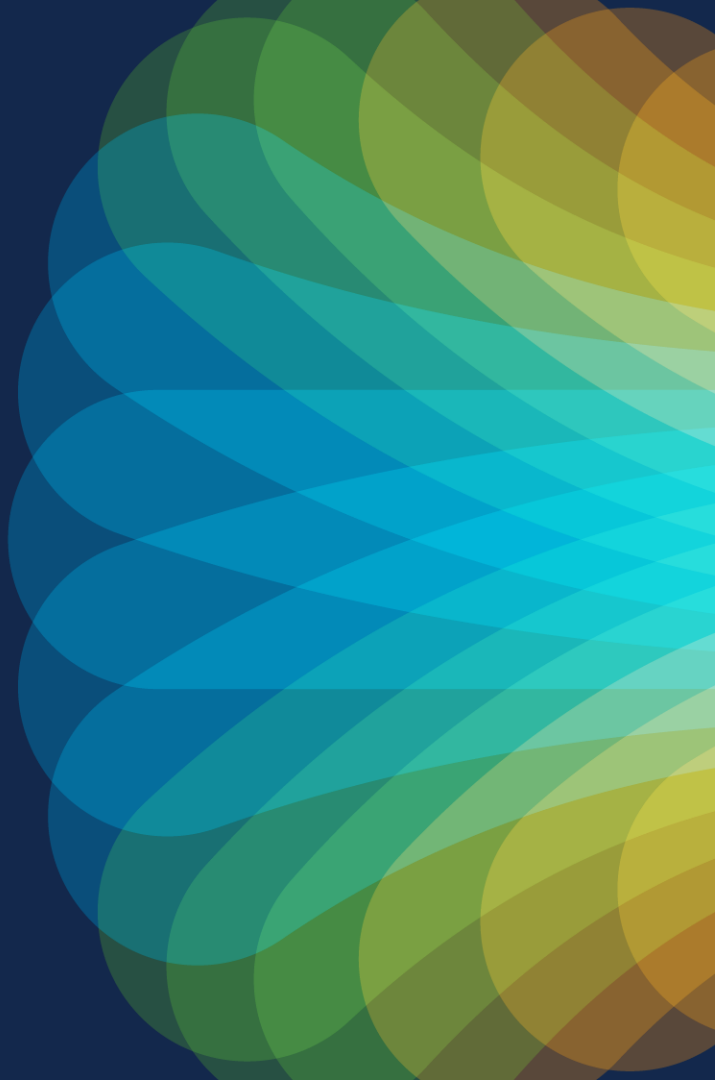

Case 1. Disjoined underlay

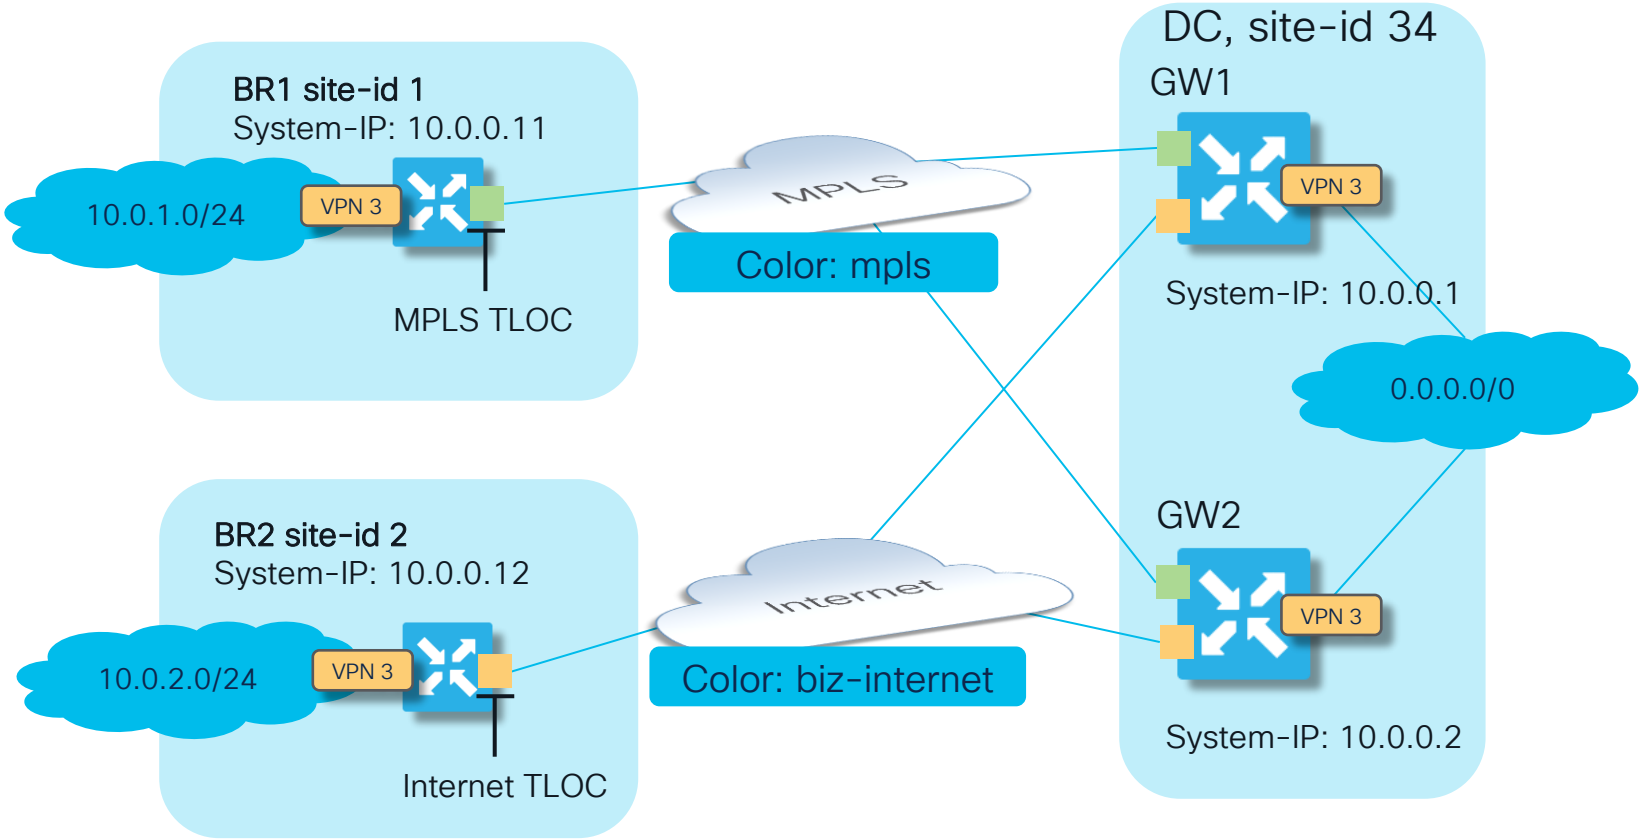

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1. Disjoined underlay

More specific route from BR2 won't be installed into RIB on BR1 and traffic will follow default route to GWs:

```
BR1#show sdwan omp routes vpn 3 10.0.2.0/24 | b PATH
                  PATH                        ATTRIBUTE
FROM PEER         ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        11     1011     Inv,U       installed   10.0.0.2         mpls             ipsec  -

BR1#show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                  PATH                        PSEUDO
FROM PEER         ID     LABEL    STATUS      KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        1011   1004     Inv,U       1       10.0.0.2         mpls             ipsec  -
10.0.0.101        1012   1004     C,I,R       1       10.0.0.2         biz-internet     ipsec  -
10.0.0.101        1071   1008     Inv,U       1       10.0.0.1         mpls             ipsec  -
10.0.0.101        1072   1008     C,I,R       1       10.0.0.1         biz-internet     ipsec  -
10.0.0.102        1355   1004     Inv,U       1       10.0.0.2         mpls             ipsec  -
10.0.0.102        1356   1004     C,R         1       10.0.0.2         biz-internet     ipsec  -
10.0.0.102        1375   1008     Inv,U       1       10.0.0.1         mpls             ipsec  -
10.0.0.102        1376   1008     C,R         1       10.0.0.1         biz-internet     ipsec  -
```

This is because there is no direct data plane tunnel formed between BR1 and BR2:

```
BR1#show sdwan omp routes vpn 3 10.0.2.0/24 | b PATH
                  PATH                        ATTRIBUTE
FROM PEER         ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        531    1011     Inv,U       installed   10.0.0.12        biz-internet     ipsec  -
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1. Disjoined underlay

Depending on EMCP hash, traffic follows default route via GW1 or GW2:

```
BR1#sh ip route vrf 3 0.0.0.0

Routing Table: 3
Routing entry for 0.0.0.0/0, supernet
  Known via "omp", distance 251, metric 0, candidate default path, type omp
  Last update from 10.0.0.12 on Sdwan-system-intf, 00:03:57 ago
  Routing Descriptor Blocks:
    10.0.0.2 (default), from 10.0.0.2, 00:03:57 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1
  * 10.0.0.1 (default), from 10.0.0.1, 00:03:57 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1

BR1#traceroute vrf 3 10.0.2.2 source 10.0.1.2 numeric
Type escape sequence to abort.
Tracing the route to 10.0.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.10.11 100 msec 1 msec 1 msec
  2 10.0.2.2 2 msec *  1 msec
BR1#traceroute vrf 3 10.0.2.2 source 10.0.1.1 numeric
Type escape sequence to abort.
Tracing the route to 10.0.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.10.12 1 msec 1 msec 0 msec
  2 10.0.2.2 1 msec *  2 msec
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1. Disjoined underlay

Failure scenario. GW2 disconnected from Internet.

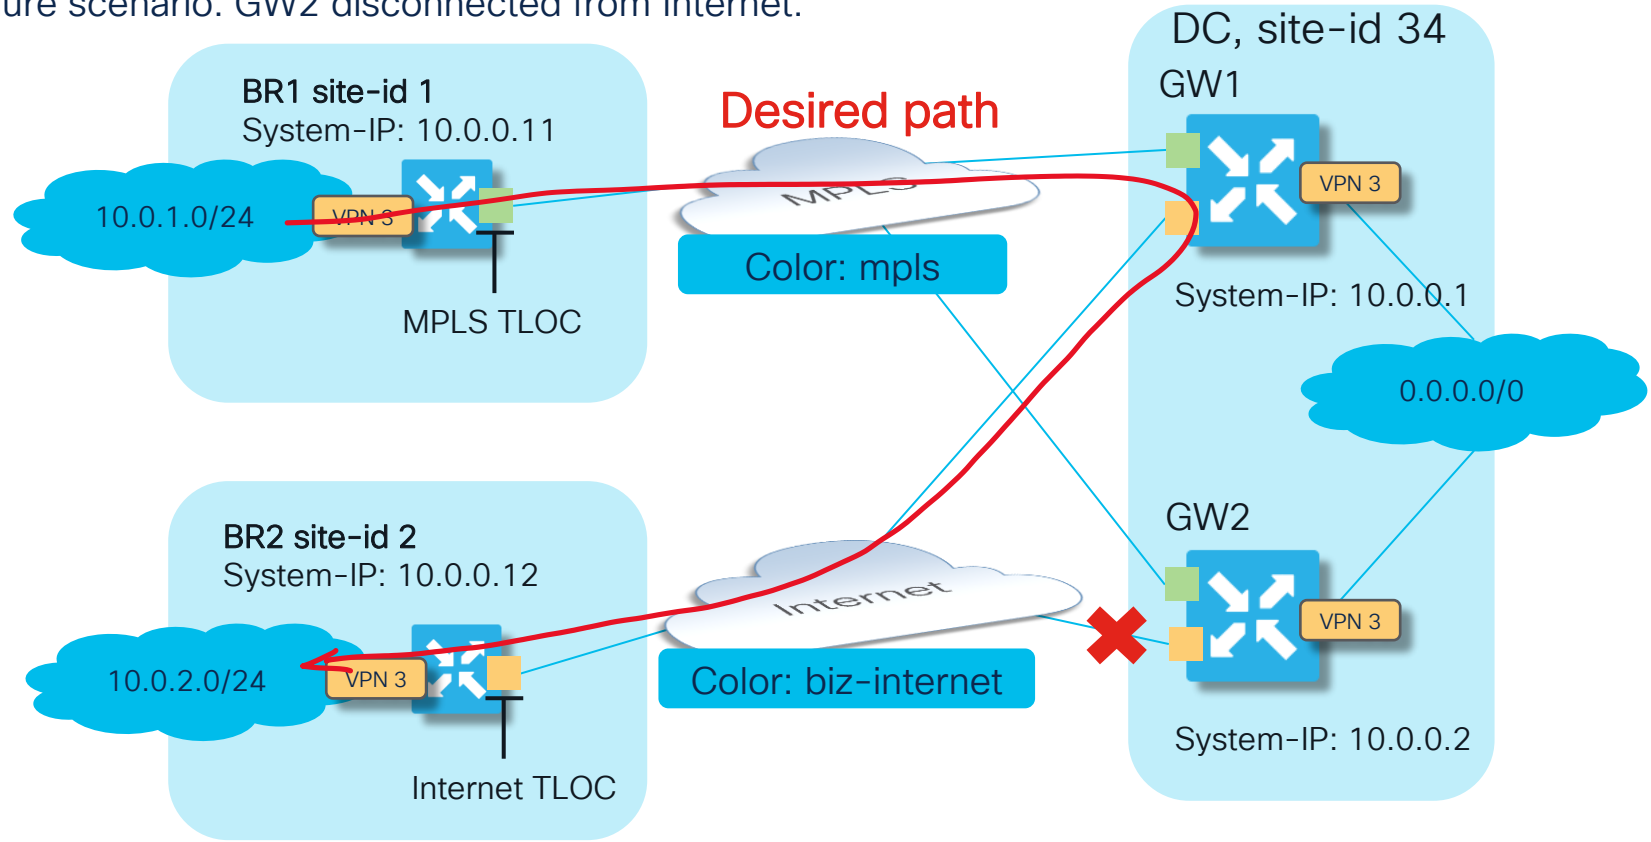

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1. Disjoined underlay

Problem: during GW2 internet failure, ~50% traffic wil be blackholed now due to ECMP:

```
BR1#traceroute vrf 3 10.0.2.2 source 10.0.1.1 numeric
Type escape sequence to abort.
Tracing the route to 10.0.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.9.12 1 msec 1 msec 0 msec
  2 192.168.9.12 !H  *  !H
BR1#traceroute vrf 3 10.0.2.2 source 10.0.1.2 numeric
Type escape sequence to abort.
Tracing the route to 10.0.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.9.11 1 msec 0 msec 1 msec
  2 10.0.2.2 2 msec *  2 msec
BR1#ping vrf 3 10.0.2.2 source 10.0.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.2.2, timeout is 2 seconds:
Packet sent with a source address of 10.0.1.1
U.U.U
Success rate is 0 percent (0/5)
BR1#ping vrf 3 10.0.2.2 source 10.0.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.2.2, timeout is 2 seconds:
Packet sent with a source address of 10.0.1.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1. Disjoined underlay.

Typical solution. Configure control policy to change TLOCs “Next-Hops” (aka hub-n-spoke)

```
policy
 control-policy CHANGE_TLOC_NH
  sequence 10
   match route
    site-list BR1
    vpn       3
   !
   action accept
    set
     tloc-list INET_TLOCS
    !
   !
  !
  sequence 20
   match route
    site-list BR2
    vpn       3
   !
   action accept
    set
     tloc-list MPLS_TLOCS
    !
   !
  !
  default-action accept
 !
!
```

```
policy
 lists
  site-list ALL_BRANCHES
   site-id 1
   site-id 2
  !
  site-list BR1
   site-id 1
  !
  site-list BR2
   site-id 2
  !
  tloc-list INET_TLOCS
   tloc 10.0.0.1 color biz-internet encap ipsec
   tloc 10.0.0.2 color biz-internet encap ipsec
  !
  tloc-list MPLS_TLOCS
   tloc 10.0.0.1 color mpls encap ipsec
   tloc 10.0.0.2 color mpls encap ipsec
  !
 !
!
apply-policy
 site-list ALL_BRANCHES
  control-policy CHANGE_TLOC_NH out
 !
!
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 1. Disjoined underlay

Typical solution - testing

Once policy applied, TLOC rewrite happens to GW TLOCs:

```
BR1#show sdwan omp routes vpn 3 10.0.2.0/24 | b PATH
                  PATH                        ATTRIBUTE
FROM PEER         ID     LABEL    STATUS      TYPE       TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        2022   1008     C,I,R       installed  10.0.0.1         mpls             ipsec  -
10.0.0.101        2023   1004     C,I,R       installed  10.0.0.2         mpls             ipsec  -
```

Traffic follows specific path to BR2 subnet:

```
BR1#sh ip route vrf 3 10.0.2.0

Routing Table: 3
Routing entry for 10.0.2.0/24
  Known via "omp", distance 251, metric 0, type omp
  Last update from 10.0.0.2 on Sdwan-system-intf, 00:03:00 ago
  Routing Descriptor Blocks:
    10.0.0.2 (default), from 10.0.0.2, 00:03:00 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1
  * 10.0.0.1 (default), from 10.0.0.1, 00:03:00 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1
```

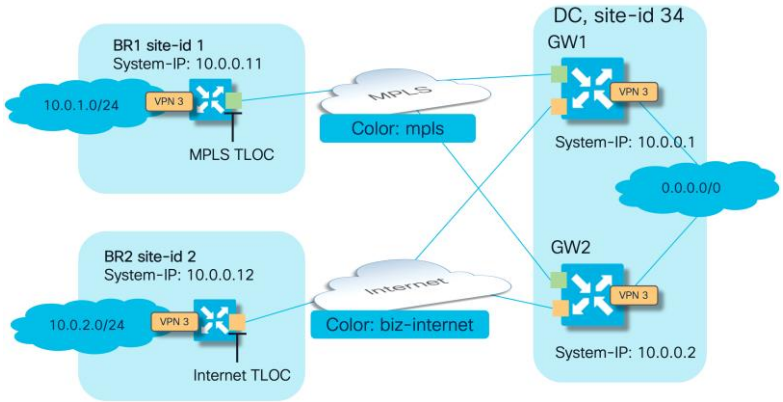

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 2: Active-standby redundancy failure with disjoined underlay

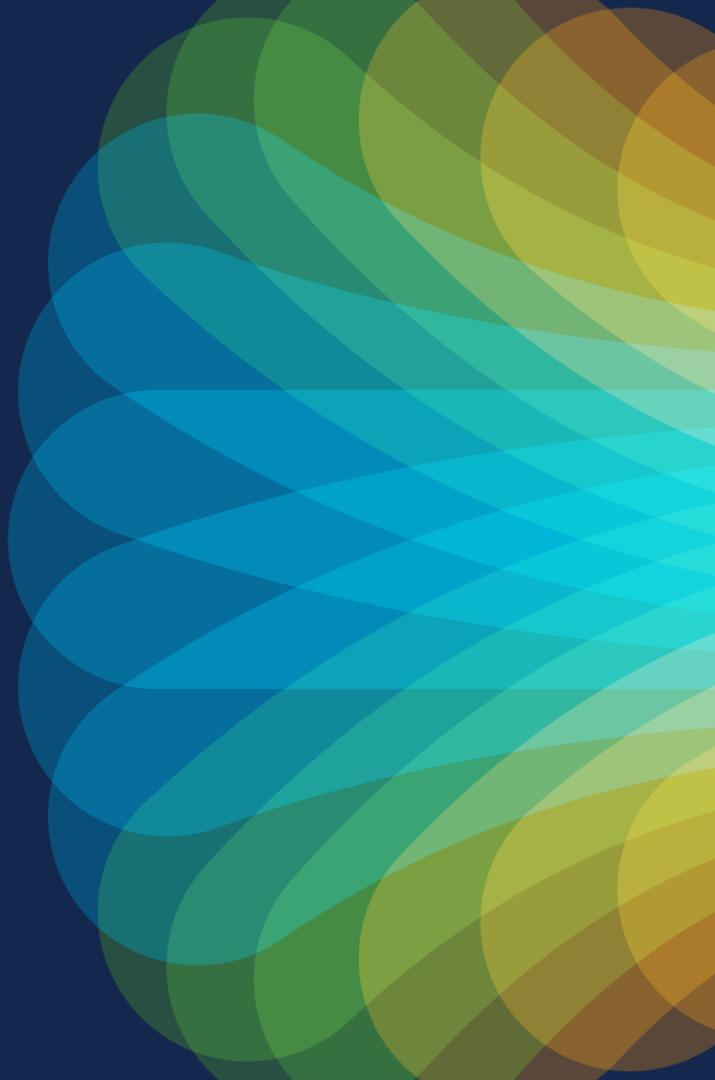

Case 2. Active-standby redundancy failure with disjoint underlay.

Failure to influence path with OMP route preference.

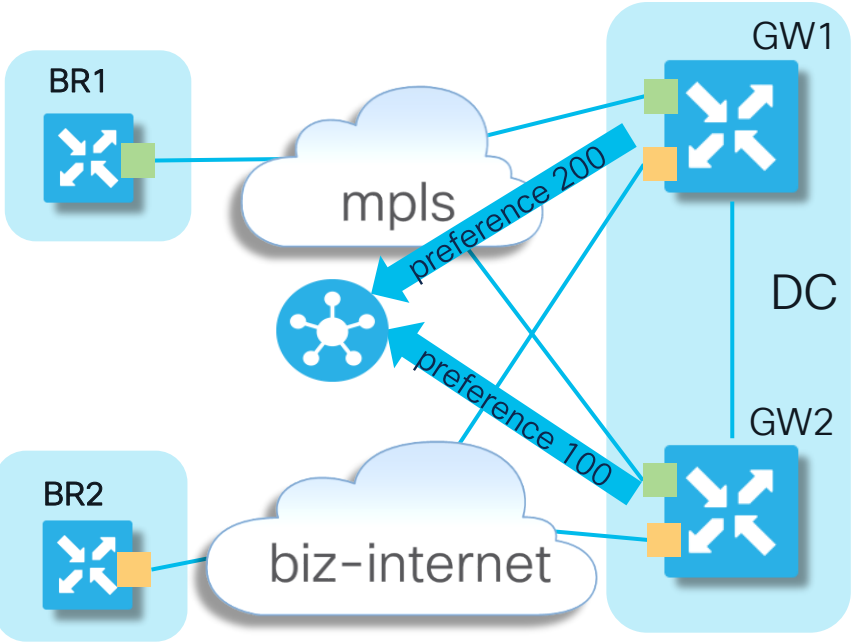

- Main objective is to ensure preferred path to DC subnets is via GW1
- OMP route preference set to 200 for routes advertised by GW1 with help of vSmart inbound policy (or, for example, instead of preference, routing protocol metric is being used to influence paths)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 2. Active-standby redundancy failure with disjoint underlay.

Original centralized control policy on vSmart:

```
policy
 lists
  site-list GW1
   site-id 1
  !
 !
!
 control-policy PREFER_GW1
  sequence 10
   match route
    site-id 1
   !
   action accept
    set
     preference 200
    !
   !
  !
  default-action accept
 !
!
apply-policy
 site-list GW1
  control-policy PREFER_GW1 in
 !
!
```

Can you see potential problem here?

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 2. Active-standby redundancy failure with disjoint underlay.

Problem.

BR1 prefers GW1 as a result of the policy:

```
BR1# show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                  PATH                      PSEUDO
FROM PEER         ID      LABEL    STATUS   KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------
10.0.0.101        1066    1008     C,I,R    1       10.0.0.1         mpls             ipsec  200
10.0.0.101        1067    1008     Inv,U    1       10.0.0.1         biz-internet     ipsec  200
10.0.0.102        2142    1008     C,R      1       10.0.0.1         mpls             ipsec  200
10.0.0.102        2143    1008     Inv,U    1       10.0.0.1         biz-internet     ipsec  200
```

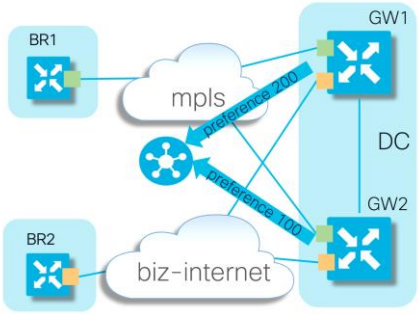

But in case of mpls link failure on GW1, there are no more valid paths remain on BR1:

```
BR1# show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                  PATH                      PSEUDO
FROM PEER         ID      LABEL    STATUS   KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------
10.0.0.101        1067    1008     Inv,U    1       10.0.0.1         biz-internet     ipsec  200
10.0.0.102        2143    1008     Inv,U    1       10.0.0.1         biz-internet     ipsec  200
```

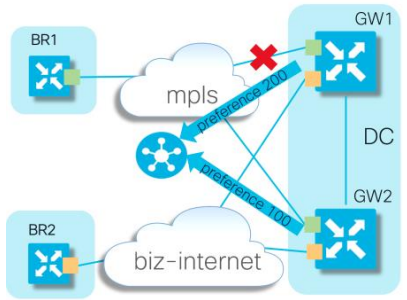

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 2. Active-standby redundancy failure with disjoint underlay.

Why problem happens here?

And BR1 can't resolve path via internet because it is connected to mpls color ony

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 2. Typical Solution.

Solution is to influence preference only **after** best path selection (i.e. outbound control policy to set preference)

```
policy
 lists
  site-list ALL_BRANCHES
   site-id 1
   site-id 2
  !
 !
!
control-policy PREFER_BR1
  sequence 10
   match route
    site-id 1
   !
   action accept
    set
     preference 200
    !
   !
  !
  default-action accept
 !
!
apply-policy
 site-list ALL_BRANCHES
  control-policy PREFER_BR1 out
 !
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 2. Typical Solution (cont.)

Testing outbound (to branches) control policy configured on vSmart to ensure branches prefer GW1.

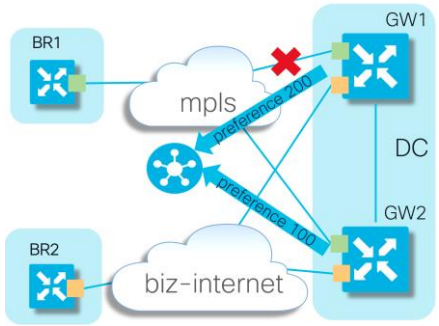

Normal pre-failure conditions:

```
BR1# show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                  PATH                        PSEUDO
FROM PEER         ID      LABEL    STATUS     KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        21      1008     C,I,R      1       10.0.0.1         mpls             ipsec  200
10.0.0.101        22      1008     Inv,U      1       10.0.0.1         biz-internet     ipsec  200
10.0.0.101        65      1004     R          1       10.0.0.2         mpls             ipsec  -
10.0.0.101        66      1004     Inv,U      1       10.0.0.2         biz-internet     ipsec  -
```

Testing failover scenario. GW1 lost MPLS link, but BR1 successfully installs backup path via GW2

```
cE1_BR1# show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                  PATH                        PSEUDO
FROM PEER         ID      LABEL    STATUS     KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        22      1008     Inv,U      1       10.0.0.1         biz-internet     ipsec  200
10.0.0.101        65      1004     C,I,R      1       10.0.0.2         mpls             ipsec  -
10.0.0.101        66      1004     Inv,U      1       10.0.0.2         biz-internet     ipsec  -
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

(Not so Well) Known Failures with Centralized Control Policies

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with “set tloc-list”

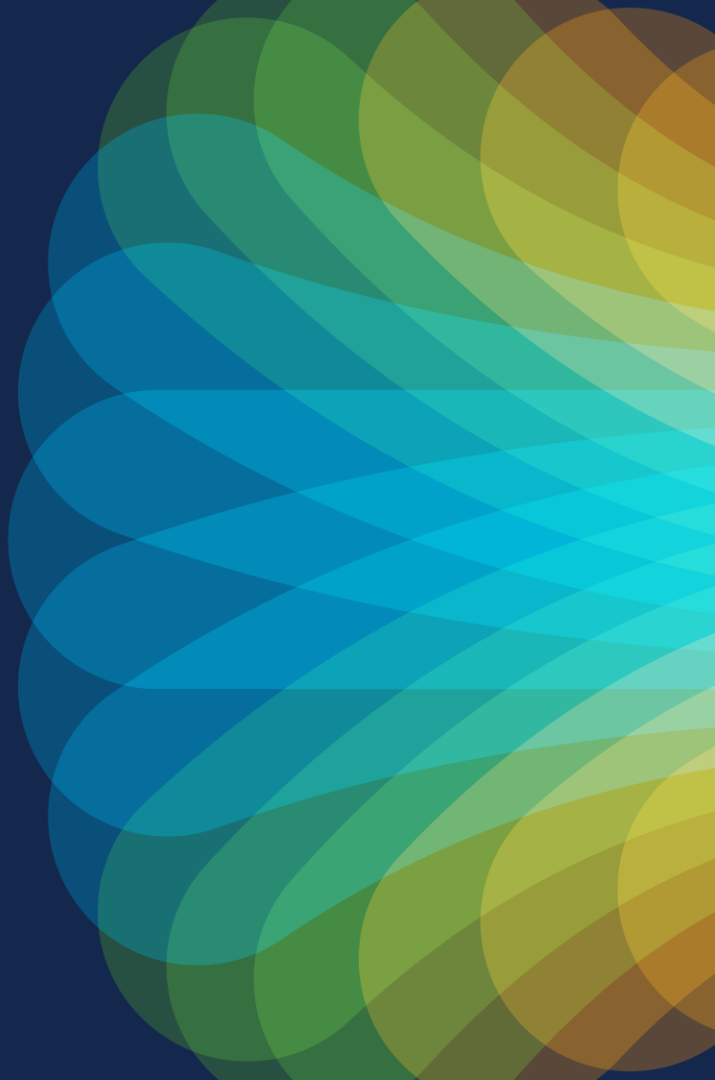

Case 3: Multi-level backup preference with ”set tloc-list”

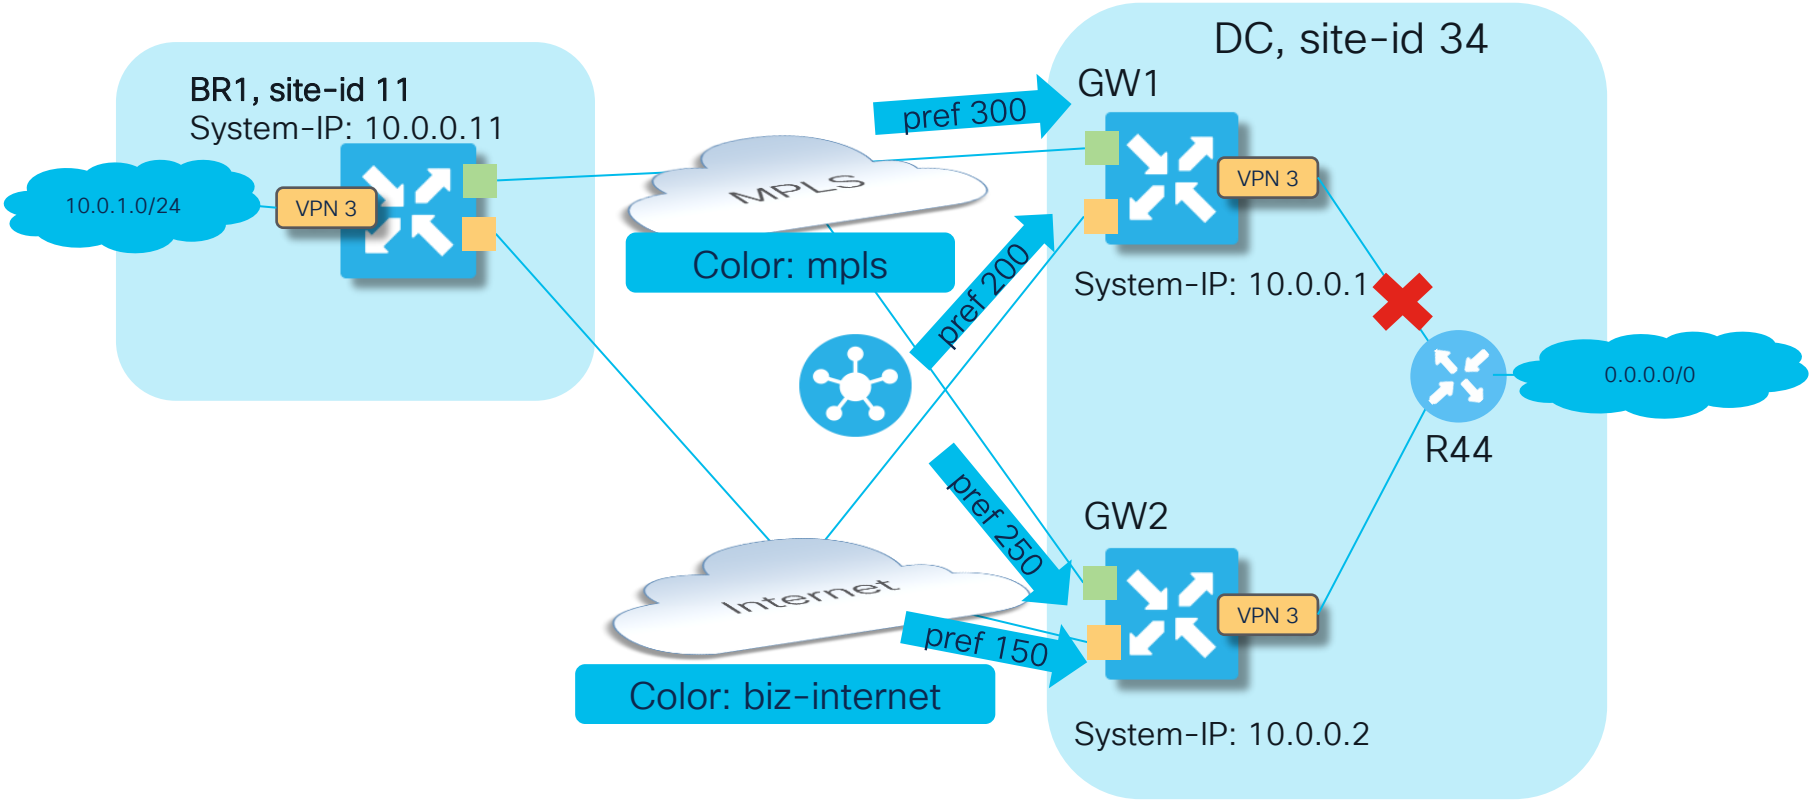

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Centralized control policy on vSmart:

```
policy
 lists
  tloc-list DC_TLOCS_W_PREF
   tloc 10.0.0.1 color mpls encap ipsec preference 300
   tloc 10.0.0.1 color biz-internet encap ipsec preference 200
   tloc 10.0.0.2 color mpls encap ipsec preference 250
   tloc 10.0.0.2 color biz-internet encap ipsec preference 150
  !
 lists
  site-list DCs
   site-id 34
  !
  site-list ALL_BRANCHES
   site-id 11-12
  !
 !
!
```

```
control-policy DC_PREFERENCES
  sequence 10
   match route
    site-list DCs
   !
   action accept
    set
     tloc-list DC_TLOCS_W_PREF
    !
   !
  !
  default-action accept
 !
!
apply-policy
 site-list ALL_BRANCHES
  control-policy DC_PREFERENCES out
 !
!
```

Can you see any problems here?

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Check routing and policy under normal conditions:

```
BR1#show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                PATH                        ATTRIBUTE
FROM PEER       ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101      1146   1008     C,I,R       installed   10.0.0.1         mpls             ipsec  300
10.0.0.101      1147   1008     R           installed   10.0.0.1         biz-internet     ipsec  200
10.0.0.101      1148   1004     R           installed   10.0.0.2         mpls             ipsec  250
10.0.0.101      1149   1004     R           installed   10.0.0.2         biz-internet     ipsec  150

BR1#show sdwan policy service-path vpn 3 interface Loopback 3 source-ip 10.0.1.1 dest-ip 10.10.10.10 protocol 6 all
Number of possible next hops: 1
Next Hop: IPsec
  Source: 192.168.9.11 12366 Destination: 192.168.9.13 12426 Local Color: mpls Remote Color: mpls Remote System IP:
10.0.0.1

BR1#show ip route vrf 3 10.10.10.10 resolve

Routing Table: 3
Routing entry for 0.0.0.0/0
  Known via "omp", distance 251, metric 0, candidate default path, type omp
  Last update from 10.0.0.1 on Sdwan-system-intf, 00:02:39 ago
  Routing Descriptor Blocks:
  * 10.0.0.1 (default), from 10.0.0.1, 00:02:39 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1
```

GW1 is preferred and it is the only path to destination

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Failover testing: GW1 disconnected from the service-side (LAN) segment:

```
GW1#sh ip cef vrf 3 10.10.10.10
0.0.0.0/0
  nexthop 10.0.34.44 GigabitEthernet4
GW1#config-t

admin connected from 127.0.0.1 using console on cE3_GW1
GW1(config)#
GW1(config)# interface GigabitEthernet4
GW1(config-if)# shutdown
GW1(config-if)# commit
Commit complete.
GW1(config-if)#end
GW1#sh ip cef vrf 3 10.10.10.10
0.0.0.0/0
  no route
GW1# sh ip ro vrf 3 10.10.10.0

Routing Table: 3
% Subnet not in table
```

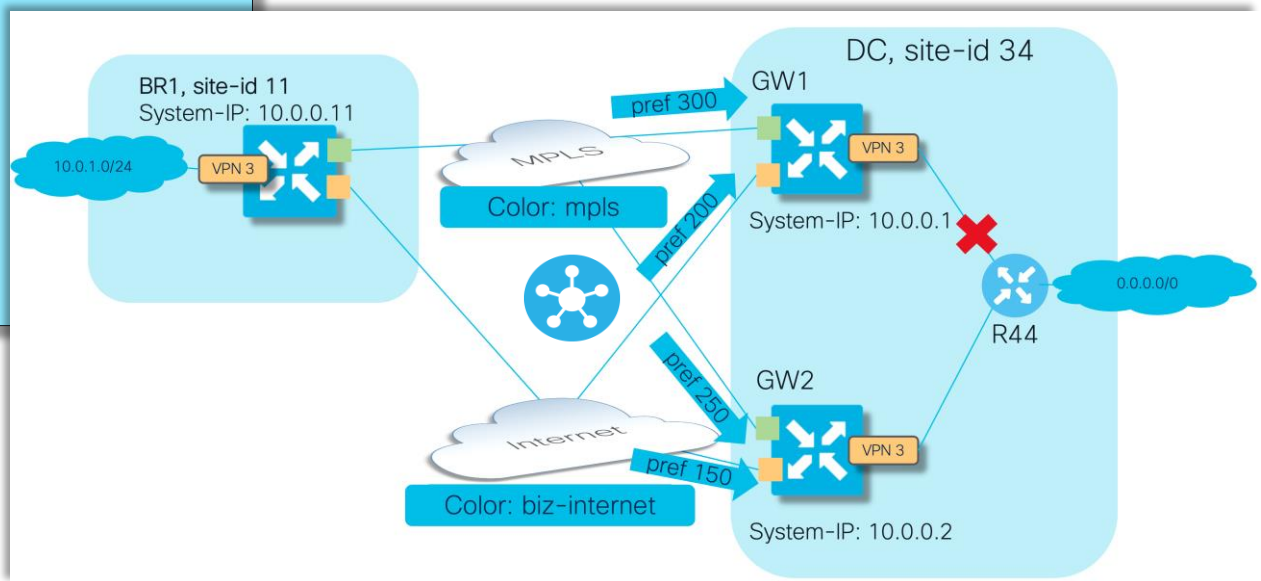

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Failover testing (cont.)

Despite that only GW2 now advertises default route and GW1 route was withdrawn from vSmart…

```
vsmart1# show omp routes vpn 3 0.0.0.0/0 received | tab | b PATH
                  PATH                        ATTRIBUTE
FROM PEER         ID     LABEL    STATUS      TYPE        TLOC IP          COLOR            ENCAP  PREFERENCE
-------------------------------------------------------------------------------------------------------------
10.0.0.2          66     1004     C,R         installed   10.0.0.2         mpls             ipsec  -
10.0.0.2          68     1004     C,R         installed   10.0.0.2         biz-internet     ipsec  -
10.0.0.102        2408   1004     C,R         installed   10.0.0.2         mpls             ipsec  -
10.0.0.102        2409   1004     C,R         installed   10.0.0.2         biz-internet     ipsec  -
```

… somehow BR1 still selects GW1 as a preferred path:

```
BR1#show sdwan policy service-path vpn 3 interface Loopback 3 source-ip 10.0.1.1 dest-ip 10.10.10.44 protocol 6 all
Number of possible next hops: 1
Next Hop: IPsec
  Source: 192.168.9.11 12366 Destination: 192.168.9.13 12426 Local Color: mpls Remote Color: mpls Remote System IP:
10.0.0.1

BR1#show ip route vrf 3 10.10.10.10 resolve

Routing Table: 3
Routing entry for 0.0.0.0/0
  Known via "omp", distance 251, metric 0, candidate default path, type omp
  Last update from 10.0.0.1 on Sdwan-system-intf, 00:11:27 ago
  Routing Descriptor Blocks:
  * 10.0.0.1 (default), from 10.0.0.1, 00:11:27 ago, via Sdwan-system-intf
      Route metric is 0, traffic share count is 1
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Failover testing (cont.)

Note that GW1 MPLS TLOC is still preferred, but order of paths has changed (*hint!*)

```
BR1#show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                PATH                      PSEUDO
FROM PEER       ID     LABEL    STATUS    KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------
10.0.0.101      12     1004     R         1       10.0.0.2         mpls             ipsec  250
10.0.0.101      13     1004     R         1       10.0.0.2         biz-internet     ipsec  150
10.0.0.101      31     1008     C,I,R     1       10.0.0.1         mpls             ipsec  300
10.0.0.101      32     1008     R         1       10.0.0.1         biz-internet     ipsec  200
```

Certainly, it leads to BR1 traffic blackholing because GW1 has no reachability to LAN anymore:

```
BR1#ping vrf 3 10.10.10.44
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.44, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Why? This is because vSmart still executes policy as instructed and sets route preference + TLOC:

```
vsmart1# show omp routes vpn 3 0.0.0.0/0 advertised detail | nomore | exclude not\ set | begin 10.0.0.11
```

```
peer    10.0.0.11
    Attributes:
     originator       10.0.0.2
     label            1004
     path-id          12
     tloc             10.0.0.2, mpls, ipsec
     site-id          34
     overlay-id        1
     preference       250
     origin-proto     static
     origin-metric    0
    Attributes:
     originator       10.0.0.2
     label            1004
     path-id          13
     tloc             10.0.0.2, biz-internet, ipsec
     site-id          34
     overlay-id        1
     preference       150
     origin-proto     static
     origin-metric    0
    Attributes:
     originator       10.0.0.2
     label            1008
     path-id          31
     tloc             10.0.0.1, mpls, ipsec
     site-id          34
     overlay-id        1
     preference       300
     origin-proto     static
     origin-metric    0
    Attributes:
     originator       10.0.0.2
     label            1008
     path-id          32
     tloc             10.0.0.1, biz-internet, ipsec
     site-id          34
     overlay-id        1
     preference       200
     origin-proto     static
     origin-metric    0
```

*Note that originator is always 10.0.0.2 (GW2)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Recap original control policy.

```
policy
 lists
  tloc-list DC_TLOCS_W_PREF
   tloc 10.0.0.1 color mpls encap ipsec preference 300
   tloc 10.0.0.1 color biz-internet encap ipsec preference 200
   tloc 10.0.0.2 color mpls encap ipsec preference 250
   tloc 10.0.0.2 color biz-internet encap ipsec preference 150
  !
 lists
  site-list DCs
   site-id 34
  !
  site-list ALL_BRANCHES
   site-id 11-12
  !
 !
!
```

```
control-policy DC_PREFERENCES
  sequence 10
   match route
    site-list DCs
   !
   action accept
    set
     tloc-list DC_TLOCS_W_PREF
    !
   !
  !
  default-action accept
 !
!
apply-policy
 site-list ALL_BRANCHES
  control-policy DC_PREFERENCES out
 !
!
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Why problem happened?

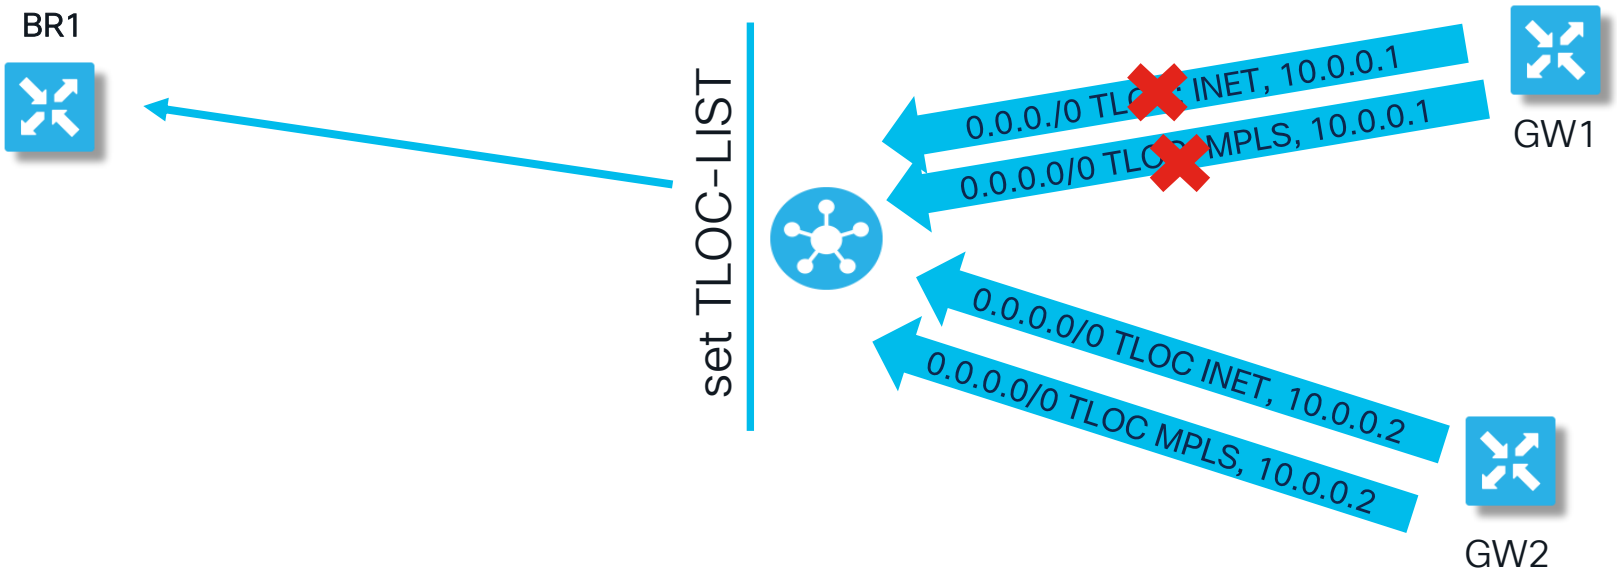

GW1 does not advertise default routes anymore, but vSmart still unconditionally rewrite route’s TLOC with list that contains GW1 TLOCs as instructed by the policy, this is expected behavior.

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

Typical solution. Set TLOC preference conditionally and only if received route has corresponding TLOC :

```
policy
 lists
  tloc-list GW1_TLOCS
   tloc 10.0.0.1 color mpls encap ipsec
   tloc 10.0.0.1 color biz-internet encap ipsec
  !
  tloc-list GW1_TLOCS_W_PREF
   tloc 10.0.0.1 color mpls encap ipsec preference 300
   tloc 10.0.0.1 color biz-internet encap ipsec preference 200
  !
  tloc-list GW2_TLOCS
   tloc 10.0.0.2 color mpls encap ipsec
   tloc 10.0.0.2 color biz-internet encap ipsec
  !
  tloc-list GW2_TLOCS_W_PREF
   tloc 10.0.0.2 color mpls encap ipsec preference 250
   tloc 10.0.0.2 color biz-internet encap ipsec preference 150
  !
```

```
apply-policy
 site-list ALL_BRANCHES
  control-policy DC_PREFERENCES_FIX out
 !
!
```

```
control-policy DC_PREFERENCES_FIX
 sequence 10
  match route
   site-list DCs
   tloc-list GW1_TLOCS
  !
  action accept
   set
    tloc-list GW1_TLOCS_W_PREF
   !
  !
 !
 sequence 20
  match route
   site-list DCs
   tloc-list GW2_TLOCS
  !
  action accept
   set
    tloc-list GW2_TLOCS_W_PREF
   !
  !
 !
 default-action accept
!
!
```

* Unlike some other available solutions, this is the best one because it won’t lead to suboptimal routing

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 3: Multi-level backup preference with ”set tloc-list” (cont.)

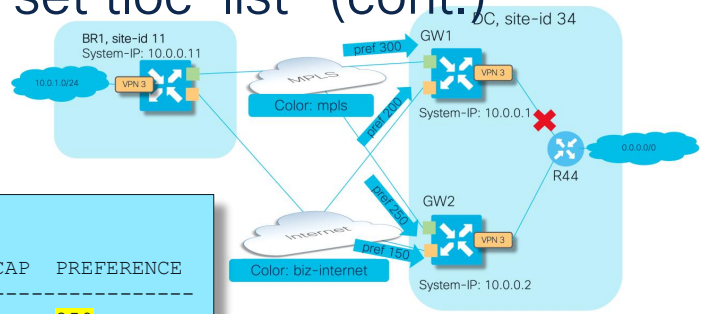

Solution testing when GW1 has the LAN link failure

```
BR1#show sdwan omp routes vpn 3 0.0.0.0/0 | b PATH
                  PATH                        PSEUDO
FROM PEER         ID      LABEL    STATUS     KEY     TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------
10.0.0.101        1166    1004     C,I,R      1       10.0.0.2         mpls             ipsec  250
10.0.0.101        1167    1004     R          1       10.0.0.2         biz-internet     ipsec  150
```

Note that there are only 2 paths remain and GW2 MPLS is preferred:

```
BR1#show sdwan policy service-path vpn 3 interface Loopback 3 source-ip 10.0.1.1 dest-ip 10.10.10.44 protocol 6 all
Number of possible next hops: 1
Next Hop: IPsec
  Source: 192.168.9.11 12366 Destination: 192.168.9.14 12406 Local Color: mpls Remote Color: mpls Remote System IP:
10.0.0.2


BR1#ping vrf 3 10.10.10.44
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.44, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Failover works as expected

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4. Traffic engineering with “set tloc-action"

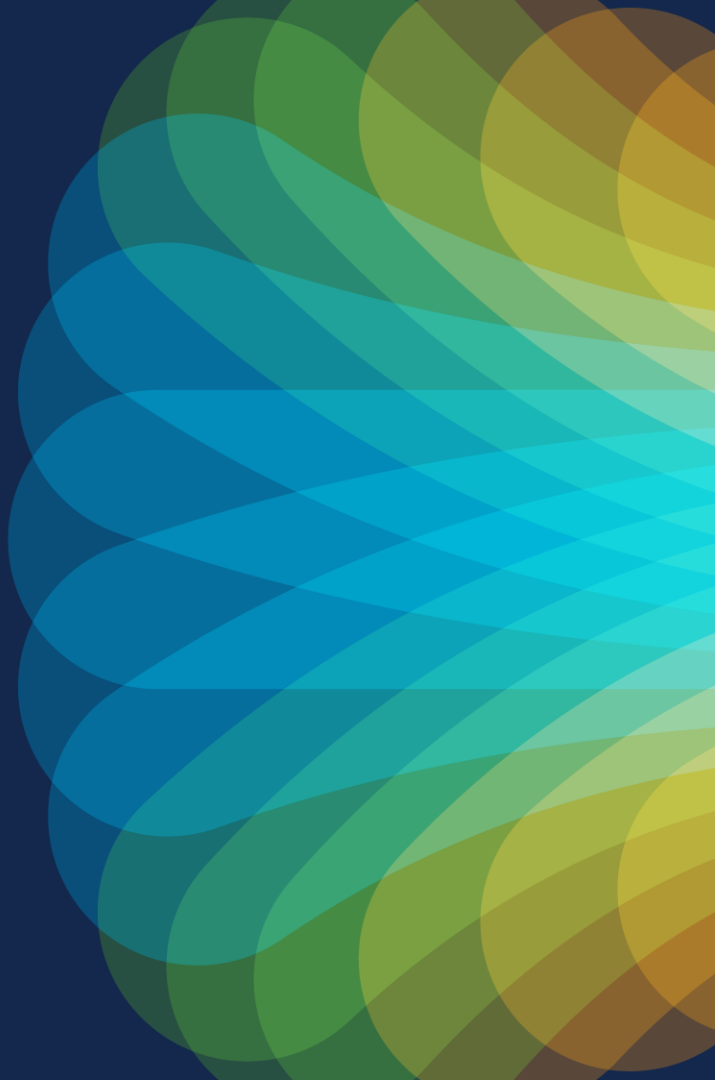

Case 4. Traffic engineering with “set tloc-action"

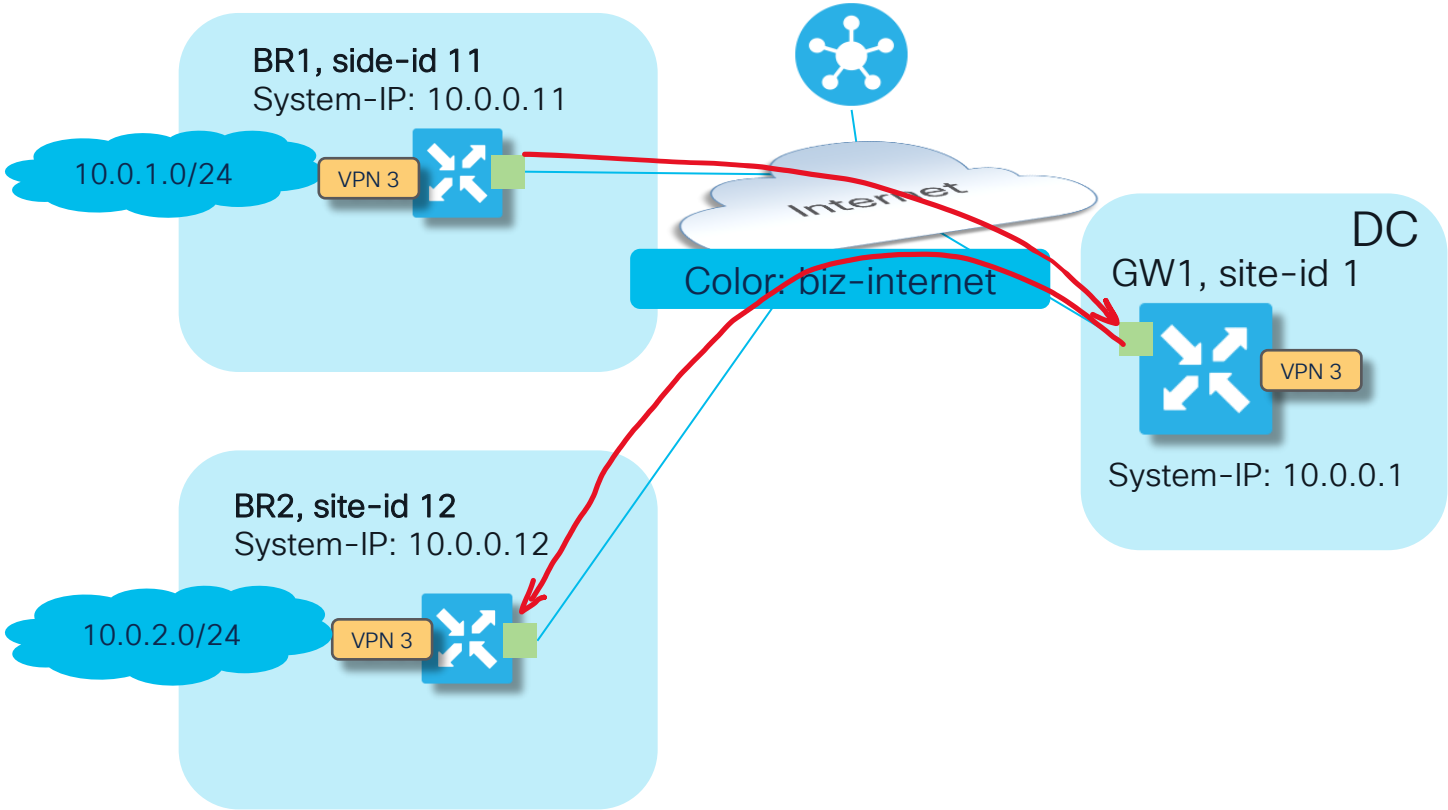

Aim here to steer traffic from BR1 to BR2 via GW1

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4. Traffic engineering with “set tloc-action"

vSmart policy to enforce traffic path via GW1:

```
policy
 lists
  site-list ALL_BRANCHES
   site-id 11-12
 !
 control-policy REDIRECT_VIA_GW1
  sequence 10
   match route
    site-list ALL_BRANCHES
   !
   action accept
    set
     tloc-action primary
     tloc 10.0.0.1 color biz-internet encap ipsec
    !
   !
  !
  default-action accept
 !
!
apply-policy
 site-list ALL_BRANCHES
  control-policy REDIRECT_VIA_GW1 out
 !
!
```

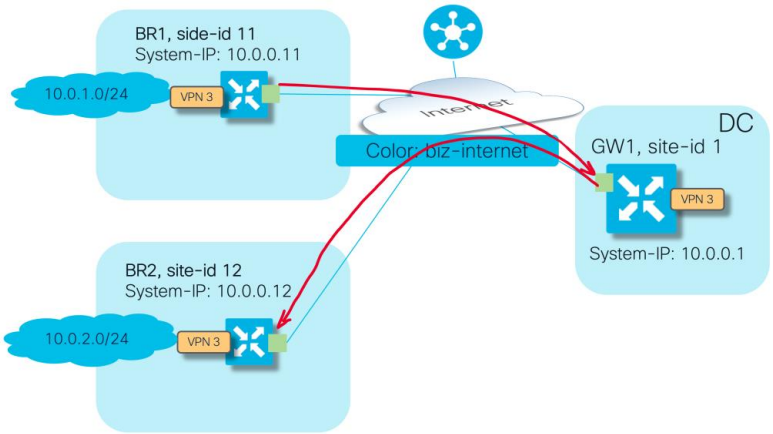

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4. Traffic engineering with “set tloc-action"

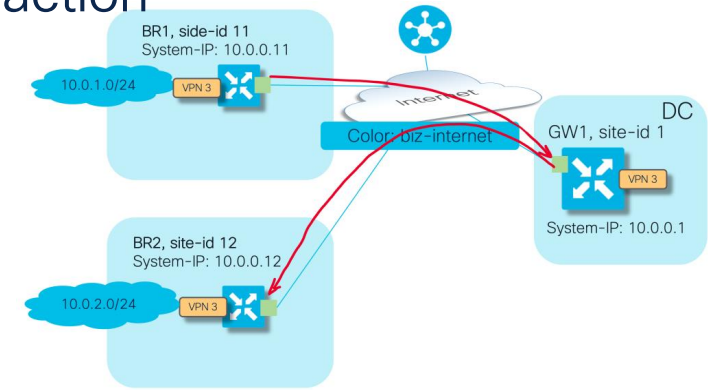

Testing and the problem.

Why traffic follows direct path (no intermediate hops there)?

```
BR1#traceroute vrf 3 10.0.2.2
Type escape sequence to abort.
Tracing the route to 10.0.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 10.0.2.2 1 msec *  2 msec
```

Why is there only one path available which is directly to BR2?

```
BR1#show sdwan policy service-path vpn 3 interface Loopback 3 source-ip 10.0.1.1 dest-ip 10.0.2.2 protocol 6 all
Number of possible next hops: 1
Next Hop: IPsec
  Source: 192.168.10.11 12366 Destination: 192.168.10.12 12366 Local Color: biz-internet Remote Color: biz-internet Remote
System IP: 10.0.0.12
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4. Traffic engineering with “set tloc-action"

Let’s check OMP routes on BR1 (unimportant attributes excluded)

```
cE1_BR1#show sdwan omp routes 10.0.2.0/24 detail | exclude not\ set|metric|overlay|label|path-id|origin
----------------------------------------------------
omp route entries for vpn 3 route 10.0.2.0/24
----------------------------------------------------
            RECEIVED FROM:
peer            10.0.0.101
status          C,I,R
    Attributes:
     originator       10.0.0.12
     type             installed
     tloc             10.0.0.12, biz-internet, ipsec
     site-id          12
            RECEIVED FROM:
peer            10.0.0.101
status          Inv,U
loss-reason     invalid
lost-to-peer    10.0.0.101
    Attributes:
     originator       10.0.0.12
     type             installed
     tloc             10.0.0.1, biz-internet, ipsec
     ultimate-tloc    10.0.0.12, biz-internet, ipsec -- primary
     site-id          12
```

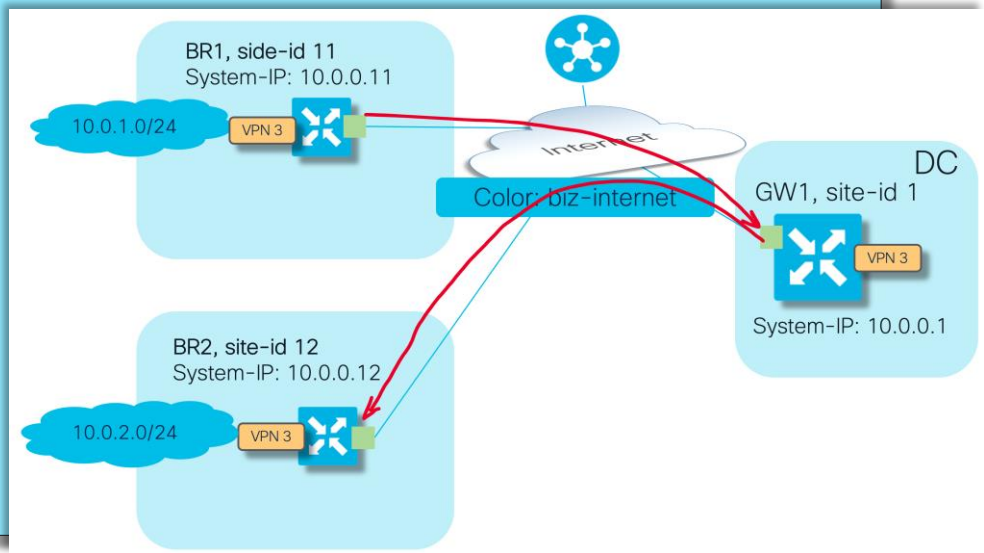

- Note that second “traffic-engineering” path via GW1 is invalid and unresolved. It also has something called **`ultimate-tloc`**
- An **`ultimate-tloc`** is the TLOC to which the intermediate hop (GW1 in this case) builds data plane tunnel (IPsec or GRE) in order to get to the final (ultimate) destination (BR2)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4. Traffic engineering with “set tloc-action"

This is an example of a few ”must know” slutions. Caused by misconfiguration (or rather lack of required config)

- If the action is **`set tloc-action`**, you must configure “**`service TE`**” in the global VRF on the intermediate router

```
GW1#sh sdwan running-config "sdwan service"
sdwan
 service TE vrf global
  !
!

GW1#show sdwan omp services | include TE
GW1#
```

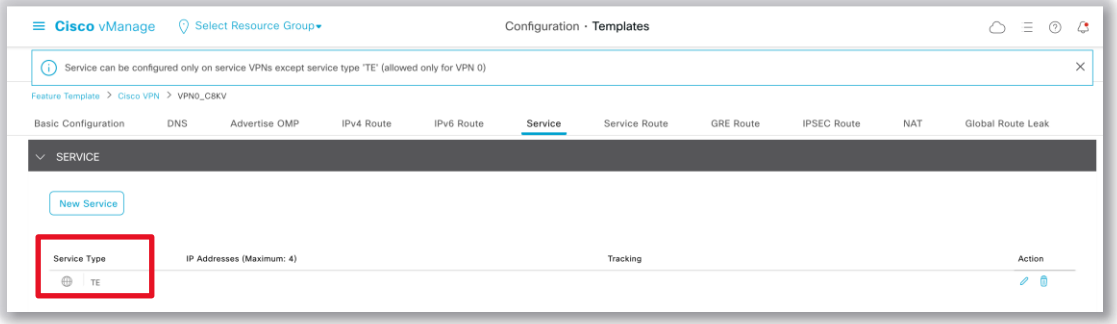

- It can not be seen with **`show sdwan omp services`** command
- By the way, this is also a pre-requisite for dynamic on-demand tunnels (ODT) to work properly

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4. Traffic engineering with “set tloc-action"

Solution testing. Route with ultimate-tloc now selected and traffic goes via GW1 as desired:

```
BR1#show sdwan omp routes 10.0.2.0/24 detail | exclude not\ set|metric|overlay|label|path-id|origin-
----------------------------------------------------
omp route entries for vpn 3 route 10.0.2.0/24
----------------------------------------------------
            RECEIVED FROM:
peer            10.0.0.101
status          R
loss-reason     tloc-action
lost-to-peer    10.0.0.101
    Attributes:
     originator       10.0.0.12
     type             installed
     tloc             10.0.0.12, biz-internet, ipsec
     site-id          12
            RECEIVED FROM:
peer            10.0.0.101
status          C,I,R
    Attributes:
     originator       10.0.0.12
     type             installed
     tloc             10.0.0.1, biz-internet, ipsec
     ultimate-tloc    10.0.0.12, biz-internet, ipsec -- primary
     site-id          12

BR1#traceroute vrf 3 10.0.2.2
Type escape sequence to abort.
Tracing the route to 10.0.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.10.13 1 msec 0 msec 1 msec
  2 10.0.2.2 1 msec *  2 msec
```

GW1

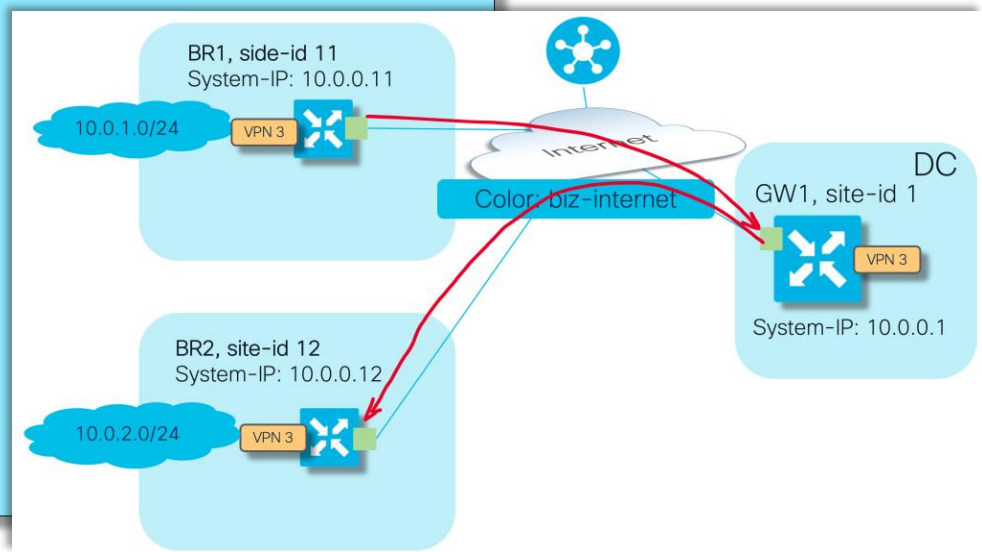

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4 ½. Traffic engineering with “set tloc-action". Disjoined underlay.

Same control policy, but BR2 connected to a different transport.

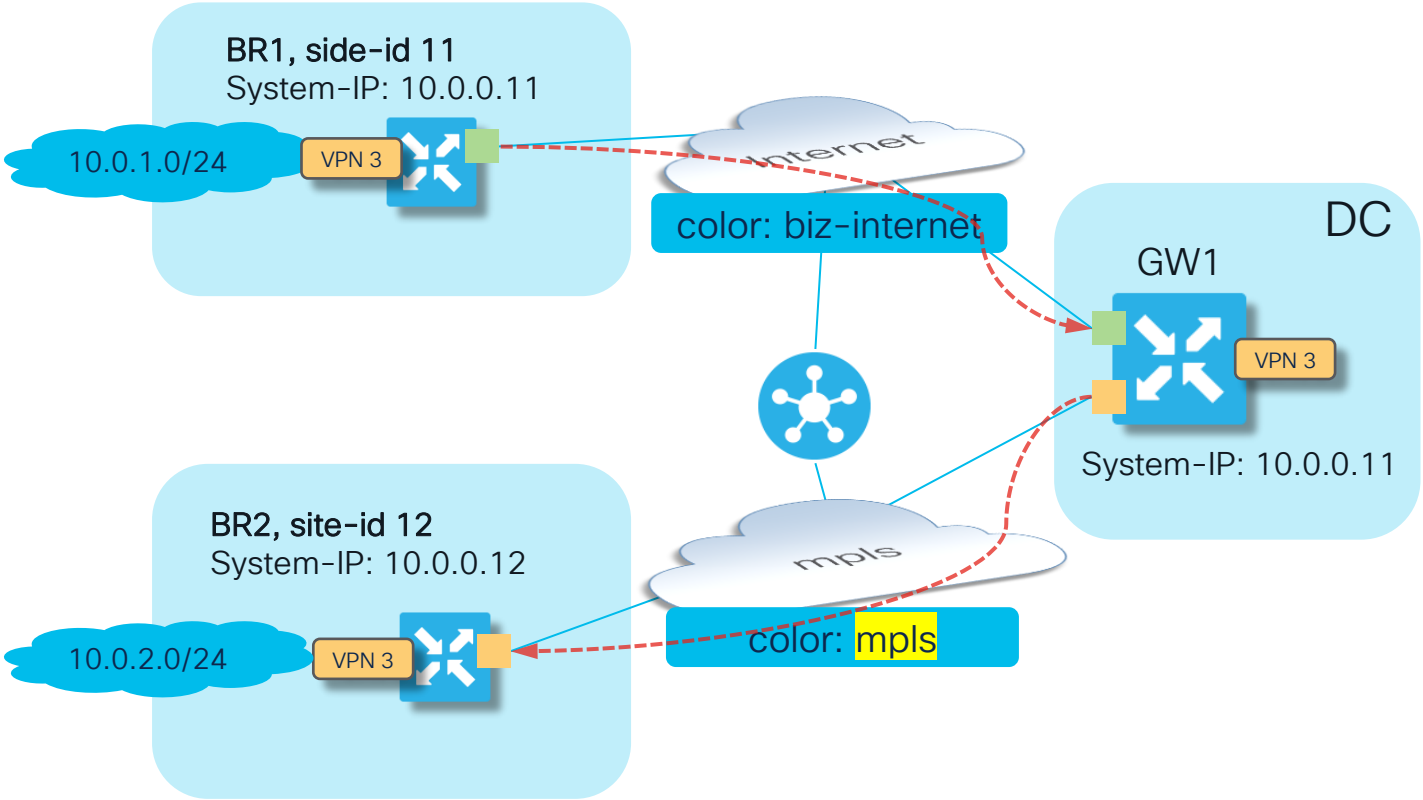

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4 ½. Traffic engineering with “set tloc-action". Disjoined underlay.

```
BR1#show sdwan omp routes 10.0.2.0/24 detail | exclude not\ set|metric|overlay|label|origin-|site|type
-----------------------------------------------------
omp route entries for vpn 3 route 10.0.2.0/24
-----------------------------------------------------
            RECEIVED FROM:
peer            10.0.0.101
path-id         1
status          Inv,U
loss-reason     tloc-action
lost-to-peer    10.0.0.101
lost-to-path-id 2
    Attributes:
     originator       10.0.0.12
     tloc             10.0.0.12, mpls, ipsec
            RECEIVED FROM:
peer            10.0.0.101
path-id         2
status          Inv,U
    Attributes:
     originator       10.0.0.12
     tloc             10.0.0.1, biz-internet, ipsec
     ultimate-tloc    10.0.0.12, mpls, ipsec -- primary
```

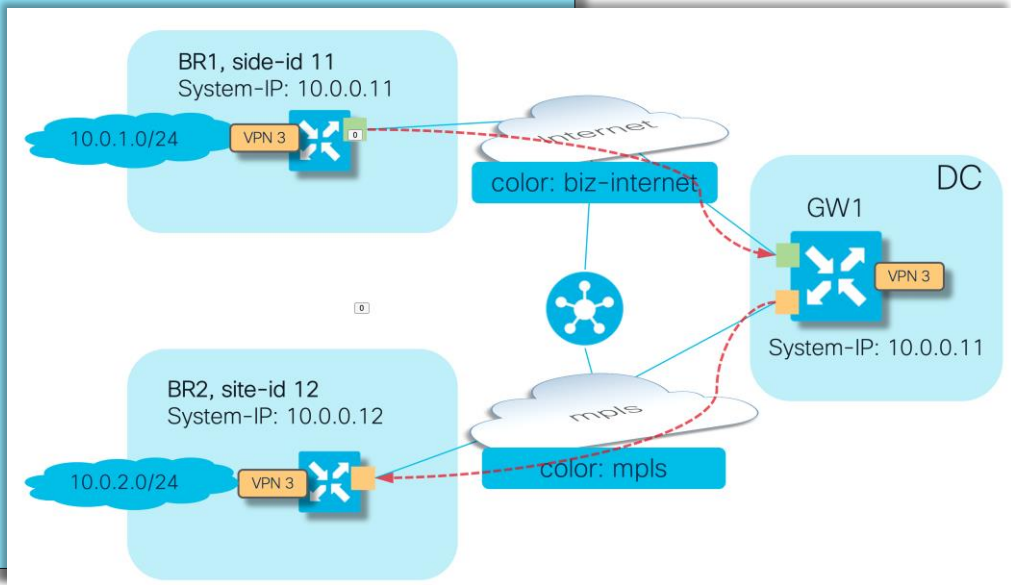

- Path 1 is unresolved because underlay is disjoined (no data plane tunnels with BR2)
- Why path 2 is unresolved and invalid?
- It is because different colors are not supported with tloc-action

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 4 ½. Traffic engineering with “set tloc-action". Disjoined underlay.

Note: tloc-action is only supported end-to-end if the transport color is the same from a site to the intermediate hop and from the intermediate hop to the final (ultimate) destination.

If the transport used to get to the intermediate hop from a site is a different color than the transport used from the intermediate hop to get to the final (ultimate) destination, then this will cause a policy failure with tloc-action.

Reference:

https://www.cisco.com/c/en/us/td/docs/routers/sdwan/command/sdwan-cr-book/config-cmd.html#r_action_1267.xml

Enhancement request: CSCvr80957

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Interesting cases with Data and AAR policies

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can’t install policy in large-scale network after reload

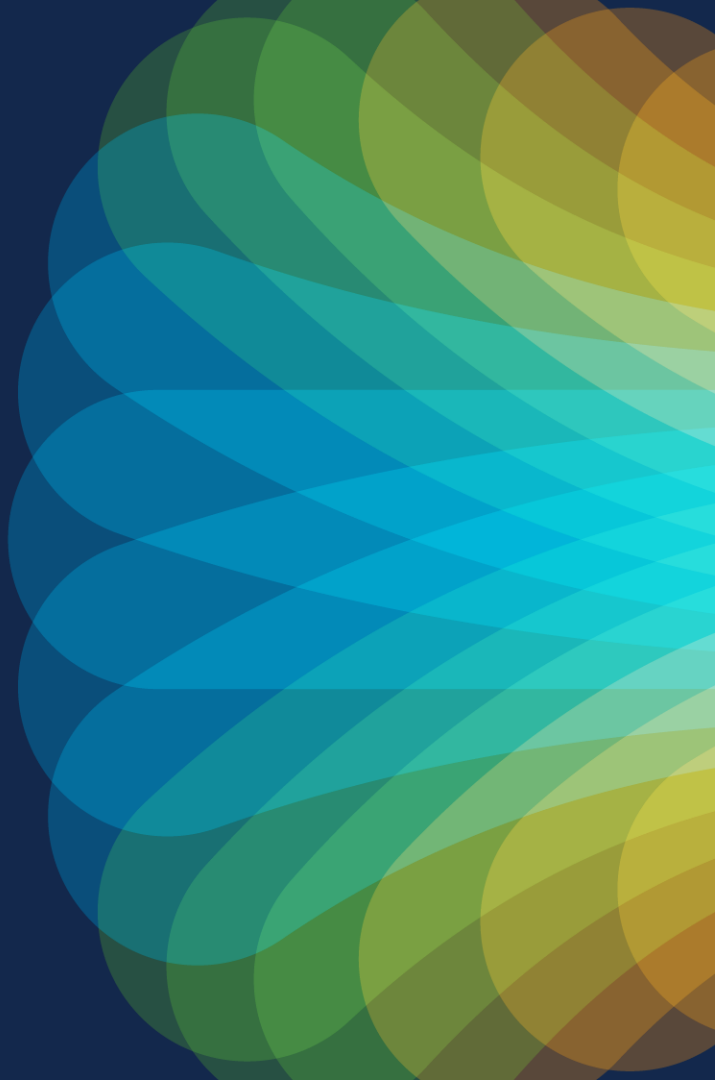

Case 5. Device can't install policy in large-scale network after reload

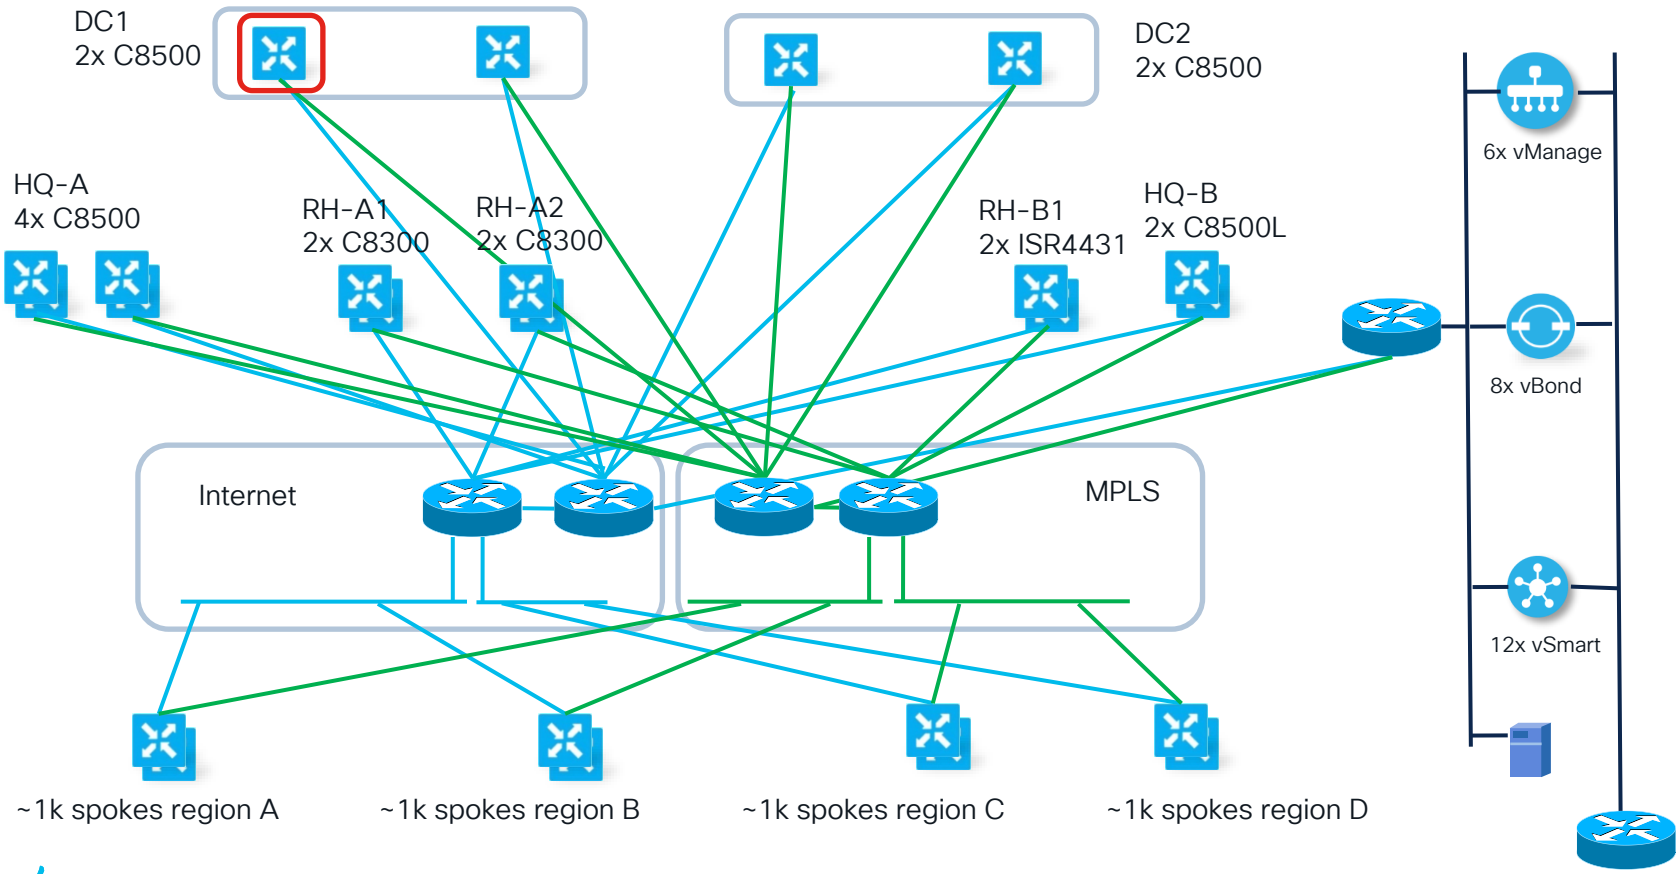

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can’t install policy in large-scale network after reload

Symptoms:

- Hub router was reloaded to perfrom software upgrade
- After the upgrade, device does not install any policy anymore
- Control connections are mostly up and running
- Hub router was sucessfully dongraded to exclude possibility of software defect with the same result: no policy installed on the device

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

Toubleshooting from vSmart side
OMP peering established and stable (hence underlying control connection as well):

```
vsmart1# show omp peers 10.0.0.101
R -> routes received
I -> routes installed
S -> routes sent

                        DOMAIN    OVERLAY   SITE
PEER            TYPE    ID        ID        ID        STATE    UPTIME           R/I/S
-------------------------------------------------------------------------------------
10.0.0.101       vedge   1         1         12        up       0:00:18:13       10/0/2
```

And policy assigned properly on vSmart:

```
vSmart1# show support omp peer peer-ip 1.1.1.101 | in -pol
        site-pol: STL_DC,STL_DC_1 route-pol-in: None route-pol-out: CPL_DC_1 data-pol-in: _VPN_LIST10_QOS_MARKING
           data-pol-out: None pfr-pol: _VPN_LIST10_APP_Route1 mem-pol: None cflowd:None
<data-policy>
<direction>from-service</direction></data-policy><app-route-policy>
</app-route-policy>
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

And translated to XML properly:

```
vsmart1# show support omp peer peer-ip 10.0.0.101 | begin "Policy received" | until "Statistics"
 Policy received: Complete
       Forwarding policy len: 4981
<data-policy>
  <name>_VPN_LIST10_QOS_MARKING</name>
  <vpn-list>
    <name>VPN_LIST10</name>
    <sequence>
      <seq-value>1</seq-value>
      <match>
        <destination-port>5000</destination-port>
      </match>
      <action>
        <action-value>accept</action-value>
        <set>
          <dscp>46</dscp>
          <forwarding-class>Queue0</forwarding-class>
        </set>
      </action>
<app-route-policy>
  <name>_VPN_LIST10_APP_Route1</name>
  <vpn-list>
    <name>VPN_LIST10</name>
    <sequence>
      <seq-value>1</seq-value>
      <match>
        <source-ip>0.0.0.0/0</source-ip>
        <destination-port>5000</destination-port>
      </match>
      <action>
        <sla-class>
          <sla-class-name>SLA_CLASS1</sla-class-name>
          <preferred-color>mpls</preferred-color>
        </sla-class>
      </action>
    </sequence>
<…>
 Statistics:
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

Troubleshooting from device side:

1. No policy chages, just hub router was reloaded so we are not checking commit changes
2. Control connnetions are up and mostly stable:

```
DC1_101#show sdwan control connections
                                                  PEER            PEER                                    CONTROLLER
PEER    PEER PEER      SITE  DOMAIN PEER          PRIV  PEER      PUB                                     GROUP
TYPE    PROT SYSTEM IP ID    ID     PRIVATE IP    PORT  PUBLIC IP PORT  ORGANIZATION      LOCAL COLOR   PROXY STATE UPTIME      ID
-------------------------------------------------------------------------------------------------------------------------------
vsmart  dtls 1.1.1.20  1     1      10.50.1.20    12346 10.50.1.20 12346 OrgName 1 - 31337  biz-internet  No    up    0:00:19:03 1
vsmart  dtls 1.1.1.21  1     1      10.50.1.21    12346 10.50.1.21 12346 OrgName 1 - 31337  biz-internet  No    up    0:00:03:35 2
vsmart  dtls 1.1.1.21  1     1      10.50.1.21    12346 10.50.1.21 12346 OrgName 1 - 31337  mpls          No    up    0:00:08:38 2
vsmart  dtls 1.1.1.27  1     1      10.50.1.27    12346 10.50.1.27 12346 OrgName 1 - 31337  mpls          No    up    0:00:08:35 8
vbond   dtls 0.0.0.0   0     0      10.50.1.10    12346 10.50.1.10 12346 OrgName 1 - 31337  biz-internet  -     up    0:00:26:19 0
vbond   dtls 0.0.0.0   0     0      10.50.1.13    12346 10.50.1.13 12346 OrgName 1 - 31337  mpls          -     up    0:00:23:20 0
vmanage dtls 1.1.1.4   1     0      10.50.1.4     12346 10.50.1.4  12346 OrgName 1 - 31337  biz-internet  No    up    0:00:19:06 0
```

as well as OMP peering (not really necessary to check because it is stable from vSmart perspective):

```
DC1_101#show sdwan omp peers
R -> routes received
I -> routes installed
S -> routes sent

TENANT                          DOMAIN    OVERLAY   SITE      REGION
ID        PEER              TYPE    ID        ID        ID        ID        STATE    UPTIME           R/I/S
-------------------------------------------------------------------------------------------------------------
0         1.1.1.20          vsmart  1         1         1         None      up       0:00:18:17       0/0/10
0         1.1.1.26          vsmart  1         1         1         None      up       0:00:14:26       59742/22939/10
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

But the mystery is that device still does not have any policy:

```
DC1_101#show sdwan omp summary | include policy
policy-sent             0
policy-received         0
```

And certainly other commands confirm the same:

```
DC1_101#show sdwan from-vsmart commit-history
summary

DC1_101#show sdwan policy from-vsmart
% No entries found.
```

From the logs it says no policy assigned and seems other vSmarts are less stable (*hint!*):

```
Mar 16 12:17:21.268: %Cisco-SDWAN-DC1_101-OMPD-3-ERRO-400002: vSmart peer 1.1.1.21 state changed to Init
Mar 16 12:17:21.268: %Cisco-SDWAN-DC1_101-OMPD-6-INFO-400005: Number of vSmarts connected : 2
Mar 16 12:17:23.268: %Cisco-SDWAN-DC1_101-OMPD-6-INFO-400007: Using empty policy from peer 1.1.1.20
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

Something strange happens on a device? Then always check QFP drop counters first!

```
DC1_101#show platform hardware qfp active statistics drop clear
Last clearing of QFP drops statistics : Thu Mar 16 13:20:11 2023


-------------------------------------------------------------------------
Global Drop Stats                          Packets                 Octets
-------------------------------------------------------------------------
Disabled                                         2                    506
Ipv6NoRoute                                      1                     56
Nat64v6tov4                                      6                    480
PuntPerCausePolicerDrops                   8504352             1625710362
SdwanImplicitAclDrop                          2844                 451300


DC1_101#show platform hardware qfp active statistics drop detail
Last clearing of QFP drops statistics : Thu Mar 16 13:20:11 2023
 (13s ago)

-------------------------------------------------------------------------------
   ID  Global Drop Stats                          Packets                 Octets
-------------------------------------------------------------------------------
   206 PuntPerCausePolicerDrops                   49419                 9442474
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can’t install policy in large-scale network after reload

Then you can use packet-trace to see dropped packets detals:

```
DC1_101#debug platform condition both
DC1_101#debug platform packet-trace drop code 206
DC1_101#debug platform packet-trace packet 1024
 Please remember to turn on 'debug platform condition start' for packet-trace to work
DC1_101#debug platform condition start


DC1_101#show platform packet-trace summary
Pkt   Input                     Output                    State  Reason
1     Te0/0/0                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
2     Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
3     Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
4     Te0/0/0                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
5     Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
6     Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
7     Te0/0/0                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
8     Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
9     Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
10    Te0/0/0                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
11    Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
12    Te0/0/1                   internal0/0/rp:0          DROP   206 (PuntPerCausePolicerDrops)
…
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

While checking packets, noticed that some of them are originated from controllers:

```
DC1_101#show platform packet-trace packet 3
Packet: 3           CBUG ID: 3
Summary
  Input     : TenGigabitEthernet0/0/0
  Output    : internal0/0/rp:0
  State     : DROP 206 (PuntPerCausePolicerDrops)
  Timestamp
    Start   : 2699999601354 ns (03/16/2023 13:22:19.296832 UTC)
    Stop    : 2700000237397 ns (03/16/2023 13:22:19.297468 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : TenGigabitEthernet0/0/0
    Output      : <unknown>
    Source      : 10.50.1.26
    Destination : 10.60.1.6
    Protocol    : 17 (UDP)
      SrcPort   : 12346
      DstPort   : 12346
  Feature: SDWAN Implicit ACL
    Action : ALLOW
    Reason : SDWAN_TUN_CTRL
```

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can’t install policy in large-scale network after reload

Facts and summary so far:

- No issues observed from vSmart perspective
- The hub router does not have policies installed but it should as per vSmart view
- The hub router dropping packets extensively with some unknown code “PuntPerCausePolicerDrops”, some of them belong to a traffic from controllers
- We can guess by the name of the drop reason that there is some policer
- Reload of the device is a trigger

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

If you search for "PuntPerCausePolicerDrops" on the Internet, the very first result will help to explain the reason and help to find corresponding commands to check drop level settings

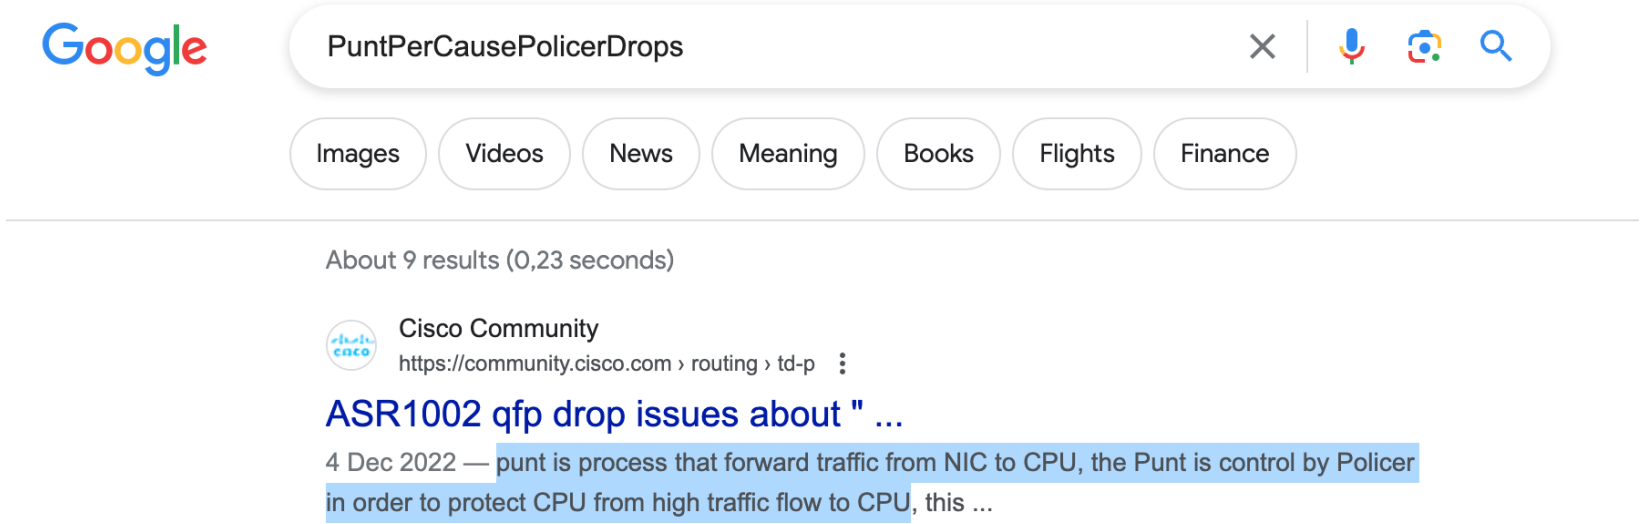

So, there is a rate limiter for punted (sent to CPU control plane) packets which is exceeded

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

... and you will find that a lot of packets dropped by this policer:

```
DC1_101#show platform software punt-policer drop-only

Per Punt-Cause Policer Configuration and Packet Counters


Punt                  Config Rate(pps)     Conform Packets            Dropped Packets           Config Burst(pkts)  Config Alert
Cause  Description    Normal   High     Normal           High      Normal           High         Normal   High     Normal   High
-------------------------------------------------------------------------------------------------------------------------------
 11    For-us data    40000    5000     230              19482005  0                14789128     40000    5000     Off      Off

DC1_101#show platform software punt-policer drop-only

Per Punt-Cause Policer Configuration and Packet Counters


Punt                  Config Rate(pps)     Conform Packets            Dropped Packets           Config Burst(pkts)  Config Alert
Cause  Description    Normal   High     Normal           High      Normal           High         Normal   High     Normal   High
-------------------------------------------------------------------------------------------------------------------------------
 11    For-us data    40000    5000     232              19607381  0                14883968     40000    5000     Off      Off
```

Why? Keep in mind there are ~4000 routers trying to establish tunnels at the same time and the hub has default settings for control plane policing

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 5. Device can't install policy in large-scale network after reload

Solution, increase punt policer:

```
DC1_101#config-t

admin connected from 127.0.0.1 using console on Router
Router(config)# platform punt-policer 11 10000 high
Router(config)# commit
Commit complete.
```

Test with reload confirms policy installed sucessully:

```
Mar 16 14:44:16.089: %Cisco-SDWAN-DC1_101-OMPD-6-INFO-400002: vSmart peer 1.1.1.26 state changed to Handshake
Mar 16 14:44:16.098: %Cisco-SDWAN-DC1_101-OMPD-5-NTCE-400002: vSmart peer 1.1.1.26 state changed to Up
Mar 16 14:44:16.098: %Cisco-SDWAN-DC1_101-OMPD-6-INFO-400005: Number of vSmarts connected : 2
Mar 16 14:44:20.260: %Cisco-SDWAN-DC1_101-OMPD-6-INFO-400007: Using policy from peer 1.1.1.20
```

```
DC1_101#sh sdwan omp summary | include policy

policy-sent             0
policy-received         4

DC1_101#sh sdwan bfd summary
sessions-total          8076
sessions-up             0
sessions-max            8076
sessions-flap           0
poll-interval           600000
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

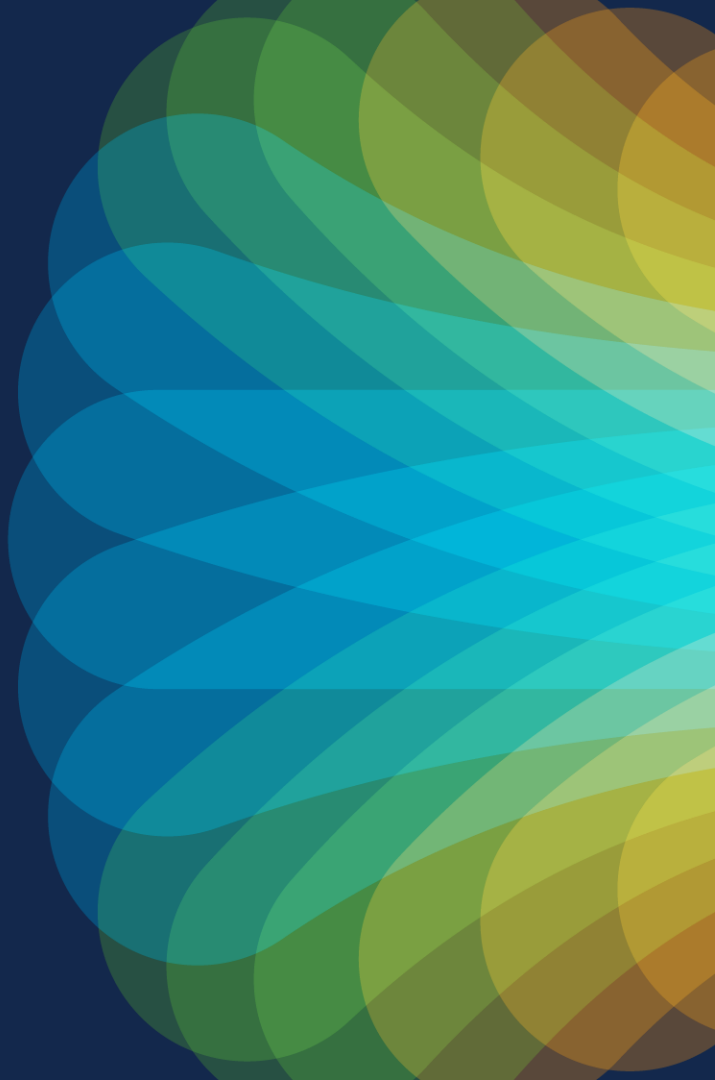

Case 6. Traffic blackholing with DIA policy

Typical symptoms:

- User traffic affected, no connections to some enterprise internal servers
- Only some specific application traffic affected, but destination is reachable with ”ping”
- Trigger is an implementation of Direct Internet Access (DIA) data policy or Cloud on Ramp (CoR) for SaaS, less usually AAR policy

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

Typical symptoms from a user perspective: no connection to internal servers (timeout)

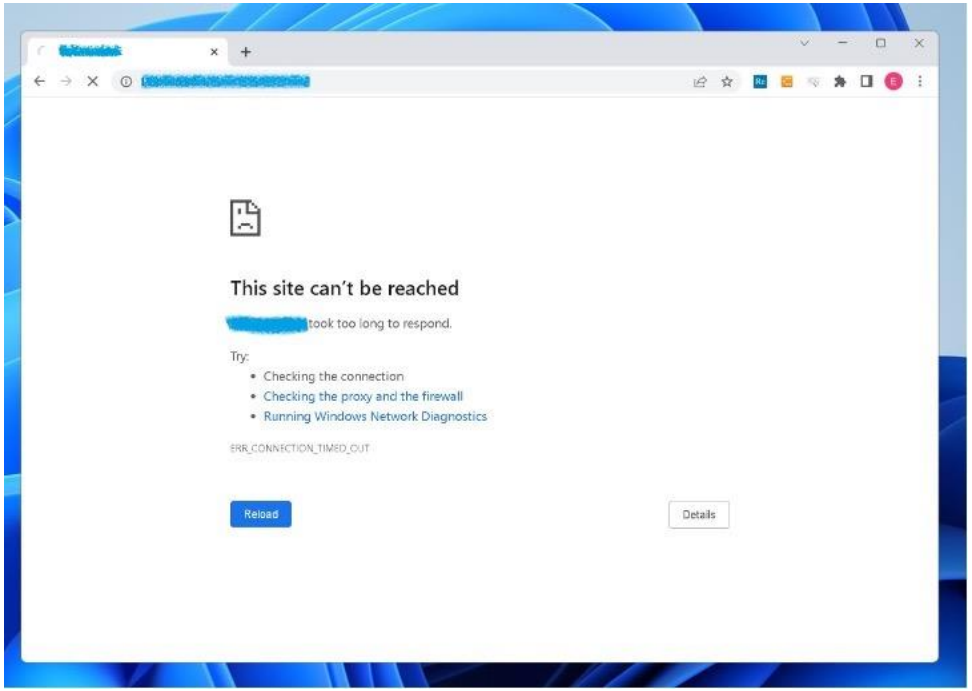

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

Data policy is very simple. Aim is to to implement DIA for Office 365, for example:

```
cE2_BR2#show sdwan policy from-vsmart
from-vsmart data-policy VPN_1_NAT
 direction from-service
 vpn-list VPN_1
  sequence 1
   match
    app-list O365
   action accept
    nat use-vpn 0
    no nat fallback
  default-action accept
from-vsmart lists vpn-list VPN_1
 vpn 1
from-vsmart lists app-list O365
 app ms-office-365
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

Just as per the troubleshooting workflow, we need ot ensure correct TLOC/next-hop/color/interface selection for the affected traffic:

```
cE1#show sdwan policy service-path vpn 1 interface GigabitEthernet 4 source-ip 192.168.4.100 dest-ip 10.0.1.12 protocol 6 dest-port 443
Next Hop: Blackhole
```

There are few typical reasons for this, the most common is lack of a route to the destination, which is not the case here:

```
cE1#sh ip route vrf 1 10.0.1.12

Routing Table: 1
Routing entry for 10.0.0.0/16
  Known via "omp", distance 251, metric 0, type omp
  Last update from 169.254.206.35 on Sdwan-system-intf, 00:02:24 ago
  Routing Descriptor Blocks:
  * 169.254.206.35 (default), from 169.254.206.35, 00:02:24 ago, via Sdwan-system-intf
      opaque_ptr 0x7FB0E6FB62A0
      Route metric is 0, traffic share count is 1
```

And ”ping” works just fine, also confirmed with “**show sdwan policy service-path**”:

```
cE1#show sdwan policy service-path vpn 1 interface GigabitEthernet 4 source-ip 192.168.4.100 dest-ip 10.0.1.12 protocol 1 app ping
Next Hop: IPsec
  Source: 192.168.10.11 12366 Destination: 192.168.9.35 12346 Local Color: biz-internet Remote Color: biz-internet Remote System IP:
169.254.206.35
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

If traffic blackholed, there is a reasonable assumption that it should be dropped on the device, right? Let's check QFP drop counters:

```
cE1#show platform hardware qfp active statistics drop clear
Last clearing of QFP drops statistics : Mon May  8 17:45:08 2023


-------------------------------------------------------------------------
Global Drop Stats                          Packets                Octets
-------------------------------------------------------------------------
BFDoffload                                     345                 29670
Disabled                                       247                 15414
Ipv4EgressIntfEnforce                            8                  1544
Ipv4NoAdj                                        6                   413
Ipv6NoRoute                                      5                   280
Nat64v6tov4                                      6                   480
SdwanDataPolicyDrop                            114                 15504
SdwanImplicitAclDrop                         11544               1984076
UnconfiguredIpv6Fia                            502                 54287

cE1#show platform hardware qfp active statistics drop
Last clearing of QFP drops statistics : Mon May  8 17:45:08 2023
 (58s ago)

-------------------------------------------------------------------------
Global Drop Stats                          Packets                Octets
-------------------------------------------------------------------------
BFDoffload                                      63                  5418
```

But there are only legitimate drops...

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

Packet trace to rescue again?

```
cE1#debug platform condition ipv4 10.0.1.12/32 both
cE1#debug platform packet-trace packet 1024 fia-trace
 Please remember to turn on 'debug platform condition start' for packet-trace to
work
cE1#debug platform condition start
cE1#show platform packet-trace summary
…
680   Tu2                       Gi4                       FWD
681   Gi4                       Gi2                       FWD
682   Gi4                       Gi2                       FWD
683   Tu2                       Gi4                       FWD
684   Tu2                       Gi4                       FWD
685   Gi4                       Gi2                       FWD
686   Gi4                       Gi2                       FWD
687   Tu2                       Gi4                       FWD
688   Tu2                       Gi4                       FWD
689   Gi4                       Gi2                       FWD
690   Gi4                       Gi2                       FWD
691   Tu2                       Gi4                       FWD
692   Gi4                       Gi2                       FWD
693   Gi4                       Gi2                       FWD
694   Tu2                       Gi4                       FWD
695   Tu2                       Gi4                       FWD
696   Gi4                       Gi2                       FWD
697   Gi4                       Gi3                       FWD
698   Gi4                       Gi3                       FWD
699   Gi4                       Gi3                       FWD
700   Gi4                       Gi3                       FWD
701   Gi4                       Gi3                       FWD
702   Gi4                       Gi3                       FWD
```

Problematic becase there are maybe multiple flows in parallel unless you know src, dst precisely

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

In a live network packet-trace may cause a lot of confusion if you don't know where to look at specifically, so NWPI is preferred because it will trace end-to-end and show all-in-one insignt:

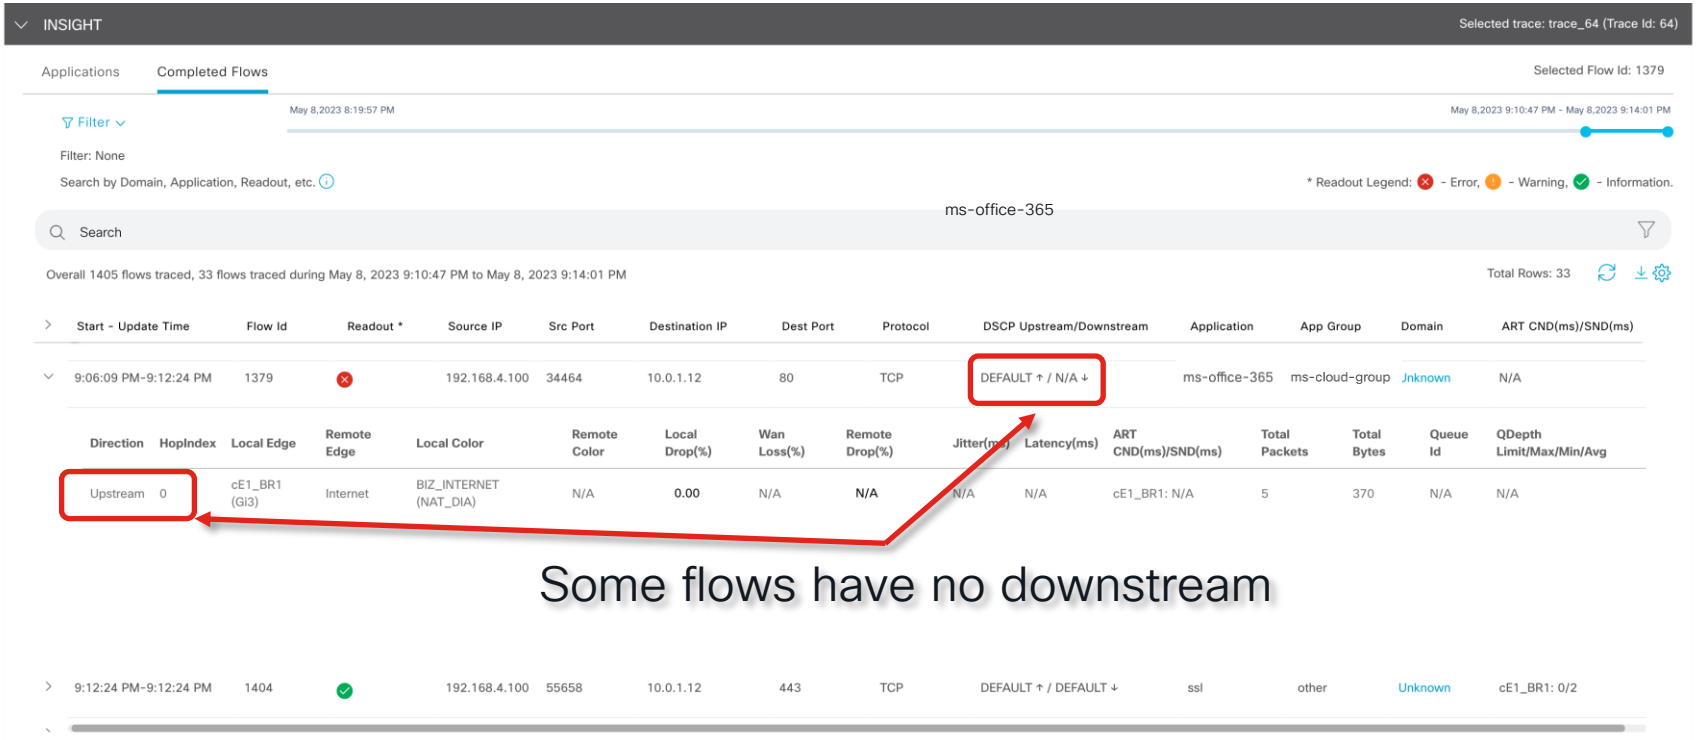

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

From ”insight - advanced view”, you can find that DIA data policy was applied:

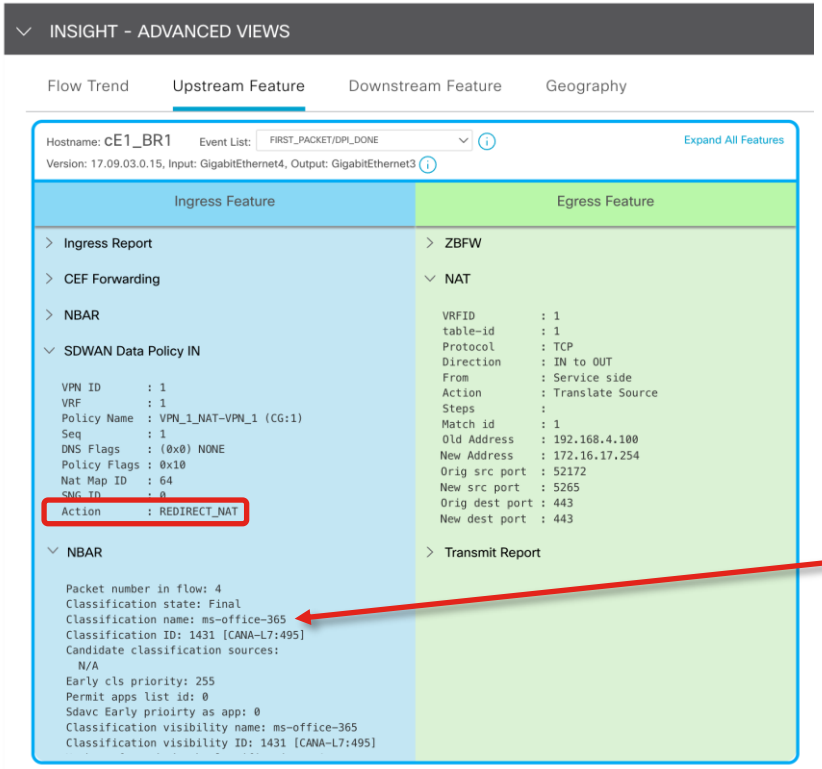

Here you can see DPI misclassified internal application as “ms-office-365” and traffic was sent to DIA circuit instead of overlay tunnel, hence, causing blackholing

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

Why misclassification happens?

Great answer is that it depends. Sometimes apps are just hard to recognize and differentiate (on-prem services vs SaaS) or it may be a bug.

But solution for the DIA case is very simple:

- Ensure RFC1918 prefixes excluded from DPI evaluation
- Inherited benefit: reduced load on a router because less traffic to be processed by DPI enginer (NBAR)

How?

- Insert data policy (sequence) prior sequence perfrorming NAT to match based on RFC1918 source addresses and accept them (accept is a final action)

```
policy
 data-policy VPN_1_NAT
  vpn-list VPN_1
   !
   sequence 1
    match
     destination-data-prefix-list RFC1918
    !
    action accept
    !
   !
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 6. Traffic blackholing with DIA policy

Keep in mind, same problem may be experienced with with CoR for SaaS:

```
Router#show sdwan policy from-vsmart
from-vsmart app-route-policy _CC1_AAR_POLICY
vpn-list CC1
sequence 41
   match
    source-ip           0.0.0.0/0
    cloud-saas-app-list office365_apps
   action
    count office365_apps_-856788698
    cloud-saas
  sequence 51
   match
    source-ip           0.0.0.0/0
    cloud-saas-app-list salesforce_apps
   action
    count salesforce_apps_-856788698
    cloud-saas
  default-action sla-class DEFAULT
```

Solution is the same, configured data policy to override CoR for SaaS

Why so complicated? I thought the same. That's why enhancement CSCvv68740 was implemented in 20.13/17.13 to exclude RFC1918 by default

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. DSCP marking not applied with AAR policy

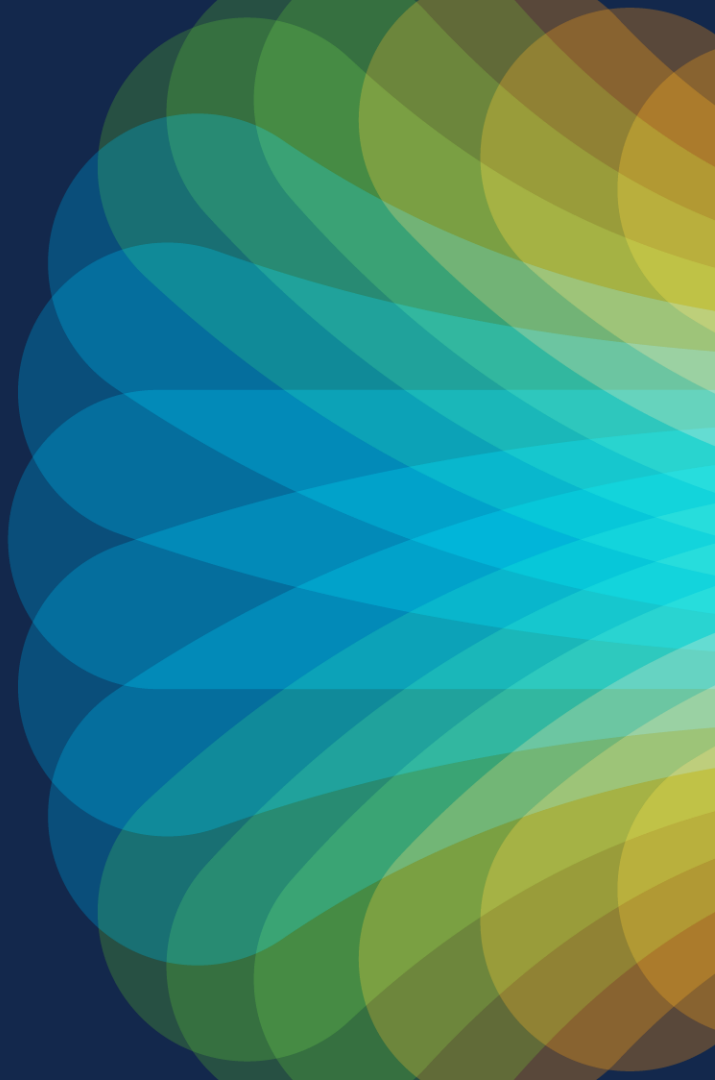

Case 7. Incorrect DSCP marking symptoms

- **`show sdwan app-fwd cflowd flows table`** shows DSCP mark as “0”

```
cE1_BR1#show sdwan app-fwd cflowd flows vpn 4
Generating output, this might take time, please wait ...
app-fwd cflowd flows vpn 4 src-ip 192.168.5.197 dest-ip
192.168.4.196 src-port 22 dest-port 37748 dscp 4 ip-proto 6
 tcp-cntrl-bits           24
 icmp-opcode              0
 total-pkts               6
 total-bytes              2064
 start-time               "Fri Dec 22 15:35:11 2023"
 egress-intf-name         GigabitEthernet4
 ingress-intf-name        GigabitEthernet3
 application              ssh
 family                   terminal
 drop-cause               "No Drop"
 drop-octets              0
 drop-packets             0
 sla-not-met              0
 color-not-met            0
 queue-id                 2
 initiator                2
 tos                      0
 dscp-output              0
 sampler-id               0
 fec-d-pkts               0
 fec-r-pkts               0
 pkt-dup-d-pkts-orig      0
 pkt-dup-d-pkts-dup       0
 pkt-dup-r-pkts           0
 pkt-cxp-d-pkts           0
 category                 0
 service-area             0
 cxp-path-type            0
 region-id                0
 ssl-read-bytes           0
 ssl-written-bytes        0
 ssl-en-read-bytes        0
 ssl-en-written-bytes     0
 ssl-de-read-bytes        0
 ssl-de-written-bytes     0
 ssl-service-type         0
 ssl-traffic-type         0
 ssl-policy-action        0
 appqoe-action            0
 appqoe-sn-ip             0.0.0.0
 appqoe-pass-reason       0
 appqoe-dre-input-bytes   0
 appqoe-dre-input-packets 0
 appqoe-flags             0
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking symptoms (2)

- But data policy configured to mark it as “10” (AF11) and received on the device:

```
cE1_BR1#show sdwan policy from-vsmart data-policy
from-vsmart data-policy SET_DSCP
 direction from-service
 vpn-list VPN_4
  sequence 1
   match
    destination-port 22
    protocol         6
   action accept
    cflowd
    set
     dscp 10
  sequence 2
   match
    source-port 22
    protocol    6
   action accept
    cflowd
    set
     dscp 10
  default-action accept
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking symptoms (3)

- Cflow template for your reference also:

```
vsmart1# show running-config policy cflowd-template
policy
 cflowd-template test-cflowd-template
  template-refresh 90
  collector vpn 0 address 192.168.10.240 port 9555 transport transport_udp
  !
 !
!
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking symptoms (4)

- FIA trace was done to confirm if DSCP 10 was set and it was

```
cE1_BR1#show platform packet-trace packet 1
Packet: 1           CBUG ID: 138
Summary
  Input     : GigabitEthernet4
  Output    : GigabitEthernet3
  State     : FWD
  Timestamp
    Start   : 15111231553111 ns (12/22/2023 15:25:22.147214 UTC)
    Stop    : 15111231650980 ns (12/22/2023 15:25:22.147311 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet4
    Output      : <unknown>
    Source      : 192.168.4.196
    Destination : 192.168.5.197
    Protocol    : 6 (TCP)
      SrcPort   : 22
      DstPort   : 44408

  <skipped>
```

```
Feature: IPV4_INPUT_FNF_FIRST
    Entry       : Input - 0x814db670
    Input       : GigabitEthernet4
    Output      : <unknown>
    Lapsed time : 1682 ns
  Feature: DEBUG_COND_APPLICATION_IN_CLR_TXT
    Entry       : Input - 0x814ca90c
    Input       : GigabitEthernet4
    Output      : <unknown>
    Lapsed time : 32 ns
  Feature: SDWAN Data Policy IN
    VPN ID      : 4
    VRF         : 2
    Policy Name : SET_DSCP-VPN_4 (CG:4)
    Seq         : 2
    DNS Flags   : (0x0) NONE
    Policy Flags : 0x408
    Policy Flags2: 0x0
    Action      : FNF
    Action      : SET_DSCP af11(10)
  Feature: SDWAN_POLICY_FIA

<rest is skipped>
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking symptoms (5)

- Which output should we trust?
- Packet capture on a remote host was done and it was confirmed that DSCP set properly by the router

```
root@user:/home/user# tcpdump -v "host 192.168.4.196" -i ens192
tcpdump: listening on ens192, link-type EN10MB (Ethernet), capture size 262144 bytes
17:01:45.554798 IP (tos 0x28, ttl 62, id 55357, offset 0, flags [DF], proto TCP (6), length 88)
    192.168.4.196.37748 > 192.168.5.197.ssh: Flags [P.], cksum 0x8f47 (correct), seq 290589212:290589248, ack 2393058474, win 501,
options [nop,nop,TS val 3705085476 ecr 4291921734], length 36
17:01:45.555261 IP (tos 0x10, ttl 64, id 25646, offset 0, flags [DF], proto TCP (6), length 152)
    192.168.5.197.ssh > 192.168.4.196.37748: Flags [P.], cksum 0x8c64 (incorrect -> 0x3cdd), seq 1:101, ack 36, win 501, options
[nop,nop,TS val 4291946734 ecr 3705085476], length 100
17:01:45.555967 IP (tos 0x28, ttl 62, id 55358, offset 0, flags [DF], proto TCP (6), length 52)
    192.168.4.196.37748 > 192.168.5.197.ssh: Flags [.], cksum 0x9a62 (correct), ack 101, win 501, options [nop,nop,TS val
3705085477 ecr 4291946734], length 0
17:01:45.710499 IP (tos 0x28, ttl 62, id 55359, offset 0, flags [DF], proto TCP (6), length 88)
    192.168.4.196.37748 > 192.168.5.197.ssh: Flags [P.], cksum 0xe590 (correct), seq 36:72, ack 101, win 501, options [nop,nop,TS
val 3705085632 ecr 4291946734], length 36
<rest is skipped>
```

* 0x28 ToS HEX = 10 DSCP decimal = AF11

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking symptoms (6)

- Let's take a look at packet-trace again

```
cE1_BR1#show platform packet-trace packet 1
Packet: 1           CBUG ID: 138
Summary
  Input     : GigabitEthernet4
  Output    : GigabitEthernet3
  State     : FWD
  Timestamp
    Start   : 15111231553111 ns (12/22/2023 15:25:22.147214 UTC)
    Stop    : 15111231650980 ns (12/22/2023 15:25:22.147311 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet4
    Output      : <unknown>
    Source      : 192.168.4.196
    Destination : 192.168.5.197
    Protocol    : 6 (TCP)
      SrcPort   : 22
      DstPort   : 44408

  <skipped>
```

```
Feature: IPV4_INPUT_FNF_FIRST
  Entry       : Input - 0x814db670
  Input       : GigabitEthernet4
  Output      : <unknown>
  Lapsed time : 1682 ns
Feature: DEBUG_COND_APPLICATION_IN_CLR_TXT
  Entry       : Input - 0x814ca90c
  Input       : GigabitEthernet4
  Output      : <unknown>
  Lapsed time : 32 ns
Feature: SDWAN Data Policy IN
  VPN ID      : 4
  VRF         : 2
  Policy Name : SET_DSCP-VPN_4 (CG:4)
  Seq         : 2
  DNS Flags   : (0x0) NONE
  Policy Flags : 0x408
  Policy Flags2: 0x0
  Action      : FNF
  Action      : SET_DSCP af11(10)
Feature: SDWAN_POLICY_FIA

<rest is skipped>
```

Data policy action "cflowd" is FNF (Flexible Net Flow), but FNF feature itself preceding data policy

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking – solution

- Is it order of operations issue?
- Yes, but there is an option available to ensure DSCP/ToS marking recorded into NetFlow data anyway.
- I did not show cflowd template view as per the router because then problem and solution would be obvious (if you attentive enough):

```
cE1_BR1#show sdwan policy from-vsmart cflowd-template
from-vsmart cflowd-template test-cflowd-template
 flow-active-timeout    600
 flow-inactive-timeout  60
 template-refresh       90
 flow-sampling-interval 1
 protocol               ipv4
 no collect-tloc-loopback
 customized-ipv4-record-fields
  no collect-tos
  no collect-dscp-output
 collector vpn 0 address 192.168.10.240 port 9555 transport transport_udp
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking – solution (2)

- Let's fix it (the feature introduced in 20.6+ specifically to address this problem)

```
vsmart1(config)# show configuration
policy
 cflowd-template test-cflowd-template
  customized-ipv4-record-fields
   collect-tos
   collect-dscp-output
  !
 !
!
vsmart1(config)# commit
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Case 7. Incorrect DSCP marking – solution (3)

... and check the output again

```
cE1_BR1#show sdwan app-fwd cflowd flows vpn 4
Generating output, this might take time, please wait ...
app-fwd cflowd flows vpn 4 src-ip 192.168.4.196 dest-ip
192.168.5.197 src-port 33418 dest-port 22 dscp 4 ip-proto 6
 tcp-cntrl-bits           24
 icmp-opcode              0
 total-pkts               26
 total-bytes              1568
 start-time               "Fri Dec 22 16:28:57 2023"
 egress-intf-name         GigabitEthernet3
 ingress-intf-name        GigabitEthernet4
 application              ssh
 family                   terminal
 drop-cause               "No Drop"
 drop-octets              0
 drop-packets             0
 sla-not-met              0
 color-not-met            0
 queue-id                 2
 initiator                1
 tos                      16
 dscp-output              10
 sampler-id               0
 fec-d-pkts               0
 fec-r-pkts               0
 pkt-dup-d-pkts-orig      0
 pkt-dup-d-pkts-dup       0
 pkt-dup-r-pkts           0
 pkt-cxp-d-pkts           0
 category                 0
 service-area             0
 cxp-path-type            0
 region-id                0
 ssl-read-bytes           0
 ssl-written-bytes        0
 ssl-en-read-bytes        0
 ssl-en-written-bytes     0
 ssl-de-read-bytes        0
 ssl-de-written-bytes     0
 ssl-service-type         0
 ssl-traffic-type         0
 ssl-policy-action        0
 appqoe-action            0
 appqoe-sn-ip             0.0.0.0
 appqoe-pass-reason       0
 appqoe-dre-input-bytes   0
 appqoe-dre-input-packets 0
 appqoe-flags             0
```


© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

As a conclusion: Typical policy faults

Typical policy faults (generic)

Always keep in mind policy processing logic:

- `default-action reject` or `default-action accept`
- wrong direction of policy application (in vs out, from-tunnel vs from-service)
- subject to policy application has already processed by previous match statement (and match in a policy is final)
- policy application scope is too narrow or too wide (e.g. site-id not specified in a sequence match statement and action is applied to the whole set of site-list defined under **`apply-policy`** section)
- simple misconfigurations and typos (e.g. a prefix missing from a prefix-list, wrong mask, wrong site-id and so on).

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Typical Control Policy specific faults

- Control policy applied on inbound direction before OMP best-path selection resulting in backup paths missing
- Unconditional TLOC rewrites (e.g. “**`set tloc-list`**” and vSmart is not aware of TLOCs state)
- Attempt to use “**`set tloc-action`**” while “**`service TE`**” is not enabled on WAN Edge
- Attempt to glue/stick together different colors with “**`set tloc-action`**”

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Typical AAR and Data policies specific faults

- Common AAR issues:
  - return traffic is asymmetric. Does not mean that AAR function improperly (feature is unidirectional)
  - equal cost paths (ECMP) missing, hence only one path available and AAR has no choice.
- Common misunderstanding:
  - by default, it may take up to 1 hour for AAR policy to change a path (**`app-route poll-interval`** 600s x **`multiplier`** 6 = 1h)
  - **`bfd poll-interval`** impacts frequency of app-route poll-interval updates (accuracy), but not AAR reaction time (convergence) as such
- Common issues AAR+Data Policy: in short, DP overrides AAR, but considers AAR SLA class match (20.6+)
- Common AAR/DP misconfig: DPI matches internal traffic (e.g. Microsoft on-prem servers) and policy sends it to DIA causing blackholing
- Fallback issues: DIA **`nat fallback`** or SIG **`sig-action fallback-to-routing`** not configured by default.
- Policy bypass because first packet match fails (Policy-Bypass-FPM-Fail): may need **`policy flow-stickiness-disable`** (17.6+ feature)
- Fragmented packets match (e.g. UDP fragments considered matching to a sequence even if there is no UDP port info available in IP fragment)

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

App-Aware Routing and Data Policy Overlap

Policy Processing when packet is subject to match in both policies

For Your Reference

Guiding Principle:
Data Policy Makes Final Decision with Consideration for AAR SLA Match

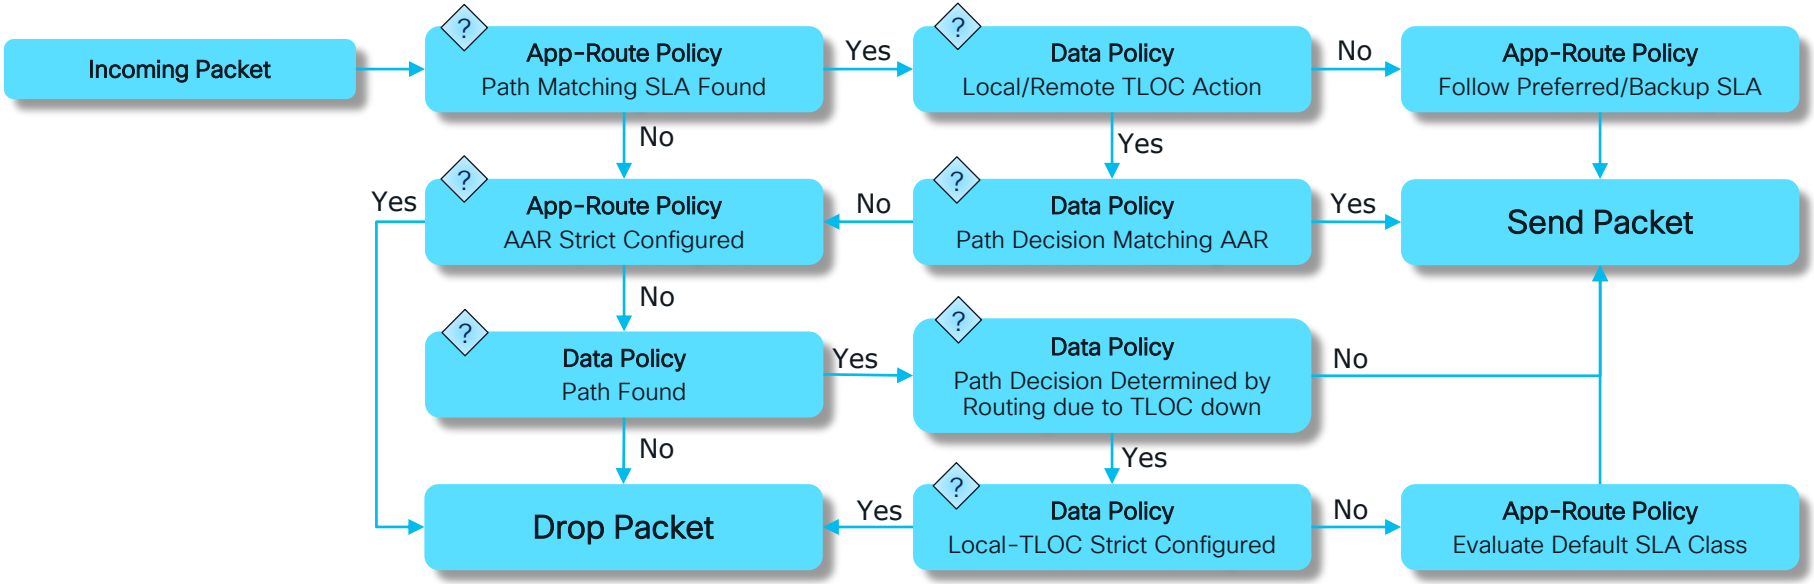

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

References and recommended resources

- Cisco Troubleshooting Tech Notes: https://www.cisco.com/c/en/us/support/routers/sd-wan/products-tech-notes-list.html
- BRKENT-3793/BRKTRS-3793 “Advanced SD-WAN Routing Troubleshooting”
- BRKTRS-3475 “Advanced Troubleshooting of CAT8k, ASR1k, ISR and SD-WAN Edge made easy”
- BRKRST-2791 “Building and Using Policies with Cisco SD-WAN”
- BRKENT-2477 “Cisco SD-WAN Troubleshooting”

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

Q&A

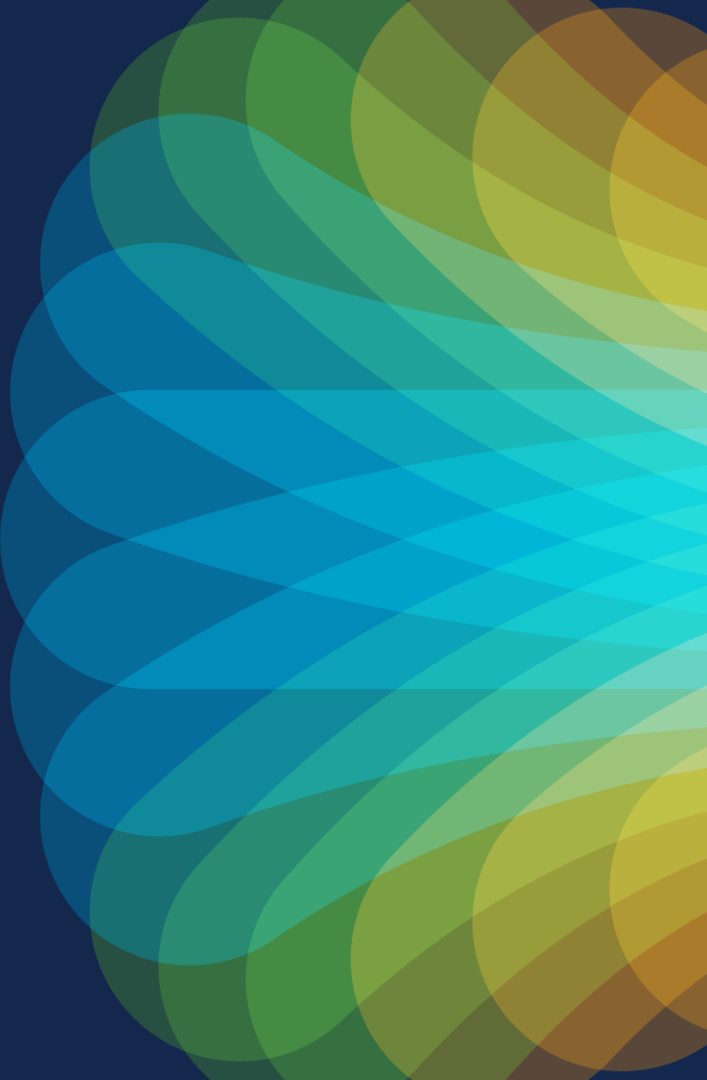

The bridge to possible

Thank you

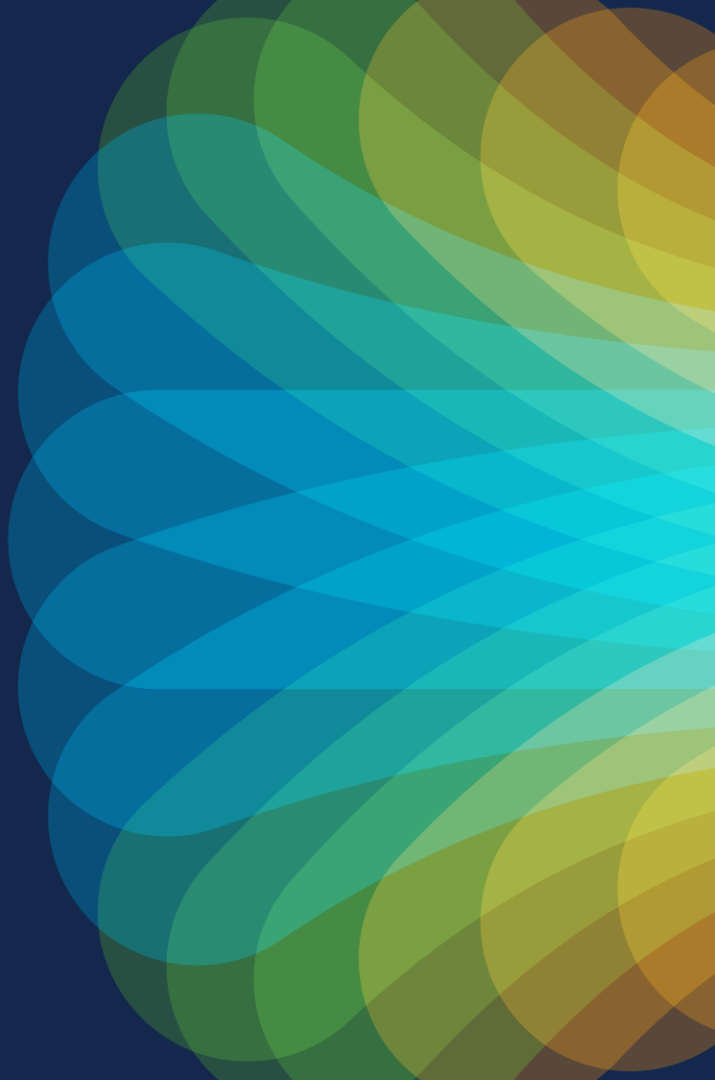

cisco Live!


cisco Live!
Let's go