Troubleshooting OMP Failure and TLOC-Action

Contents

Introduction

This document describes troubleshooting Overlay Management Protocol (OMP) failure scenarios and best practices to provide network resilience in Cisco SD-WAN.

Prerequisites

Requirements

Cisco recommends that you have knowledge of the Cisco Software Defined Wide Area Network (SD-WAN) solution.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco IOS Catalyst SD-WAN Manager aka vManage
- Cisco IOS Catalyst SD-WAN Validator aka vBond
- Cisco IOS Catalyst SD-WAN Controller aka vSmart
- vEdge Devices

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure you understand any command's potential impact.

OMP Overview

As you know, the Cisco SD-WAN Edge device only shares the routes with the Catalyst SD-WAN Controller. For a route to be valid and installed in its Forwarding table:

- The next hop Transport Locator (TLOC) must be reachable, that is, the Edge device must have a Valid route for the TLOC.

- The TLOC to which it points is active. For a TLOC to be active, an active Bidirectional Forwarding (BFD) session must be associated with that TLOC. BFD sessions are established by each device which creates a separate BFD session with each of the remote TLOCs. If a BFD session becomes inactive, the Cisco Catalyst SD-WAN Controller removes all the OMP routes that point to that TLOC from the forwarding table.
- The OMP route must be calculated as best.

Though all these statements are logical and straightforward there is a significant difference between OMP and traditional routing Protocols like Enhanced Interior Gateway Routing Protocol (EIGRP) and Open Shortest Path First (OSPF) during failure scenarios.

EIGRP Failure Scenario

In the next Network, there are three sites, namely Site1, Site3, and Site4 having routers RTR1/RTR2, RTR3, and RTR4 respectively with a single WAN Connection. The traditional routing Protocol EIGRP is run over IPSec and IP1, IP2, IP3, and IP4 are the WAN interface IP addresses in respective locations.

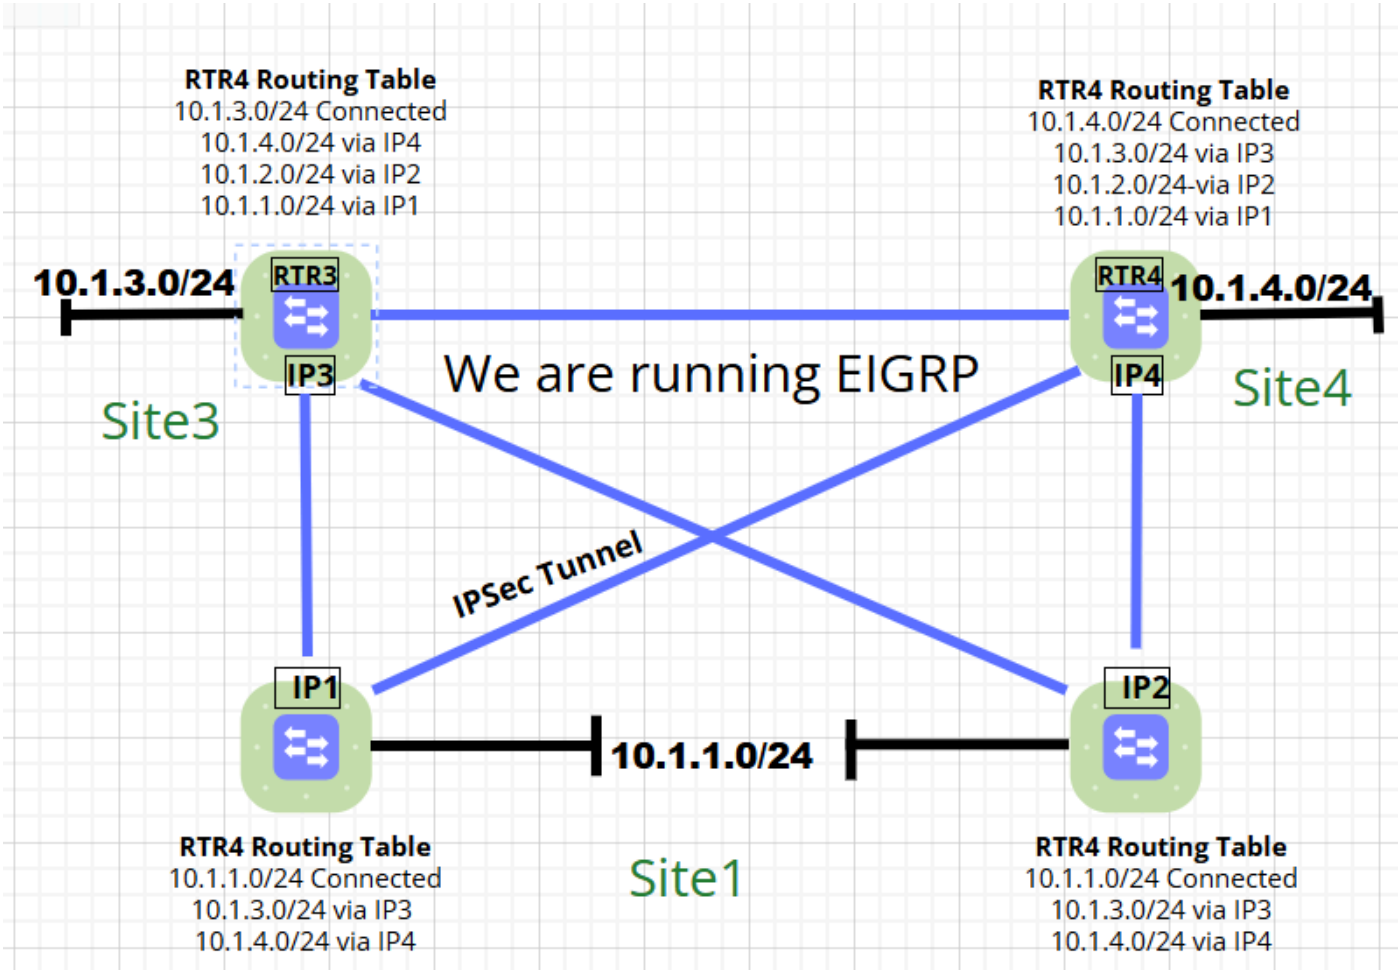

The network must be broken down, with a focus on RTR3 and RTR4 for now. On RTR3, the route to reach the 10.1.4.0/24 is via a direct tunnel between RTR3-RTR4. If the tunnel goes down, how does EIGRP react in this case? As soon as the tunnel goes down, EIGRP will run and send a query to neighboring routers for the 10.1.4.0/24 network, and based on replies received, it examines it and installs the new path for the destination in the routing table post the best path calculation.

This is a very simple explanation of the traditional routing protocol convergence process. So overall traditional routing protocols like EIGRP are capable of performing network recompute:

- When the current route to a destination goes down
- When there are no feasible successors for a destination
- When a topology change occurs

OMP Failure Scenario

The two failure scenarios for OMP are discussed here:

1. Direct Failure
2. Indirect Failure

Direct Failure

In the next Topology, there are three sites with a single transport connection.

Site	Router	Transport Locator (TLOC)	System IP	Subnet
SIte1	vEdge-1 vEdge-2	T1 T2	1.1.1.1 2.2.2.2	10.1.1.0/24
Site-3	vEdge-3	T3	3.3.3.3	10.1.3.0/23
Site-4	vEdge-4	T4	4.4.4.4	10.1.4.0/24

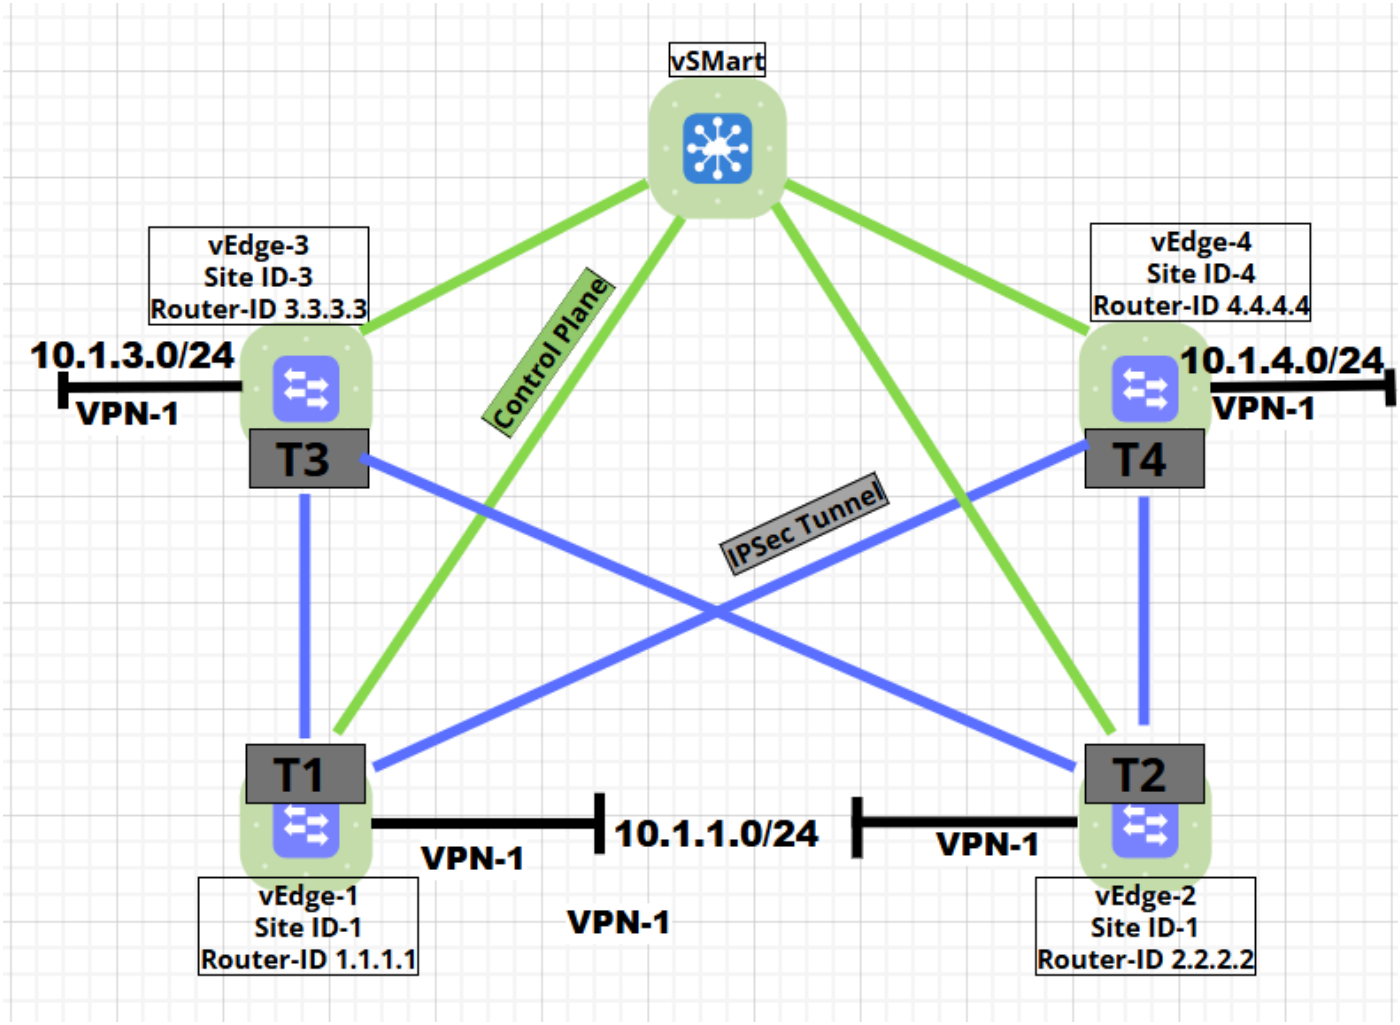

Assume that everything is set to default on the Catalyst SD-WAN Controller. The vEdge devices share the routing information directly with the Catalyst SD-WAN Controller, and the controller shares it with all vEdge devices. The next topology shows the routing table for all routers:

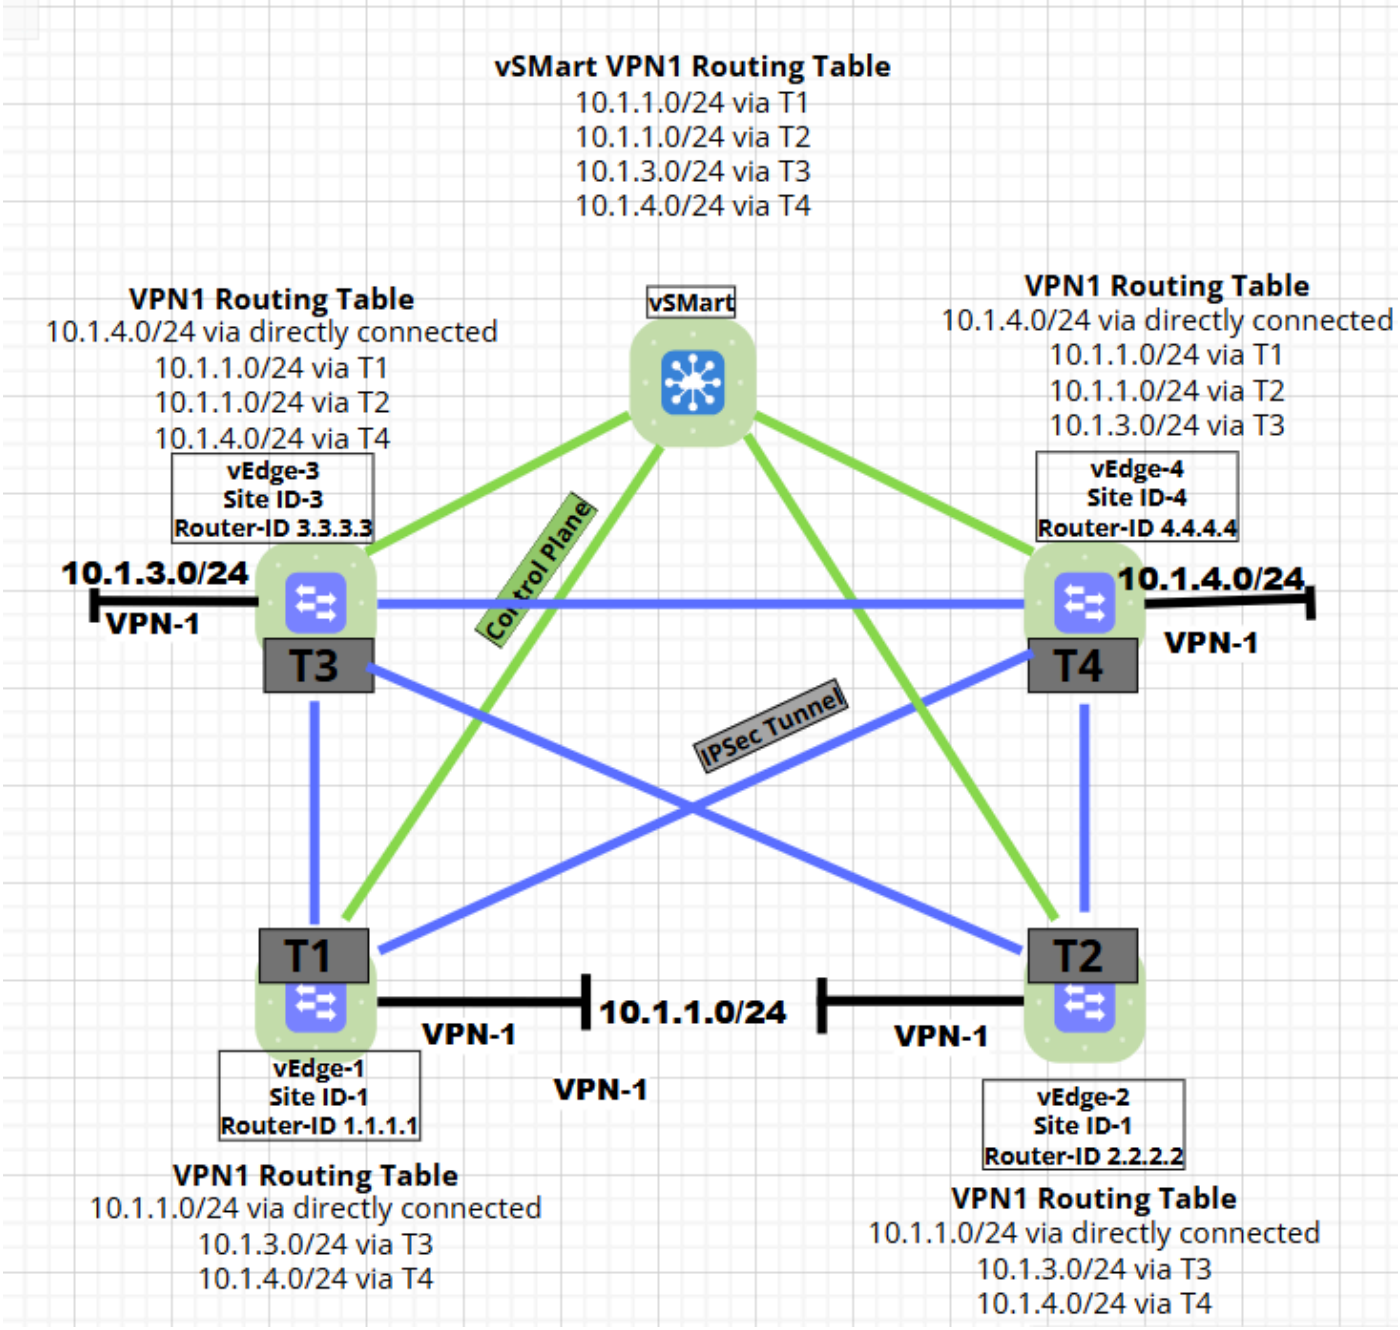

Currently, all BFD sessions are up.

```
vEdge-DC1# show bfd sessions
                                       SOURCE TLOC      REMOTE TLOC
SYSTEM IP        SITE ID  STATE       COLOR            COLOR            SOURCE IP
-----------------------------------------------------------------------------------------------------
1.1.1.1                1        up          mpls             mpls             60.1.1.1
2.2.2.2                1        up          mpls             mpls             60.1.1.1
4.4.4.4                2        up          mpls             mpls             60.1.1.1


vEdge-DC1# show omp routes vpn 20 | t
Code:
C   -> chosen
I   -> installed
Red -> redistributed
Rej -> rejected
```

```
L   -> looped
R   -> resolved
S   -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA  -> On-demand inactive
U   -> TLOC unresolved

                                        PATH                      ATTRIBUTE
VPN    PREFIX              FROM PEER    ID     LABEL    STATUS    TYPE       TLOC IP          COLOR
-------------------------------------------------------------------------------------------------------
20     10.1.1.0/24         2.2.2.2      43     1005       C,I,R          installed      1.1.1.1
                                       2.2.2.2          37     1006       C,I,R            installed
20     10.1.3.0/24         0.0.0.0      66     1005     C,Red,R        installed      3.3.3.3
20     10.1.4.0/24         2.2.2.2      45     1006       C,I,R          installed       4.4.4.4
```

If the connectivity between vEdge3 and vEdge4 is disabled, when the tunnel goes down, both vEdge3 and vEdge4 will have their BFD sessions go down as well. This causes them to mark the respective routes as 'Invalid' and 'TLOC Unresolved'. You can see that in the next output:

```
vEdge3# show bfd sessions
                                       SOURCE TLOC      REMOTE TLOC
SYSTEM IP        SITE ID  STATE       COLOR            COLOR            SOURCE IP
-----------------------------------------------------------------------------------------------------
1.1.1.1                1        up          mpls             mpls                     60.1.1.1
2.2.2.2                1        up          mpls             mpls                     60.1.1.1
4.4.4.4                4        down     mpls             mpls                     60.1.1.1


vEdge3# show omp routes vpn 20 | t
Code:
C   -> chosen
I   -> installed
Red -> redistributed
Rej -> rejected
L   -> looped
R   -> resolved
S   -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA  -> On-demand inactive
U   -> TLOC unresolved

                                        PATH                      ATTRIBUTE
VPN    PREFIX              FROM PEER    ID     LABEL    STATUS    TYPE       TLOC IP          COLOR
-------------------------------------------------------------------------------------------------------
1      10.1.1.0/24         2.2.2.2      43     1005       C,I,R             installed       1.1
                                       2.2.2.2          37     1006       C,I,R            installe
1    10.1.3.0/24           0.0.0.0      66     1005       C,Red,R         installed       3.3.3.
1      10.1.4.0/24         2.2.2.2      45     1006       Inv,U             installed       4.4
```

Indirect Failure

In order to understand 'Indirect Failure', assume the control policy is defined to change the next hop on vEdge3 for route 10.1.4.0/24 via vEdge2, and on vEdge4 the next hop for 10.1.3.0/24 is changed to vEdge1. In other words, for traffic between vEdge 3 and 4, vEdge 2 and 1 were inserted as intermediate hops. You can see that in the next diagram:

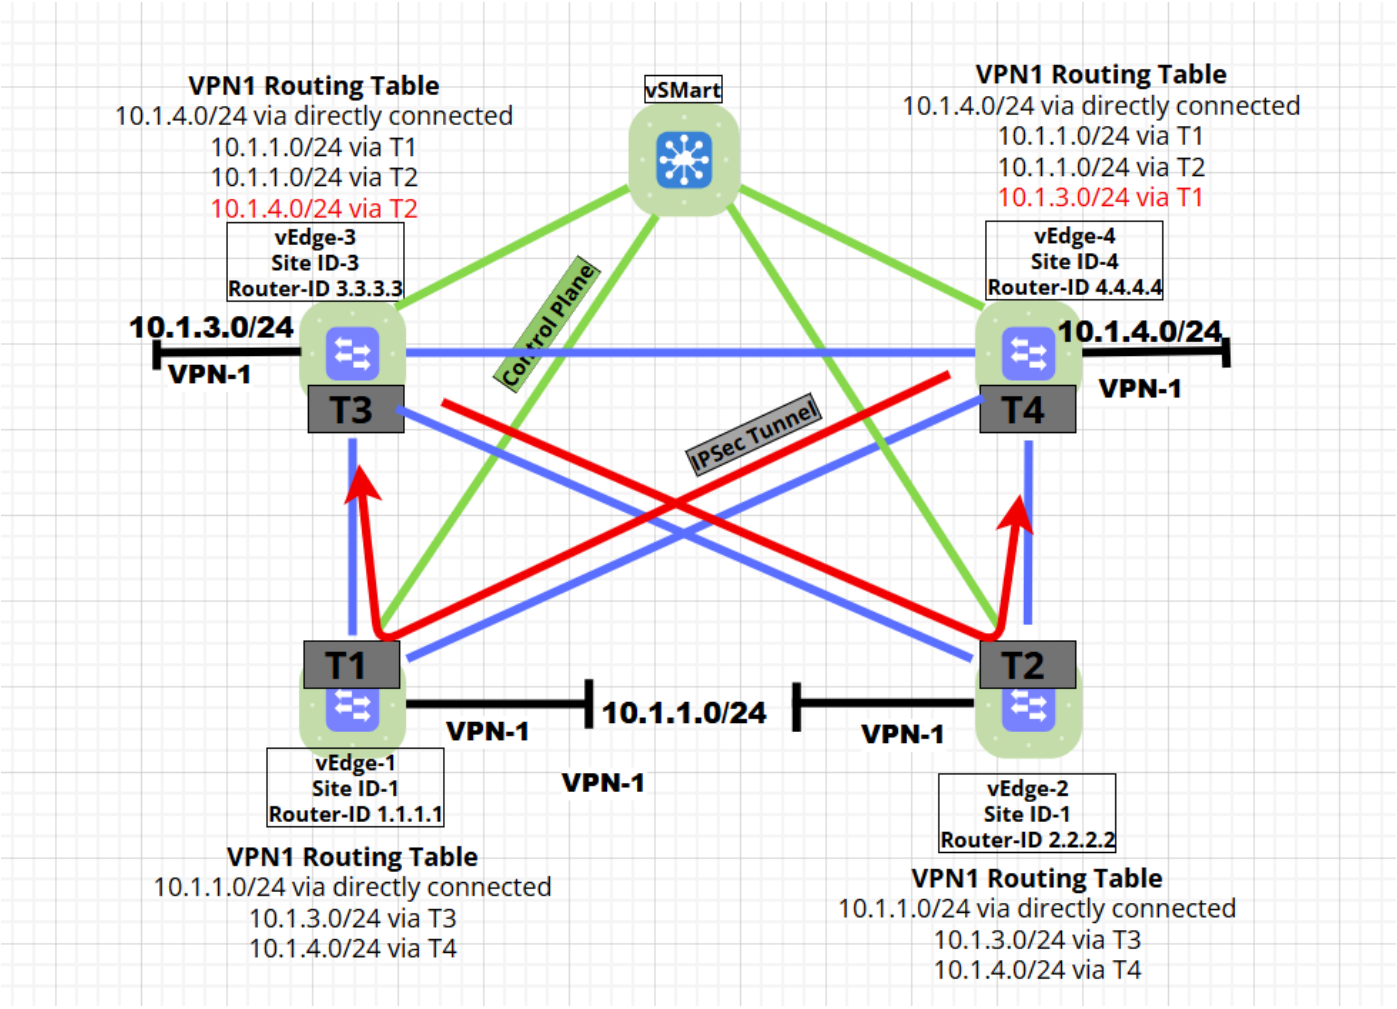

If there is a network failure that results in connectivity loss between vEdge2 and vEdge4, while the overlay tunnel between T2-T4 is down, vEdge3 still has a valid route for 10.1.4.0 via T2. Hence it sends traffic to vEdge2. vEdge2 does not have a valid Tunnel with vEdge4, so routes will no longer be active on it therefore dropping the traffic.

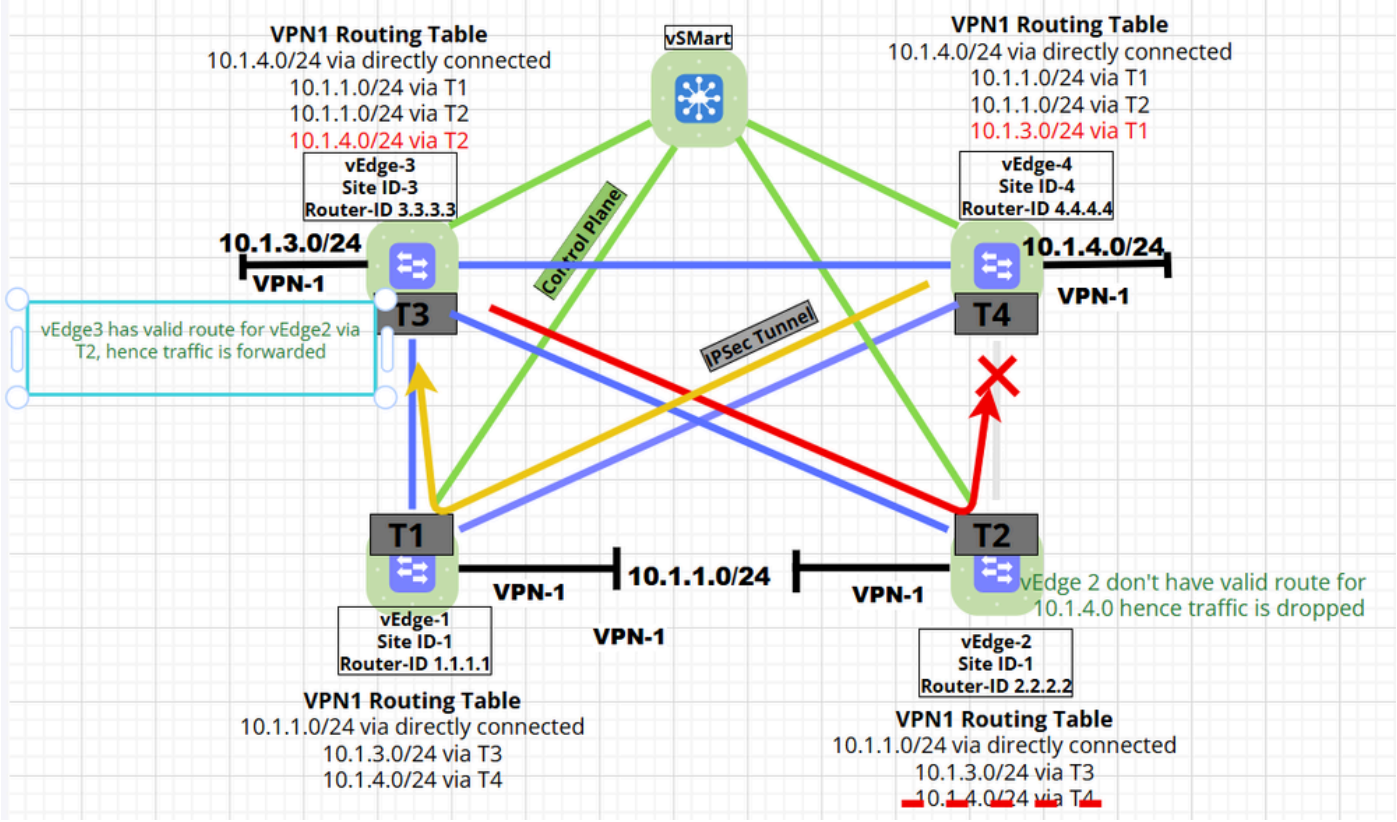

Based on the earlier logs and tests, it can be concluded that:

- With OMP, there is no auto-discovery of routing peers and next-hops
- There is no topology recalculation when a tunnel goes down
- The OMP routes to a destination prefix never change when a tunnel goes down. The only change that happens is reachability to the next hop, that is, TLOC.
- In case of direct overlay failure, a tunnel redundancy with Multiple Tunnels to the same destination must be provided.
- Extra care must be taken while introducing intermediate hop/hops in the overlay path and tunnel redundancy must be provided in order to avoid the lack hole of traffic.

So now you know that by default, OMP does not recompute or reroute on overlay failure. In order to overcome this problem, you can enable a feature called 'TLOC-Action' via a Control Policy.

TLOC-Action

- In Cisco SD-WAN, a 'TLOC Action' within a control policy allows the insertion of an intermediate hop (TLOC) to be used for traffic forwarding while maintaining visibility into the complete path from source to destination. This means setting the TLOC action option enables the Cisco Catalyst SD-WAN Controller to perform end-to-end tracking of the path to the ultimate destination device. If that path goes down, the controller informs the WAN edge routers that received this OMP route.
- It provides a backup path in case of primary link failure, thus enhancing network resilience and fault tolerance within the SD-WAN overlay network. It is a way to control how traffic is routed through the network by manipulating the TLOCs used for reaching a destination.
- When a TLOC Action is defined in a policy, it instructs the SD-WAN controller to insert an intermediate TLOC into the route calculation, meaning traffic will first go to this specified 'backup' location before reaching the final destination if necessary.
- This is particularly useful for scenarios where you want to ensure connectivity even if a primary link goes down, by automatically rerouting traffic through a different path (via the specified TLOC).

On the next topology, let's focus on **vEdge2, vEdge3, and vEdge4** to understand it better. Currently, no policy is defined and data traffic for 10.1.4.0/24 on vEdge3 is traversing over a direct tunnel between **T3** and **T4**.

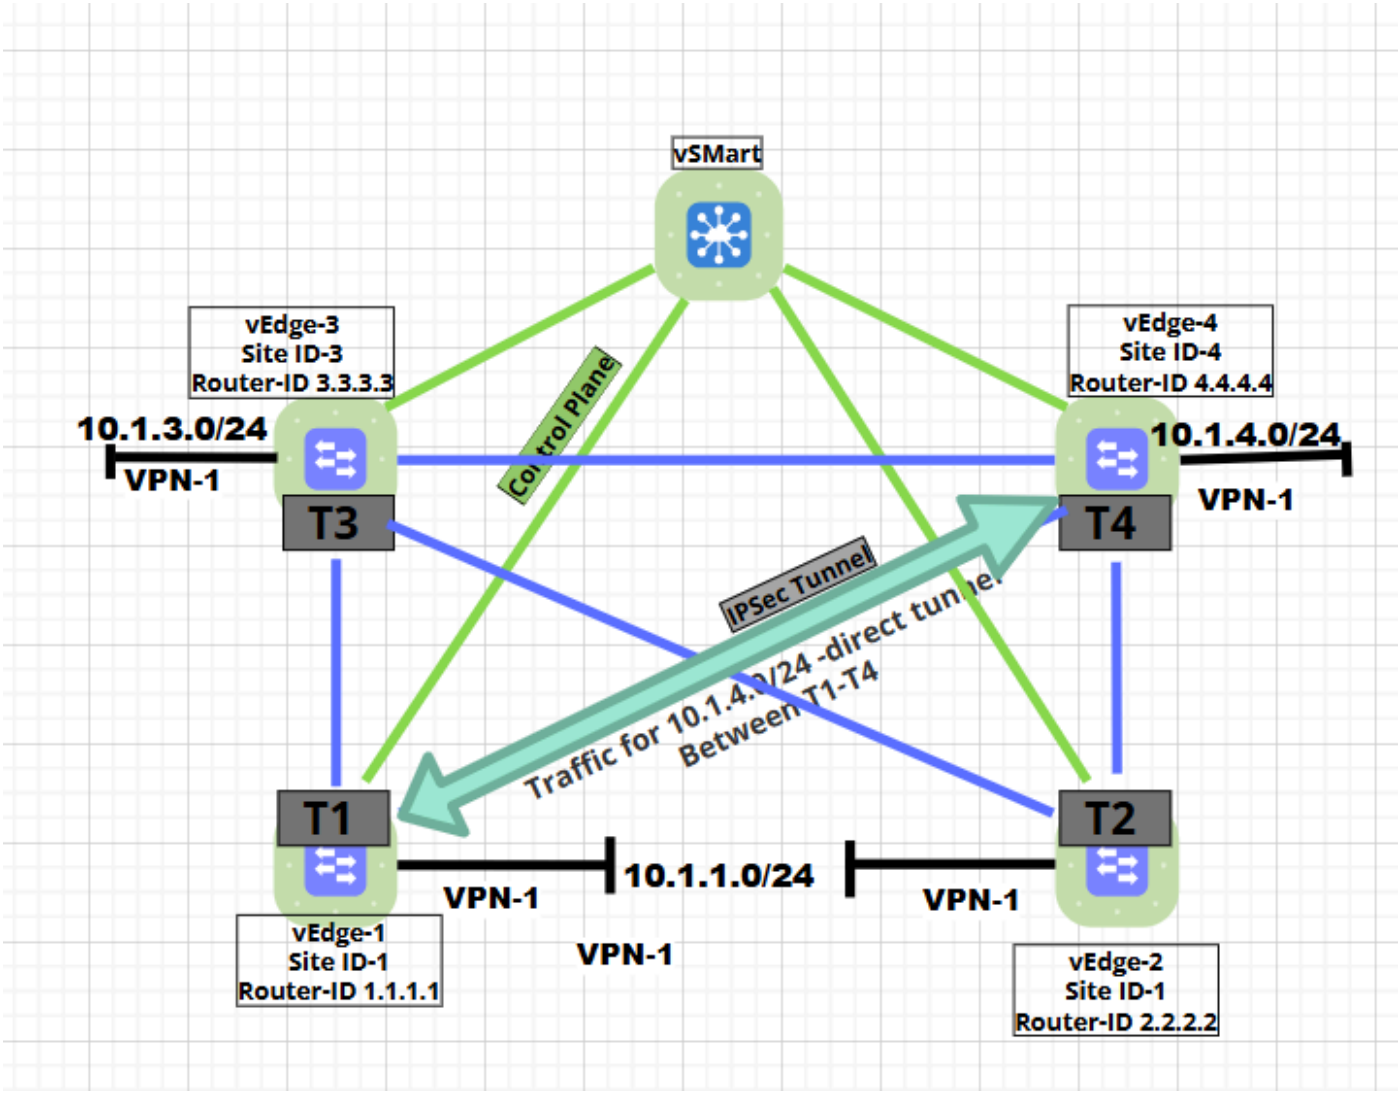

In order to provide fault tolerance and network resilience, the Control policy is configured to reroute traffic through a different path (via the specified TLOC).

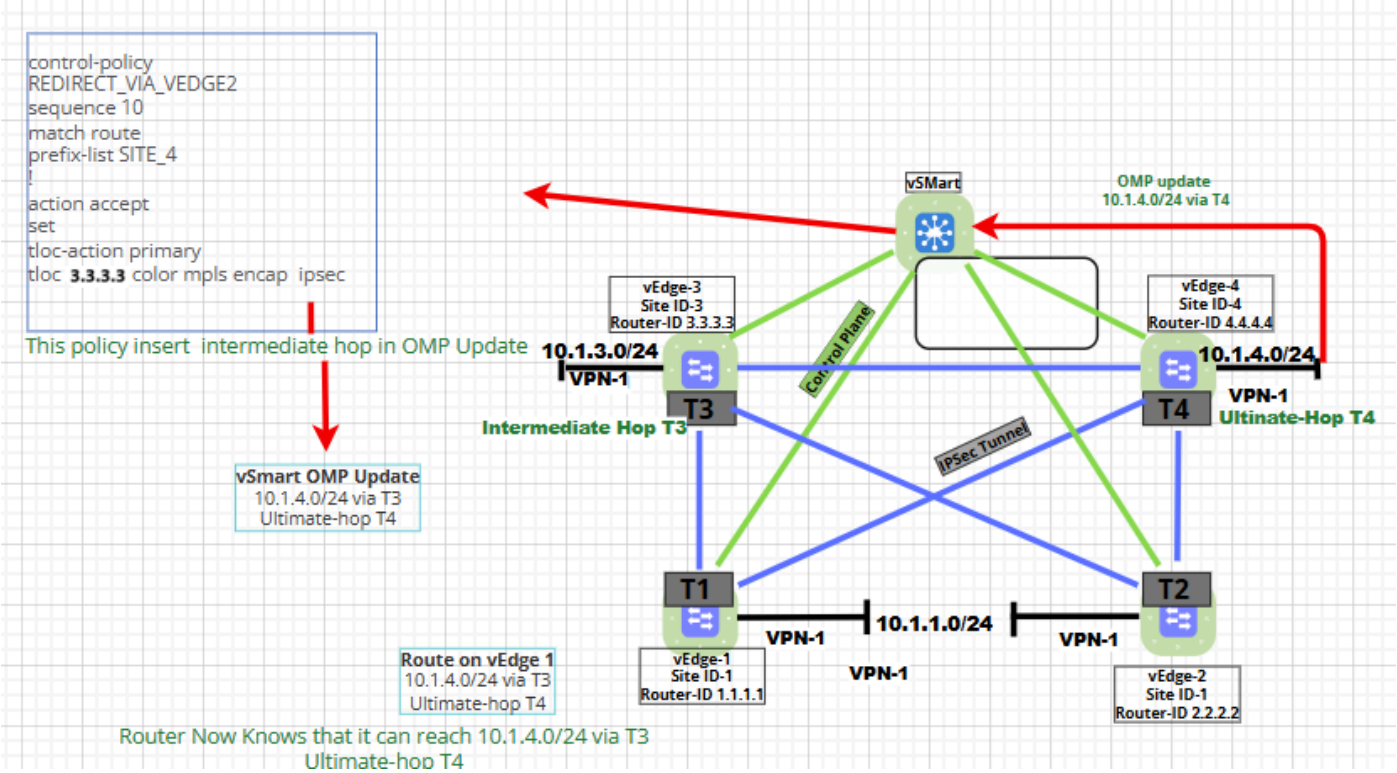

- vEdge4 sends OMP update for its directly connected network 10.1.4.0/24 with next-hop T4 to Catalyst SD-WAN Controller as '10.1.4.0/24 via T4'.
- This route matches a control policy configured on the SD-WAN Controller and sets a new TLOC and TLOC-Actions as per the policy defined on it, that is, It inserts the new 'intermediate TLOC'.
- The controller advertises the OMP route to vEdge1 with two next-hops now - intermediate TLOC (T3, 3.3.3.3) and ultimate TLOC (next-hop-T4 of the original route). This gives intelligence to vEdge1 that the destination prefix 10.1.4.0/24 is reachable via T2 and via T4.

Now based on the TLOC-Action defined vEdge1 forward the traffic for 10.1.4.0/24. So you can define these four types of TLOC-Actions in control plane Policy:

1. **Strict (default) -** The 'TLOC-Action strict' defines that the traffic between vEdge1 and vEdge4 must be via T3 (intermediate hop) and traffic must drop if the tunnel between vEdge1 and vEdge4 goes down.
2. **Primary -** The 'TLOC-Action primary' defines that the traffic between vEdge1 and vEdge4 goes through the intermediate hop T3 (3.3.3.3) and if this overlay tunnel goes down, SD-WAN Controller informs vEdge1 and traffic routed over the direct tunnel to T4.
3. **Backup -** The 'TLOC-Action backup' defines that the traffic between vEdge1 and vEdge4 goes directly to the ultimate LOC (next-hop -T4 of the original route) and if the direct overlay tunnel between vEdge1 and vEdge4 goes down, the SD-WAN Controller informs vEdge1 and traffic goes via Intermediate hop T3.
4. **Equal-Cost Multi-Path (ECMP) -** The 'TLOC-Action ECMP' specifies that in normal circumstances, the communication between vEdge1 and vEdge4 is load-balanced through the intermediate hop T3 and Ultimate-hop T4.