Cisco Intelligent Traffic Director Deployment Guide with Cisco Firepower

Deployment Guide

November 2016

Contents

Introduction

Cisco® Intelligent Traffic Director (ITD) is a multi-terabit Layer 4 load-balancing and traffic steering solution on the Cisco Nexus® 5000, 6000, 7000, and 9000 Series Switches. The traffic director is application-specific integrated circuit (ASIC) based and provides scalable load distribution of traffic to a group of servers or service appliances.

In scalable Next Generation Intrusion Prevention Solutions (NGIPS), the traffic director can balance loads to Cisco Firepower® 7000 and 8000 series appliances configured as stand-alone nodes. When scaling NGIPS, the largest IPS throughput is achieved by creating a Firepower 8300 stack. The largest Firepower 8300 stack is an 8390 rated at 60 Gbps. To reach higher speeds, multiple Firepower 8300 stacks must be utilized.

The combination of Intelligent Traffic Director and Firepower allows for both scale up and scale out architecture. The Firepower 8300 series can be scaled up from an 8350 to an 8390. Additional Firepower nodes can be added dynamically utilizing Intelligent Traffic Director's non-disruptive node insertion for scale out.

This document discusses designs that utilize ITD's symmetrical flow load balancing capabilities paired with the Firepower appliances to achieve both scale up as well as scale out architecture.

Deployment Methods

Traffic can be distributed to NGIPS appliances with Intelligent Traffic Director using the follow topology:

- Intelligent Traffic Director in a dual-VDC sandwich mode with virtual port-channels (vPC) with two Cisco Nexus 7700 series switches. This design uses two virtual device contexts (VDC) per switch with each VDC having an interface connecting to each Firepower node. The Firepower devices filter traffic traversing the two VDCs.

This design can also be used with separate physical Cisco Nexus 7000/7700 switches instead of VDCs if desired. The Cisco Nexus 7700 switches in this design are deployed in L3 vPC mode. For purposes of this document, "ITD service node" will refer to a Cisco Nexus 7000/7700 switch or a VDC running on a Nexus 7000/7700 switch.

The design as discussed in this document use Firepower 8300 devices operating as NGIPS in a Layer 2 inline mode. Each Firepower 8300 will function as a service node to ITD. For purposes of this document, "Firepower node" will refer to 8350 through 8390 appliances, regardless of how many Firepower devices are deployed in a stack configuration.

ITD natively requires a Layer 3 next-hop to which it redirects traffic. As such, the ITD service node devices will be creating Layer 2 adjacency over vPC to establish L3 next hops. In addition, Layer 3 IGP adjacency traversing the service-nodes is required to exchange route tables across the dual VDC sandwich design.

These designs are discussed here using the Cisco Nexus 7000 series switches as an example. In particular, Cisco Nexus 7700 platform switches running Cisco NX-OS software release 7.2(1)D1(1). The Cisco Nexus 5000, 6000, and 9000 Series switches currently are not in the scope of this design as they do not support L3 adjacency over vPC at this time, which is a required feature.

The configurations presented in the scenarios here are only relevant snippets that are unique to each method.

The only configurations that are presented in this document are the portions relevant to Nexus ITD plus Firepower NGIPS. Full configurations and configurations for other features of the device (for example, vPC) are beyond the scope of this document.

In the figures in this document, blue links and labels typically refer to Firepower node “inside” connections and configurations, and red links and labels refer to Firepower node “outside” connections and configurations.

Design and Deployment Considerations

Number of Intelligent Traffic Director Services

An Intelligent Traffic Director service is an instance of the feature that defines the traffic distribution parameters for a particular direction of the traffic flow. If both directions of a flow need to be redirected, you typically need to configure two ITD services: one for the forward traffic flow and one for the return traffic flow. In order for Firepower services such as NGIPS, Advanced Malware Protection (AMP), and application visibility control (AVC) to work both forward and return traffic must pass through the same Firepower node to maintain stateful inspection.

Logical Overview

The logical flow of the traffic can be viewed as three relevant components. The example in Figure 1 illustrates traffic flows in both directions for reference.

The first component is ingress traffic into the ITD Forwarding-Service. This deployment guide assumes as best practice that there will be two Layer 3 routers / switches forwarding traffic to the ITD service in a mesh configuration. Layer 3 ECMP will be used to load balance traffic to the ITD Nexus service nodes. The importance of this component will be covered in detail in the next section entitled, "Flow Symmetry."

The second component is ITD in a dual-VDC sandwich mode with vPC with two Nexus 7700 Series switches. The dual-VDC sandwich design allows traffic to be rerouted in the event of a failure in a Nexus 7700 switch, Firepower service node, or an individual link failure. It is a key component needed for high availability.

The third component is the Firepower 8300 service node(s). These nodes are deployed with four 40GE interfaces in one logical group. Two interfaces are dedicated for Outside traffic and two for Inside traffic.

Figure 1. Cisco Intelligent Traffic Director and Cisco Firepower Deployment: Logical View

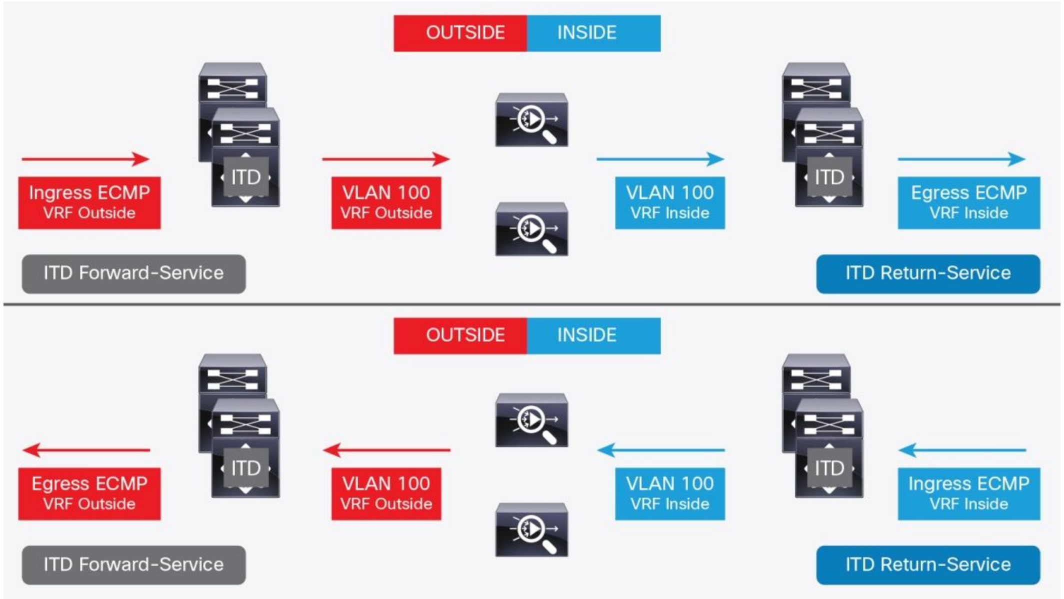

In the logical view in Figure 1, traffic that is originating from Outside to Inside will use an L3 ECMP hash to distribute flows evenly across the two Nexus service nodes running the ITD Forward-Service. The ITD Forward-Service will then forward traffic across vPC VLAN 100 which is carried to a single service-node. The service-node will then transparently forward the traffic out on vPC VLAN 100 to the Nexus 7700 switches running the ITD Return-Service nodes. Last, the Nexus 7700 switches will forward traffic out using L3 ECMP.

Figure 2. Flow Symmetry Considerations with Cisco ITD and Cisco Firepower Deployment

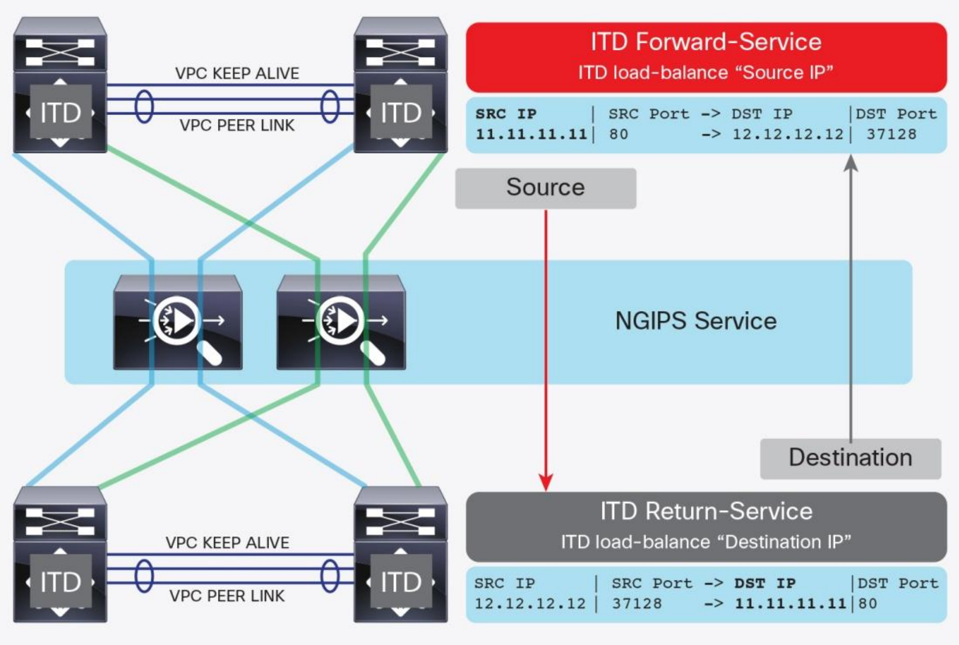

```
SRC IP     | SRC Port -> DST IP     |DST Port
11.11.11.11| 80       -> 12.12.12.12| 37128
```

```
SRC IP      | SRC Port -> DST IP     |DST Port
12.12.12.12 | 37128    -> 11.11.11.11|80
```

Flow Symmetry

NGIPSs typically inspect traffic flows in both forward and return directions. Due to the stateful nature of the inspection, flow symmetry must be maintained during normal operation. Flow symmetry can be achieved using inherent IP address persistence and the deterministic nature of the ITD algorithms. ITD will use IP SA in one direction and IP DA in the opposite direction. This ensures that flows are hashed to the same NGIPS node.

Typical ITD configuration with NGIPS uses one ITD service for the forward flow and one ITD service for the return flow. Configuring these two services in such a way that the value of the load-balance parameter remains the same for both services helps ensure that flow symmetry is maintained.

In Figure 2, the source IP address of the forward flow and the destination IP address of the reverse flow remain constant. Choosing appropriate load-balance parameters for the ITD Service that remain constant helps ensure flow symmetry in both directions.

Scaling the Architecture

The combination of Intelligent Traffic Director and Firepower NGIPS allows the NGIPS architecture to scale in two directions--up and out.

Scale Up

The Firepower 8300 itself allows a scale up design for NGIPS. The 8300 series appliances can increase the amount of traffic inspection by stacking between one and four 8300 devices. The stack uses their combined resources in a single, shared, configuration.

Figure 3. Scale Up Methodology

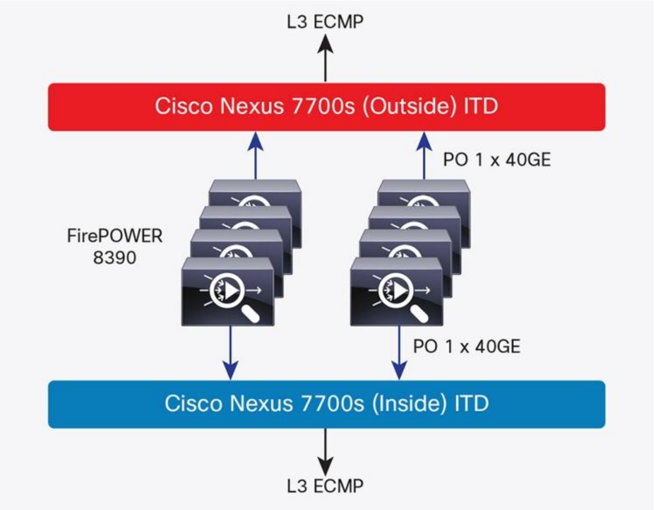
L3 ECMP
Cisco Nexus 7700s (Outside) ITD
PO 1 x 40GE
FirePOWER 8390
PO 1 x 40GE
Cisco Nexus 7700s (Inside) ITD
L3 ECMP

In the stack, one device is designated as the primary which is connected to the outside and inside network segments. The secondary devices are cabled to the primary and provide additional resources. A four node stack configuration is referenced as an 8390 appliance and can provide 60 Gbps of NGIPS with all services.

- The Firepower 8300 appliance supports both 10GE and 40GE interfaces. Cisco recommends using 40GE network modules available for the 8300 appliance.
- Pre-provisioning your 8300 primary device with four 40GE interfaces allows you to cable the external connectivity once. Pre-provisioning also eliminates the need to add additional external network capacity as you grow the stack size and consequentially the throughput.

Scale Out

Scaling out is achieved by deploying multiple Firepower 8300 service nodes (Figure 4). Utilizing the ITD service, it is possible to scale the number of Firepower 8300 appliance nodes out to the theoretical limits of the Nexus 7700 chassis. There are several design considerations that should be taken into account.

- Though it is possible to use either 10GE or 40GE interfaces in this design, consideration must be given if 10GE interfaces are chosen. To reach the maximum throughput of a Firepower 8390 a large number of 10GE interfaces will be required for the Outside and Inside port-channels. In this design we used the Nexus 7700 F3 40GE line cards and the Firepower 40GE 2 port NetMod. This simplifies the design, eases troubleshooting, and achieves the maximum theoretical scale.
- Intelligent Traffic Director's default behavior for load balancing is to provide a fairly equal traffic volume based on a sufficiently large Source Address / Destination Address (SA/DA) sample size. It is recommend that all 8300 appliances be of equal size. That is, all nodes in an ITD service should be of the same model type (8350, 8360, 8370, 8390).

- The number of ITD service nodes is recommended to be deployed as a power of 2 (e.g. 2, 4, 8, 16, …) in order to achieve equality in the traffic distribution of the nodes. More detail will be provided later in the section "ITD Load Balancing".

Figure 4. Scale Out Methodology

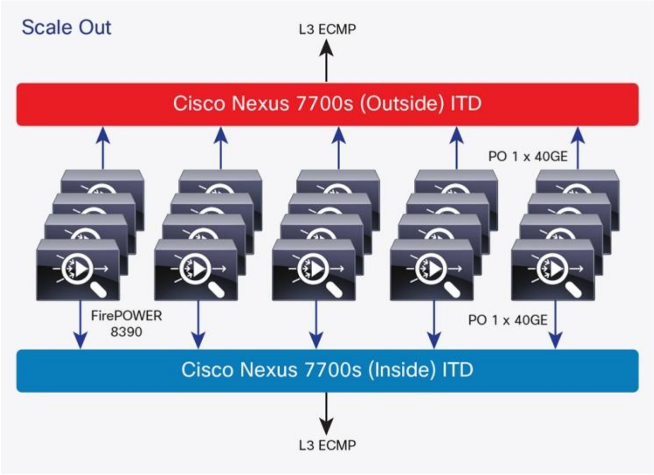
Scale Out
L3 ECMP
Cisco Nexus 7700s (Outside) ITD
PO 1 x 40GE
FirePOWER
8390
PO 1 x 40GE
Cisco Nexus 7700s (Inside) ITD
L3 ECMP

Design and Deployment Technical

ITD: Virtual Port Channels and L3

In this design, two Nexus 7700 Series Switches were used to create a dual-VDC sandwich design as seen in Figure 5. This design achieves both switch and link redundancy, however, this design cannot control how ingress and egress flows will be delivered in to the dual-VDC sandwich. With multiple Firepower nodes complexity in maintaining symmetrical traffics flows increases as multiple paths exist in the dual-VDC sandwich.

To compensate for the inability to control which switch ingress flows will pass through for forwarding and return traffic, VPCs will be used to normalize asynchronous flows.

As seen in Figure 5, the Outside VPC service nodes (SW1-VDC1 & SW2-VDC1) will create VPC 101 illustrated in blue to Firepower node 1. A mirrored VPC 101 will be created with the Inside VPC service nodes (SW1-VDC2 & SW2-VDC2). Additionally, VLAN 101 is created with SVI on each ITD node. These SVI will be used as the target in the ITD node configuration and destination of ITD probes.

This template will form the basis of connectivity for each additional Firepower node as shown with Firepower node2, VLAN 102, and VPC 102 illustrated in green. So for for each Firepower node, both a dedicated VPC and VLAN are created. In addition, each, each ITD service node would create OSPF adjacencies with its ITD peer nodes.

ITD is a layer 4 load-balancing and traffic steering service. As such, it uses Policy Based Routing (PBR) to load-balance traffic across the Firepower nodes, which are inline and transparent to the traffic flows. However, L3 routing is still required between the Outside and Inside ITD service nodes so that global routing tables can be exchanged through the dual-VDC Sandwich. In this design, OSPF is the designated IGP and is enabled on VLAN 101 and 102. Additionally, L3 over VPC requires in the VPC domain configuration the **layer 3 peer-router** command. In the Nexus 7000 series of switches NXOS version 7.2 is required for this feature.

Figure 5. Flow Diagram

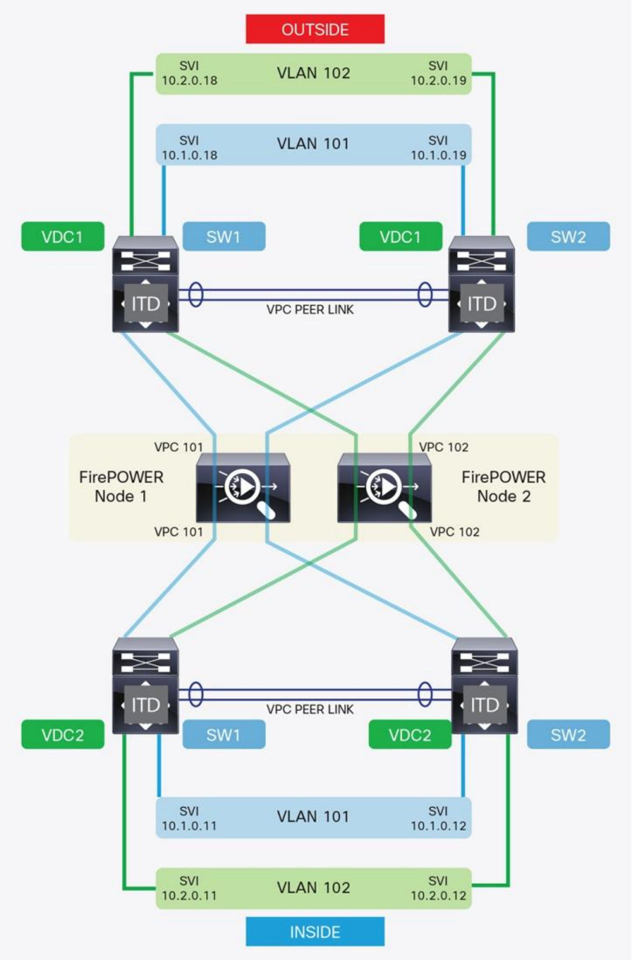

ITD: Load Balancing

The ITD service uses a combination of PBR with customized hashing to achieve its load-balancing and traffic steering capabilities. This section covers the two components--nodes (next-hop) and hashing (buckets)-- that the ITD service utilizes. This process is illustrated for SW1-VDC1 in Figure 6.

The first component is the node configuration which defines the next-hop behavior that the ITD service uses to route traffic. The previous section covered the creation of VLANs with their associated SVI. The peer ITD node IP value will be the SVI IP that was created in the ITD peer switch. Table 1, below, illustrates SW1-VDC1 and its peer switch configuration SW1-VDC2.

Table 1. ITD Configuration Snippet

SW1-VDC1	SW1-VDC2
itd device-group FP_INSIDE	itd device-group FP_INSIDE
probe icmp	probe icmp
node ip **10.1.0.11**	node ip **10.1.0.18**
node ip **10.2.0.11**	node ip **10.2.0.18**

For each node configured in the ITD device-group, a corresponding next-hop entry is created for PBR purposes. ITD will then create the underlying route-maps, track objects, and IP SLA automatically once the ITD service is started. The **show itd INSIDE statistics** command shows this next-hop configuration as the assigned and original nodes. These values should match unless one of the node IPs becomes unreachable. Table 2 illustrates an example of **show itd INSIDE statistics**.

Table 2. Show ITD INSIDE Statistics

```
Service           Device Group                VIP/mask              #Packets
-------------------------------------------------------------------------------------------
INSIDE            FP_INSIDE                                         0            0%

Traffic Bucket         Assigned to       Mode         Original Node       #Packets
---------------        ------------      ----------   ----------          ---------
INSIDE_itd_bucket_1     10.1.0.11        Redirect     10.1.0.11           14467813 (84.48%)

Traffic Bucket         Assigned to       Mode         Original Node       #Packets
------------------      -------------    -----        --------------      ------------
INSIDE_itd_bucket_2     10.2.0.11        Redirect     10.2.0.11           2656900  (15.52%)
```

The second component is the hashing buckets used to separate traffic. In this design, ITDs load-balances by using a hashing algorithm (buckets) based either on SA or DA IP address. The default value is to hash on the last octet of the IP address using a wild card mask (Table 3, 4). The ITD service will then create a minimum number of buckets based on the last octet so that every node will have a minimum of one bucket assigned.

Figure 6. ITD Algorithm

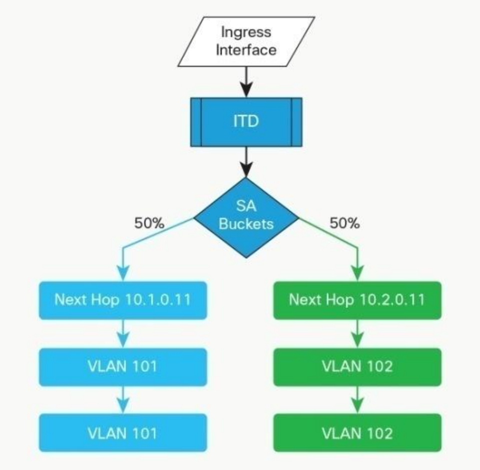

Table 3. Configuration SW1-VDC1

Node	Bucket	Firepower
SW1-VDC2	SA -> x.x.x.0-127	Node 1
SW2-VDC2	SA -> x.x.x.128-255	Node 2

Table 4. Configuration SW1-VDC2

Node	Bucket	Firepower
SW1-VDC2	DA -> x.x.x.0-127	Node 1
SW2-VDC2	DA -> x.x.x.128-255	Node 2

When determining optimal load-balancing, it is ideal that every node gets assigned an equal number of buckets. Reference Table 5 for the optimal combinations of number of Firepower nodes and the configured hashing buckets in the ITD service itself. The number of nodes is a recommendation not a requirement.

Table 5. Optimal Load-Balancing Configuration

Nodes	2	4	8	16
Buckets	[1, 2]	[1, 2, 4]	[1, 2, 4, 8]	[1, 2, 4, 8, 16]

Firepower 8300 Configuration

The Firepower nodes in this design require only basic configuration adjustments. The three components that must be configured are interface pairs, inline sets, and link propagation.

It is recommended to use the Firepower 40GE 2-port NetMods to achieve full capacity on a Firepower stack.

First, two interface inline pairs must be created on each node. In this design, inline pairs were created between [P1, P2] and [P3, P4]. See Figure 7.

Figure 7. Firepower Node

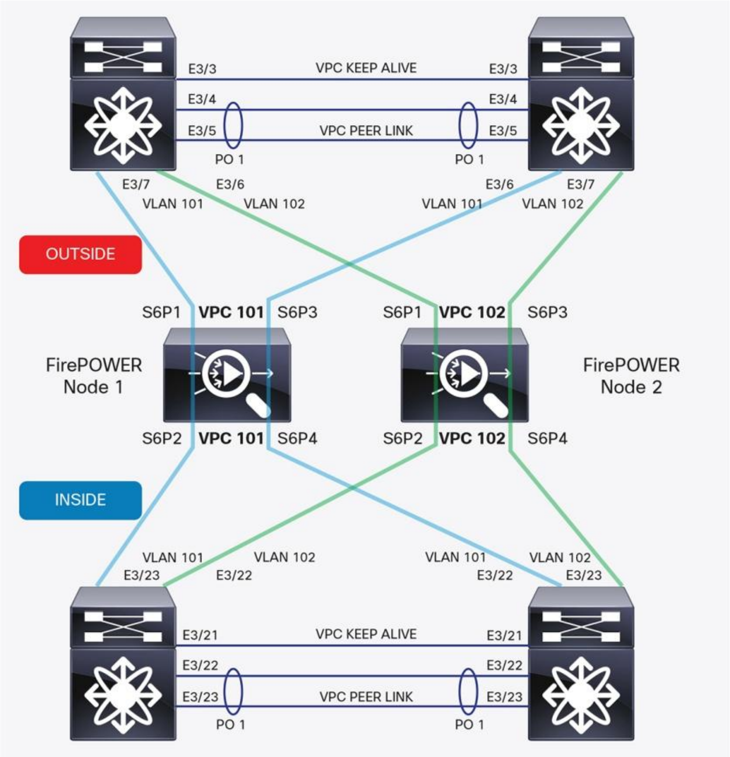

Second, an Inline set must be created. An inline set is a grouping of one or more inline interface pairs on a single Firepower node. Adding multiple inline interface pairs to the same inline interface set allows the system to identify the inbound and outbound traffic as part of the same traffic flow. In this design, both inline pairs consisting of [P1, P2] [P3, P4] were added to an inline set.

Third, link state propagation needs to be configured on each interface pair. Link state propagation is a feature of the Firepower appliance for inline sets configured in bypass mode so both pairs of an inline set track state. Link state propagation automatically shuts down the second interface in the inline interface pair when one of the interfaces in an inline set goes down.

When the previously shut down interface comes back up, the second interface automatically comes back up. In other words, if the link state of one interface changes, the appliance senses the change and updates the link state of the other interface to match it. Note that Firepower node may require up to 4 seconds to propagate link state changes however.

The configuration as outlined above is shown in Figure 8.

Figure 8. Inline Set

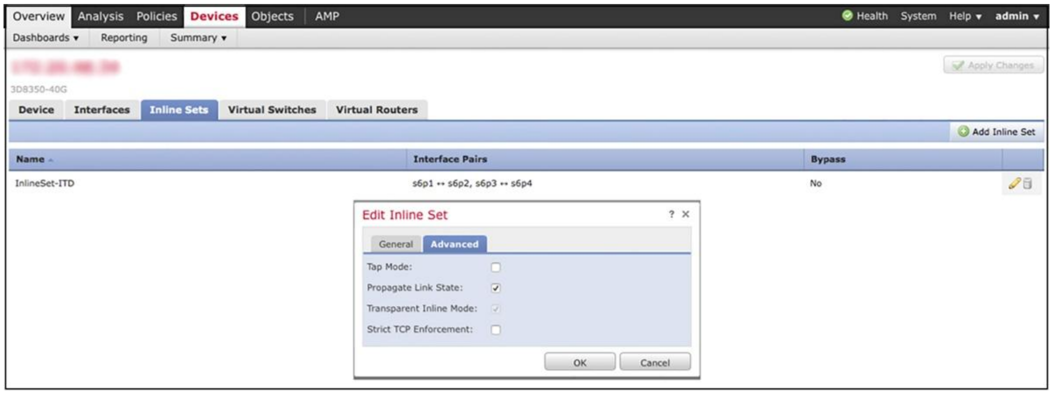

[redacted]
3D8350-40G

Name	Interface Pairs	Bypass
InlineSet-ITD	s6p1 ↔ s6p2, s6p3 ↔ s6p4	No

Tap Mode: □
Propagate Link State: ☑
Transparent Inline Mode: ☑
Strict TCP Enforcement: □

Failure Scenarios

Link-State Failure

If an individual link fails within the VPC in the dual-VDC sandwich, the link will be either shut by the N7700 switch that detects the failure or, secondarily, the Firepower node with link state detection will shut down both interfaces in the inline interface pair. The link state detections should occur within 4 seconds of detection.

With the interface failure, the N7700 ITD node will reroute traffic over the VPC peer link. The flow will then be forwarded to the same Firepower node as the hashing algorithm is used.

It is important to note that the ITD node configuration must match in order to maintain traffic flow symmetry during this failure scenario. The node list order impacts the hashing algorithm so the ITD device-groups must be identical. This ensures that hashing algorithms match on the ITD service nodes.

ITD Probe Failure

There are several situations that can cause the failure of the IP SLA for an individual ITD service node. The failure can be as simple as the SVI was shut or as major as a total switch failure.

In the event that an IP SLA fails, the ITD service will remove that specific node out of the next hop service rotation. As example in our design topology (Figure 5) VLAN 101 SVI (10.2.0.11) on SW1-VDC2 has failed. The ITD service will than rebalance traffic across the remaining ITD peer nodes (Table 6).

Table 6. Show ITD INSIDE Statistics

```
Service           Device Group               VIP/mask              #Packets
---------------------------------------------------------------------------------------------
INSIDE            FP_INSIDE                                        0             0%

Traffic Bucket          Assigned to      Mode       Original Node    #Packets
---------------         --------------   -----      --------------   ---------
INSIDE_itd_bucket_1     10.1.0.11        Redirect   10.1.0.11        14467813 (83.69%)

Traffic Bucket          Assigned to      Mode       Original Node    #Packets
---------------         --------------   -----      --------------   ---------
INSIDE_itd_bucket_2     10.1.0.11        Redirect   10.2.0.11        2818951  (16.31%)
```

In this design we used the default ITD service times. The ITD service, however, has configurable timers for the peer node monitoring.

Please consult the Nexus ITD Configuration Guide referenced at the end of this document for additional details on how to configure the various ITD service timers.

Firepower Node Failure

When a Firepower node is comprised of a stack of multiple devices (8360 -8390), it is important to understand the failure method.

When a secondary device is added to the primary 8350 a stacked configuration is established. In a stacked node, the devices act like a single, shared configuration. If the primary device fails, no traffic is passed to the secondary devices. Health alerts are generated indicating that the stacking heartbeat has failed on the secondary devices.

If the secondary device in a stack fails, inline sets with configurable bypass enabled go into bypass mode on the primary device. For all other configurations, the system continues to load balance traffic to the failed secondary device. In either case, a health alert is generated to indicate loss of link.

The Firepower appliance support a failsafe mode, where traffic is allowed to bypass detection and continue through the device. Managed devices monitor internal traffic buffers and bypass detection if those buffers are full.

The Firepower appliance devices also support a bypass mode to configure how the relays in the inline interfaces respond when an interface fails. The bypass mode allows traffic to continue to pass through the interfaces. The non-bypass mode blocks traffic.

Design and implementation must take into account whether the Firepower nodes will be configured in failsafe mode and or bypass mode. The configuration will dictate in a Firepower failure whether ITD service node will see a VPC link go down or failure of the IP SLA probes.

It is recommended that for scalability customers either deploy excess capacity in the Firepower node to accommodate a node failure or provision a N+1 node design. ITD service also offers the capability of a hot standby node.

Configuration: Nexus 7700 ITD

Configuration of ITD Node SW1-VDC1

```
feature itd

interface Ethernet3/1
  description to VDCL3-3
  ip address 10.0.128.25/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Ethernet3/2
  description to VDCL3-4
  ip address 10.0.128.29/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Vlan101
  description INSIDE_FP_VLAN
  ip address 10.1.0.18/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface Vlan102
  description INSIDE_FP_VLAN
  ip address 10.2.0.18/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface port-channel11
  description To_ITD-FP-1_PChannelInside
  switchport
  switchport access vlan 101
  vPC 101

interface port-channel12
  description To_ITD-FP-2_PChannelInside
  switchport
  switchport access vlan 102
  vPC 102

ip access-list INSIDE_itd_bucket_1
  10 permit ip 1.1.1.0 255.255.255.127 any
ip access-list INSIDE_itd_bucket_2
  10 permit ip 1.1.1.128 255.255.255.127 any
```

```
track 1 interface Ethernet3/1 line-protocol
track 2 interface Ethernet3/2 line-protocol
track 3 ip sla 10003 reachability
 delay up 30 down 30
track 4 ip sla 10004 reachability
 delay up 30 down 30

ip sla 10003
  icmp-echo 10.1.0.11
    frequency 10
ip sla schedule 10003 life forever start-time now

ip sla 10004
  icmp-echo 10.2.0.11
    frequency 10
ip sla schedule 10004 life forever start-time now

route-map INSIDE_itd_pool pbr-statistics
route-map INSIDE_itd_pool permit 10
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_1
  set ip next-hop verify-availability 10.1.0.11 track 3
route-map INSIDE_itd_pool permit 11
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_2
  set ip next-hop verify-availability 10.2.0.11 track 4

itd device-group FP_INSIDE
  probe icmp
  node ip 10.1.0.11
  node ip 10.2.0.11

itd INSIDE
  device-group FP_INSIDE
  ingress interface Eth3/1
  ingress interface Eth3/2
  failaction node reassign
  load-balance method src ip
  no shut
```

Configuration of ITD Node SW1-VDC2

```
feature itd

interface Ethernet3/17
  description to VDCL3-1
  ip address 10.0.0.17/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Ethernet3/18
  description to VDCL3-2
  ip address 10.0.0.21/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Vlan101
  description OUTSIDE_FP_VLAN
  ip address 10.1.0.11/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface Vlan102
  description OUTSIDE_FP_VLAN
  ip address 10.2.0.11/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface port-channel11
  description To_ITD-FP-1_PChannelInside
  switchport
  switchport access vlan 101
  vPC 101

interface port-channel12
  description To_ITD-FP-2_PChannelInside
  switchport
  switchport access vlan 102
  vPC 102

ip access-list INSIDE_itd_bucket_1
  10 permit ip any 1.1.1.0 255.255.255.127
ip access-list INSIDE_itd_bucket_2
  10 permit ip any 1.1.1.128 255.255.255.127
```

```
track 1 interface Ethernet3/17 line-protocol
track 2 interface Ethernet3/18 line-protocol
track 3 ip sla 10003 reachability
 delay up 30 down 30
track 4 ip sla 10004 reachability
 delay up 30 down 30

ip sla 10003
  icmp-echo 10.1.0.18
    frequency 10
ip sla schedule 10003 life forever start-time now

ip sla 10004
  icmp-echo 10.2.0.18
```

```
    frequency 10
ip sla schedule 10004 life forever start-time now

route-map INSIDE_itd_pool pbr-statistics
route-map INSIDE_itd_pool permit 10
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_1
  set ip next-hop verify-availability 10.1.0.18 track 3
route-map INSIDE_itd_pool permit 11
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_2
  set ip next-hop verify-availability 10.2.0.18 track 4

itd device-group FP_INSIDE
  probe icmp
  node ip 10.1.0.18
  node ip 10.2.0.18

itd INSIDE
  device-group FP_INSIDE
  ingress interface Eth3/17
  ingress interface Eth3/18
  failaction node reassign
  load-balance method dst ip
  no shut
```

Configuration of ITD Node SW2-VDC1

```
feature itd

interface Ethernet3/1
  description to VDCL3-4
  ip address 10.0.0.5/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Ethernet3/2
  description to VDCL3-3
  ip address 10.0.0.1/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Vlan101
  description INSIDE_FP_VLAN
  ip address 10.1.0.19/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface Vlan102
  description INSIDE_FP_VLAN
  ip address 10.2.0.19/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0


interface port-channel11
  description To_ITD-FP-1_PChannelInside
  switchport
  switchport access vlan 101
  vPC 101

interface port-channel12
  description To_ITD-FP-2_PChannelInside
  switchport
  switchport access vlan 102
  vPC 102

ip access-list INSIDE_itd_bucket_1
  10 permit ip 1.1.1.0 255.255.255.127 any
ip access-list INSIDE_itd_bucket_2
  10 permit ip 1.1.1.128 255.255.255.127 any
```

```
track 1 interface Ethernet3/2 line-protocol
track 2 interface Ethernet3/1 line-protocol
track 3 ip sla 10003 reachability
 delay up 30 down 30
track 4 ip sla 10004 reachability
 delay up 30 down 30

ip sla 10003
  icmp-echo 10.1.0.12
    frequency 10
ip sla schedule 10003 life forever start-time now
ip sla 10004
  icmp-echo 10.2.0.12
    frequency 10
ip sla schedule 10004 life forever start-time now

route-map INSIDE_itd_pool pbr-statistics
route-map INSIDE_itd_pool permit 10
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_1
  set ip next-hop verify-availability 10.1.0.12 track 3
route-map INSIDE_itd_pool permit 11
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_2
  set ip next-hop verify-availability 10.2.0.12 track 4
```

```
itd device-group FP_INSIDE
  probe icmp
  node ip 10.1.0.12
  node ip 10.2.0.12

itd INSIDE
  device-group FP_INSIDE
  ingress interface Eth3/2
  ingress interface Eth3/1
  failaction node reassign
  load-balance method src ip
  no shut
```

Configuration of ITD Node SW2-VDC2

```
feature itd

interface Ethernet3/17
  description to VDCL3-2
  ip address 10.0.0.29/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Ethernet3/18
  description to VDCL3-1
  ip address 10.0.0.25/30
  ip router ospf 1 area 0.0.0.0
  ip policy route-map INSIDE_itd_pool

interface Vlan101
  description OUTSIDE_FP_VLAN
  ip address 10.1.0.12/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface Vlan102
  description OUTSIDE_FP_VLAN
  ip address 10.2.0.12/24
  ip ospf cost 1
  ip router ospf 1 area 0.0.0.0

interface port-channel11
  description To_ITD-FP-1_PChannelInside
  switchport
  switchport access vlan 101
  vPC 101

interface port-channel12
```

```
  description To_ITD-FP-2_PChannelInside
  switchport
  switchport access vlan 102
  vPC 102

ip access-list INSIDE_itd_bucket_1
  10 permit ip any 1.1.1.0 255.255.255.127
ip access-list INSIDE_itd_bucket_2
  10 permit ip any 1.1.1.128 255.255.255.127
```

```
track 1 interface Ethernet3/17 line-protocol
track 2 interface Ethernet3/18 line-protocol
track 3 ip sla 10003 reachability
 delay up 30 down 30
track 4 ip sla 10004 reachability
 delay up 30 down 30

ip sla 10003
  icmp-echo 10.1.0.19
    frequency 10
ip sla schedule 10003 life forever start-time now
ip sla 10004
  icmp-echo 10.2.0.19
    frequency 10
ip sla schedule 10004 life forever start-time now

route-map INSIDE_itd_pool permit 10
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_1
  set ip next-hop verify-availability 10.1.0.19 track 3
route-map INSIDE_itd_pool permit 11
  description auto generated route-map for ITD service INSIDE
  match ip address INSIDE_itd_bucket_2
  set ip next-hop verify-availability 10.2.0.19 track 4

itd device-group FP_INSIDE
  probe icmp
  node ip 10.1.0.19
  node ip 10.2.0.19

itd INSIDE
  device-group FP_INSIDE
  ingress interface Eth3/17
  ingress interface Eth3/18
  failaction node reassign
  load-balance method dst ip
  no shut
```

For More Information

- Firepower 8000 Installation Guide
- Firepower 8000 Stack Configuration Guide
- Nexus 7000 ITD Configuration Guide

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C07-738152-00 11/16