Validate Intelligent Traffic Director Using Source IP Method

Contents

Introduction

This document describes the way to check both the control plane and data plane of the Intelligent Traffic Director (ITD) in the Nexus 9K platform.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Nexus NX-OS Software.
- Access List (ACL).
- IP Service Level Agreement (IP SLA).
- Policy-Based Routing (PBR).
- Intelligent Traffic Director (ITD).

Components Used

The information in this document is based on Cisco Nexus 9000 with NXOS version 10.2(5).

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 that you understand the potential impact of any command.

Configure

ITD is an intelligent, hardware-based, multi-terabit solution that allows you to build a scalable architecture for Layer 3 and Layer 4 traffic distribution, load balancing, and redirection.

Network Diagram

- Computers: 4
- Servers: 4
- Layer 2 switches: 2
- Nexus 9K switches: 1

> **Note**: The layer 2 switches merge the connection between servers/computers and Nexus 9K switches, where ITD is running.

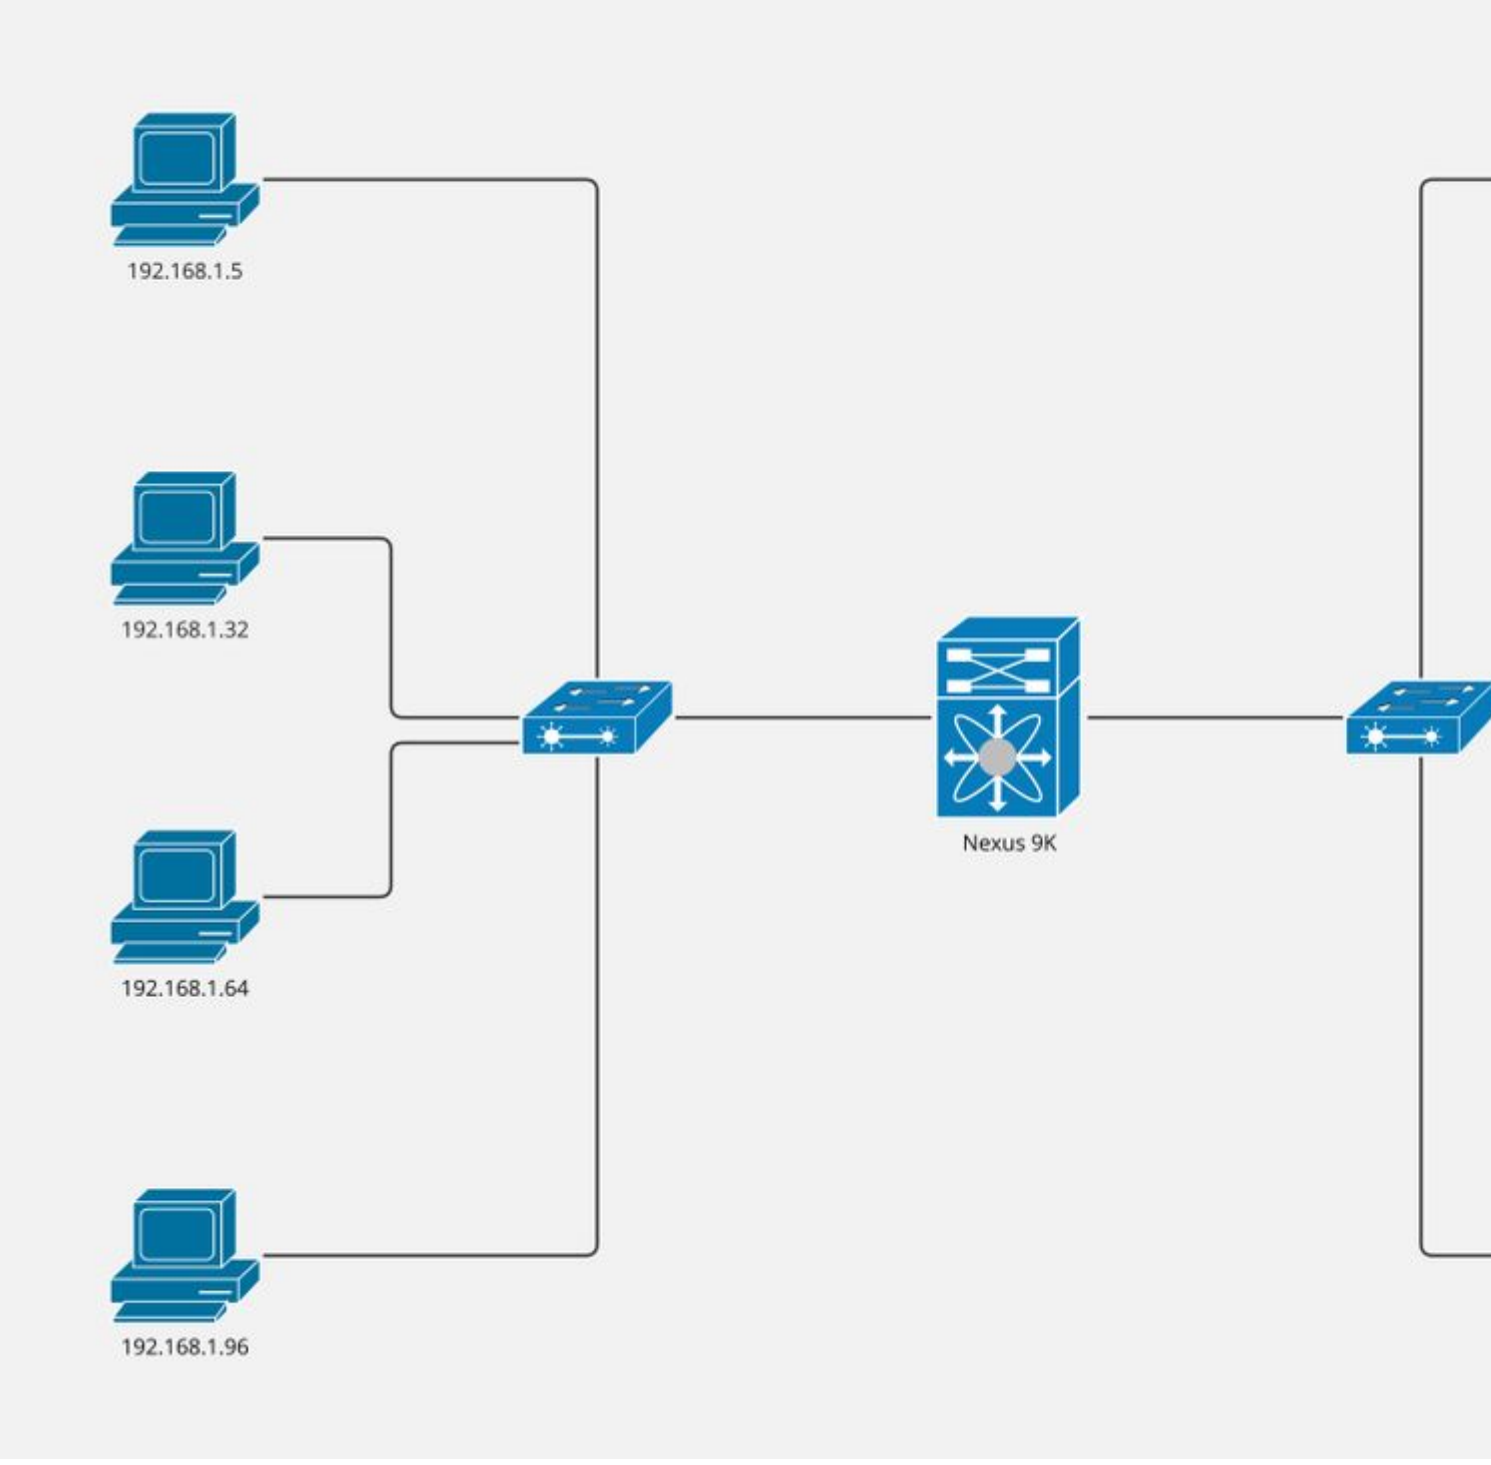

Configurations

This configuration is implemented on Nexus 9K switch and covers the ITD part only. Computers are part of VLAN 5 and servers are part of VLAN 10, while Nexus 9K is the gateway of both VLANs.

Step 1. Enable required features.

```
feature itd
feature sla sender
feature sla responder
feature pbr
```

Step 2. Define device-group, pointing to nodes, and indicate the probe method.

```
itd device-group CiscoGroup
  probe icmp frequency 5 timeout 5 retry-down-count 2 retry-up-count 2
  node ip 10.1.1.5
  node ip 10.1.1.6
  node ip 10.1.1.7
  node ip 10.1.1.8
```

Step 3. Define service. Users can indicate either virtual ip or include an access list.

> **Note**: NX-OS divides the last octet by the number of buckets indicated. For this example, we divide 256 by 8. Hence, each bucket includes 32 hosts.

```
Bucket 1: From x.x.x.0 to x.x.x.31
Bucket 2: From x.x.x.32 to x.x.x.63
Bucket 3: From x.x.x.64 to x.x.x.95
Bucket 4: From x.x.x.96 to x.x.x.127
Bucket 5: From x.x.x.128 to x.x.x.159
Bucket 6: From x.x.x.160 to x.x.x.191
Bucket 7: From x.x.x.192 to x.x.x.223
Bucket 8: From x.x.x.224 to x.x.x.255
```

```
itd CiscoService
  device-group CiscoGroup
  virtual ip 192.168.255.1 255.255.255.255 advertise enable
  ingress interface Vlan5
  failaction node reassign
  load-balance method src ip buckets 8
  no shut
```

Step 4. Enable statistics for defined service.

```
itd statistics CiscoService
```

Verify

1. Confirm ITD service is active and node status is ok.

```
Nexus# show itd CiscoService brief

Legend:
 C-S(Config-State): A-Active,S-Standby,F-Failed
 ST(Status): ST-Standby,LF-Link Failed,PF-Probe Failed,PD-Peer Down,IA-Inactive,SH-Shut,HD-Hold-down

Name           LB Scheme  Status   Buckets    Interface
-------------- ---------- -------- ---------  --------------
CiscoService   src-ip     ACTIVE   8          Vlan5

Source Interface
----------------


Device Group                                        Probe  Port  VRF
-------------------------------------------------- ----- ------ ---------
CiscoGroup                                          ICMP
Virtual IP                              Netmask/Prefix Protocol     Port
-------------------------------------------------- ------------ ----------
192.168.255.1 / 255.255.255.255                        IP           0



Node  IP                 Cluster-id C-S WGT Probe Port     Probe-IP   STS
------------------------ ---------- -- --- ---- ----- --------------- --
1               10.1.1.5             A   1 ICMP                       OK
2               10.1.1.6             A   1 ICMP                       OK
3               10.1.1.7             A   1 ICMP                       OK
4               10.1.1.8             A   1 ICMP                       OK
```

2. Confirm route policy was created and associated with the ingress interface.

> **Note**: Route-policy is applied in running-config (under ingress interface) until NXOS 10.1(2), starting NXOS 10.2(1) you can find route-policy associated in the system with the **show ip policy** command.

```
Nexus# show ip policy
Interface                Route-map            Status     VRF-Name
Vlan5                    CiscoService_itd_pool Active     default
```

3. Confirm the route-map was generated properly, associated with the expected nodes, and indicates the

tracking (IP SLA).

> **Note**: Is expected to have one route-map entry per bucket.
> Is expected to have one track per node.

```
Nexus# show route-map dynamic CiscoService_itd_pool
route-map CiscoService_itd_pool, permit, sequence 10
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_1
  Set clauses:
    ip next-hop verify-availability 10.1.1.5 track 2 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 11
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_2
  Set clauses:
    ip next-hop verify-availability 10.1.1.6 track 3 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 12
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_3
  Set clauses:
    ip next-hop verify-availability 10.1.1.7 track 4 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 13
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_4
  Set clauses:
    ip next-hop verify-availability 10.1.1.8 track 5 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 14
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_5
  Set clauses:
    ip next-hop verify-availability 10.1.1.5 track 2 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 15
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_6
  Set clauses:
    ip next-hop verify-availability 10.1.1.6 track 3 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 16
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_7
  Set clauses:
    ip next-hop verify-availability 10.1.1.7 track 4 [ UP ]  force-order
route-map CiscoService_itd_pool, permit, sequence 17
  Match clauses:
    ip address (access-lists): CiscoService_itd_vip_1_bucket_8
  Set clauses:
    ip next-hop verify-availability 10.1.1.8 track 5 [ UP ]  force-order
```

4. Confirm access list was generated properly and has the expected ip match condition.

Note: ACLs are generated in running-config until NXOS 9.3(2), starting NXOS 9.3(3) we can find ACLs in system using **show ip access-list dynamic** command.

```
Nexus# show ip access-lists CiscoService_itd_vip_1_bucket_1 dynamic

IP access list CiscoService_itd_vip_1_bucket_1
        10 permit ip 1.1.1.0 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_2 dynamic

IP access list CiscoService_itd_vip_1_bucket_2
        10 permit ip 1.1.1.32 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_3 dynamic

IP access list CiscoService_itd_vip_1_bucket_3
        10 permit ip 1.1.1.64 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_4 dynamic

IP access list CiscoService_itd_vip_1_bucket_4
        10 permit ip 1.1.1.96 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_5 dynamic

IP access list CiscoService_itd_vip_1_bucket_5
        10 permit ip 1.1.1.128 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_6 dynamic

IP access list CiscoService_itd_vip_1_bucket_6
        10 permit ip 1.1.1.160 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_7 dynamic

IP access list CiscoService_itd_vip_1_bucket_7
        10 permit ip 1.1.1.192 255.255.255.31 192.168.255.1/32
513E-A-15-C9336C-FX-2-1# show ip access-lists CiscoService_itd_vip_1_bucket_8 dynamic

IP access list CiscoService_itd_vip_1_bucket_8
        10 permit ip 1.1.1.224 255.255.255.31 192.168.255.1/32
```

5. Ensure access lists are programmed in TCAM.

```
Nexus# show system internal access-list input entries detail | begin "VLAN 5"
               VLAN 5  :
               =========

INSTANCE 0x0
---------------

  Tcam 1 resource usage:
  ----------------------
  LBL B = 0x1
   Bank 1
   ------
     IPv4 Class
       Policies: PBR(CiscoService_itd_vip_1_bucket_8)
       Netflow profile: 0
       Netflow deny profile: 0
       Entries:
```

```
        [Index] Entry [Stats]
        --------------------
  [0x0000:0x0002:0x0002] permit ip 0.0.0.0/0 224.0.0.0/4  routeable 0x1  [0]
  [0x0002:0x0004:0x0004] permit ip 0.0.0.0/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x0003:0x0005:0x0005] permit ip 0.0.0.32/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x0004:0x0006:0x0006] permit ip 0.0.0.64/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x0005:0x0007:0x0007] permit ip 0.0.0.96/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000b:0x000d:0x000d] permit ip 0.0.0.128/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000c:0x000e:0x000e] permit ip 0.0.0.160/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000d:0x000f:0x000f] permit ip 0.0.0.192/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000e:0x0010:0x0010] permit ip 0.0.0.224/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000f:0x0011:0x0011] permit ip 0.0.0.0/0 0.0.0.0/0  routeable 0x1  [0]


  L4 protocol cam entries usage: none

  No mac protocol cam entries are in use


INSTANCE 0x1
---------------

  Tcam 1 resource usage:
  ----------------------
  LBL B = 0x1
   Bank 1
   ------
     IPv4 Class
       Policies: PBR(CiscoService_itd_vip_1_bucket_8)
       Netflow profile: 0
       Netflow deny profile: 0
       Entries:
         [Index] Entry [Stats]
         ---------------------
  [0x0000:0x0002:0x0002] permit ip 0.0.0.0/0 224.0.0.0/4  routeable 0x1  [0]
  [0x0002:0x0004:0x0004] permit ip 0.0.0.0/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x0003:0x0005:0x0005] permit ip 0.0.0.32/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x0004:0x0006:0x0006] permit ip 0.0.0.64/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x0005:0x0007:0x0007] permit ip 0.0.0.96/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000b:0x000d:0x000d] permit ip 0.0.0.128/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000c:0x000e:0x000e] permit ip 0.0.0.160/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000d:0x000f:0x000f] permit ip 0.0.0.192/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000e:0x0010:0x0010] permit ip 0.0.0.224/0.0.0.224 192.168.255.1/32  routeable 0x1  [0]
  [0x000f:0x0011:0x0011] permit ip 0.0.0.0/0 0.0.0.0/0  routeable 0x1  [0]


  L4 protocol cam entries usage: none

  No mac protocol cam entries are in use
```

6. Send traffic to the virtual IP address and confirm counters (number of packets) increase for the expected node.

```
Nexus# show itd CiscoService statistics

Service                       Device Group                          VIP/mask
-------------------------------------------------------------------------------------------------------------
CiscoService                      CiscoGroup                    192.168.255.1 / 255.255.255.255
```

```
Traffic Bucket                          Assigned to                 Mode                Orig
---------------                         --------------              -----               ----
CiscoService_itd_vip_1_bucket_1         10.1.1.5                    Redirect            10.1
CiscoService_itd_vip_1_bucket_5         10.1.1.5                    Redirect            10.1

Traffic Bucket                          Assigned to                 Mode                Orig
---------------                         --------------              -----               ----
CiscoService_itd_vip_1_bucket_2         10.1.1.6                    Redirect            10.1
CiscoService_itd_vip_1_bucket_6         10.1.1.6                    Redirect            10.1

Traffic Bucket                          Assigned to                 Mode                Orig
---------------                         --------------              -----               ----
CiscoService_itd_vip_1_bucket_3         10.1.1.7                    Redirect            10.1
CiscoService_itd_vip_1_bucket_7         10.1.1.7                    Redirect            10.1

Traffic Bucket                          Assigned to                 Mode                Orig
---------------                         --------------              -----               ----
CiscoService_itd_vip_1_bucket_4         10.1.1.8                    Redirect            10.1
CiscoService_itd_vip_1_bucket_8         10.1.1.8                    Redirect            10.1
```