Cisco Intelligent Traffic Director Deployment Guide with Cisco ASA

Contents

Introduction

Cisco® Intelligent Traffic Director (ITD) is a multiterabit Layer 4 load-balancing, traffic steering, and clustering solution on the Cisco Nexus® 5000, 6000, 7000, and 9000 Series Switches. ITD is performed in hardware on application-specific integrated circuits (ASICs) and provides scalable load distribution of traffic to a group of servers or service appliances.

In scalable firewall solutions, Intelligent traffic director can distribute traffic to Cisco Adaptive Security Appliances (ASAs) configured as clusters or to standalone nonclustered ASA devices. A typical standalone deployment uses a failover pair. Intelligent traffic director can load-balance to two or multiple failover pairs.

Alternatively, ITD can load-balance to nonfailover (active-active) firewalls so that each firewall can actively process a percentage of the traffic. ITD health probes and ASA redundant interfaces can be used to help ensure high availability and failover.

This document discusses scenarios that use Intelligent Traffic Director in conjunction with multiple ASA and firewall devices that operate in active-active nonfailover modes.

Deployment Methods

Traffic can be distributed to firewalls with Intelligent Traffic Director using the following topologies:

- Intelligent Traffic Director in one-arm mode with “firewall on a stick”: In a “Firewall on a stick” mode the ASA devices perform traffic filtering and Inter-VLAN Routing by splitting the single interface into virtual subinterfaces. This design uses a single virtual device context (VDC) with a single IEEE 802.1q interface (or IEEE 802.1q port channel) connecting to the ASA devices.
- Intelligent Traffic Director in two-arm mode with single VDC: This design uses a single VDC with two separate (access or trunk) interfaces connecting to the ASA devices. The ASA devices filter traffic traversing the two interfaces. Traffic can segregated on the switch by Virtual Routing and Forwarding (VRF) instances to help ensure that traffic is inspected by the firewalls.
- Intelligent Traffic Director in dual-VDC sandwich mode without virtual port-channels (single Cisco Nexus 7000 Series Switch): This design uses two VDCs, each with an interface connecting to the ASA devices. The ASA devices filter traffic traversing the two VDCs. This design can also be used with separate Cisco Nexus 7000 Series Switches instead of VDCs if desired.
- Intelligent Traffic Director in dual-VDC sandwich mode with virtual port-channels (two Cisco Nexus 7000 Series Switches): This design uses two VDCs per switch, each with an interface connecting to the ASA devices. The ASA devices filter traffic traversing the two VDCs. This design can also be used with separate Cisco Nexus 7000 Series Switches instead of VDCs if desired. Two Cisco Nexus 7000 Series Switches are deployed in vPC mode.
- Intelligent Traffic Director with Layer 3 ASA clustering: This design can use any of the preceding topologies, but for the ASA, a cluster is configured with the cluster-control links (CCLs), with the ASA in Layer 3 mode.

All the deployments discussed in this document use ASA devices operating in routed Layer 3 mode because Intelligent traffic director natively requires a Layer 3 IP address to which it redirects traffic. These designs are discussed here using Cisco Nexus 7000 Series Switches as an example—in particular, Cisco Nexus 7700 platform switches—running Cisco NX-OS Software Release 7.2(0)D1(1). However these will work equivalently with Nexus 9000 and Nexus 5000/6000 Series switches as well, albeit without VDC’s since they are supported only on the Nexus 7000. The Nexus 5500/5600/6000 series switches currently do not support Standby/Backup node functionality.

The configurations presented in the scenarios here are only relevant snippets that are unique to each method. Full configurations and configurations for other features of the device (for example, vPC) are beyond the scope of this document.

In scenarios in which both non-vPC and vPC connections are possible, the example shows a vPC; however, single connections can be configured in a similar fashion.

In the figures in this document, blue links and labels typically refers to ASA “inside” connections and configurations, and red links and labels refers to ASA “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. Because an ASA device has different inside and outside interface IP addresses, you also need to configure two different traffic director device groups with the corresponding inside and outside IP addresses.

Additional Cisco ASA VLANs

The traffic director's forward and return services are attached to the inside and outside VLAN switch virtual interfaces (SVIs) on the Cisco Nexus switch. Because a security application such as a firewall needs to examine all traffic, no traffic filtering is configured in the traffic director. As a result, any traffic that reaches the SVI will be redirected to the corresponding ASA interfaces.

If the ASA interfaces are configured on the same subnets -facing inside and outside VLANs(10,20) on the switch, the inspected traffic returning to the switch from the firewall will be redirected back to the ASA again due to the presence of an ITD service on that VLAN. Hence, you need to use a pair of separate VLANs to prevent traffic looping between the firewalls and the Cisco Nexus switches.

The example in Figure 1 shows VLANs 10 and 20 as the inside and outside interfaces to the source and destination on the network, and VLANs 100 and 200 as the interfaces to the ASA devices to help ensure loop-free traffic. Note that physical Layer 3 interfaces can be used as traffic director ingress interfaces if required.

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

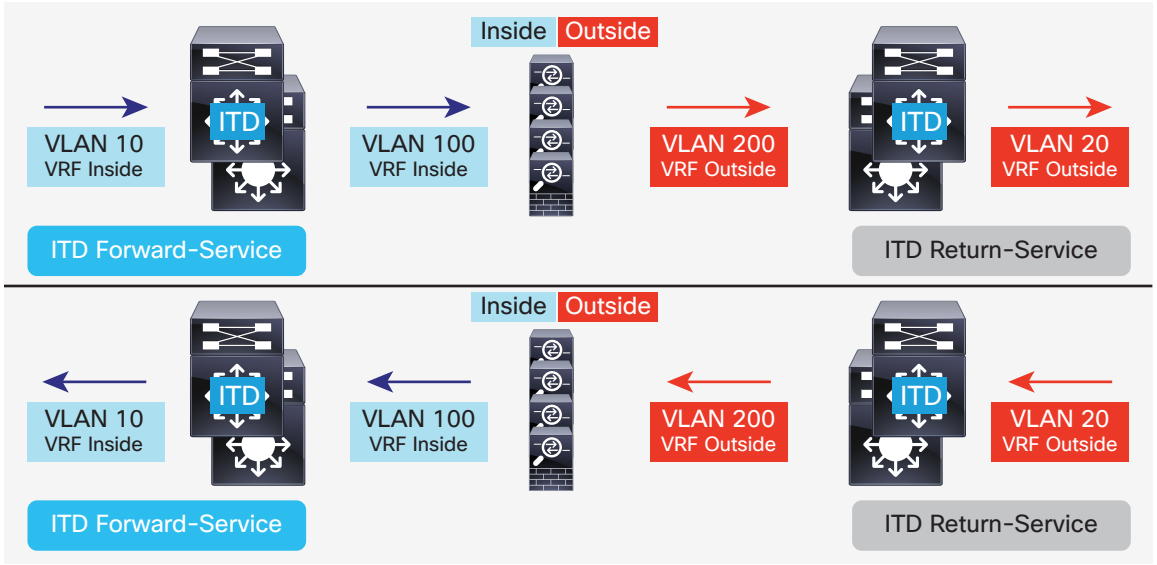

Flow Symmetry

Firewalls typically inspect traffic flows in both forward and return directions. Due to the stateful nature of the inspection, flow symmetry usually must be maintained during normal operation of firewalls that are not clustered. Even with clustered firewalls, asymmetry of traffic flows results in increased redirection over the CCLs. Asymmetry also adds unnecessary overhead to the firewalls and adversely affects performance.

Flow symmetry can be achieved using the inherent IP address persistence and the deterministic nature of the ITD algorithms. Typical ITD configuration with firewalls uses one ITD service for the forward flow and one 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.

Figure 2. Flow Symmetry Considerations with Cisco Intelligent Traffic Director and Cisco ASA Deployment

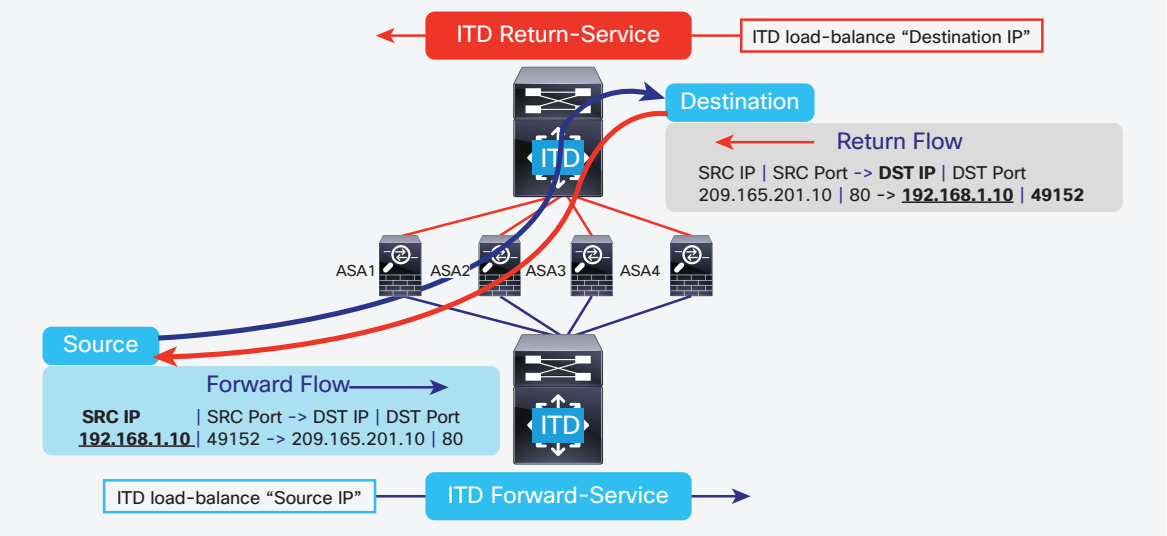

Figure 3. Cisco ASA Failure Scenario (Without peer-VDC node-state sync)

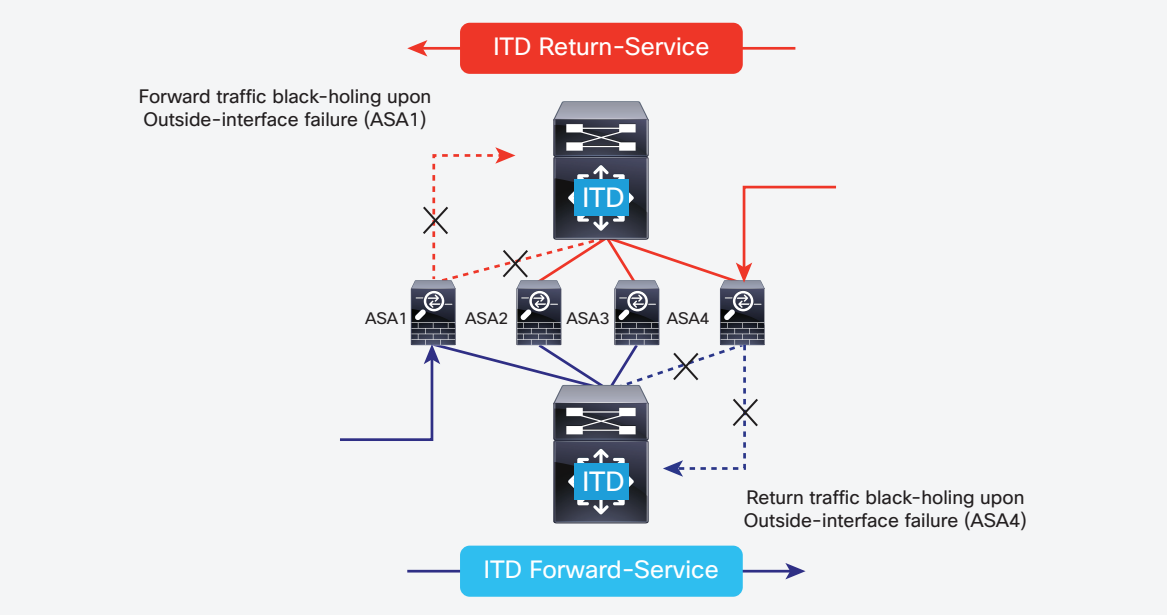

Link-Failure Scenario

When one of the interfaces of the ASA (inside or outside) fails, the ingress traffic on the other interface may be silently discarded, or "black-holed," because the egress interface is down.

ITD's **node-state synchronization** feature can resolve this problem by synchronizing the node states across the ITD services and removing the remote side of the ASA from the corresponding ITD service. Failover options in ITD will then move the traffic buckets to other firewalls as configured. (Figure 3).

The **peer-VDC node-state sync** feature currently is supported only in the dual VDC non-vPC (single switch) topology using the Cisco Nexus 7000 Series and with the Sandwich mode topology in the Nexus 9000 Series switches. The firewall-on-a-stick implementation (single link or vPC) does not experience this problem because the inside and outside interfaces on the ASA belong to the same physical (or virtual) interface. Clustering would also prevent this problem because it helps ensure that the ASA is fully taken out of service in the event of a single-side failure.

Cisco Intelligent Traffic Director and Cisco ASA Deployment

Single VDC: Cisco Intelligent Traffic Director with Firewall on a Stick (One-Arm Mode)

In firewall-on-a-stick deployments, vPC (or single-port) trunks typically are used to connect the ASA devices to the switches. In this configuration, the inside and outside interfaces are IEEE 802.1q subinterfaces in VLANs 100 and 200, and the switches have two VLANs and SVIs each in the inside and outside contexts without a physical port separation between them (Figure 4).

Figure 4. Firewall on a Stick (with vPC) Deployment

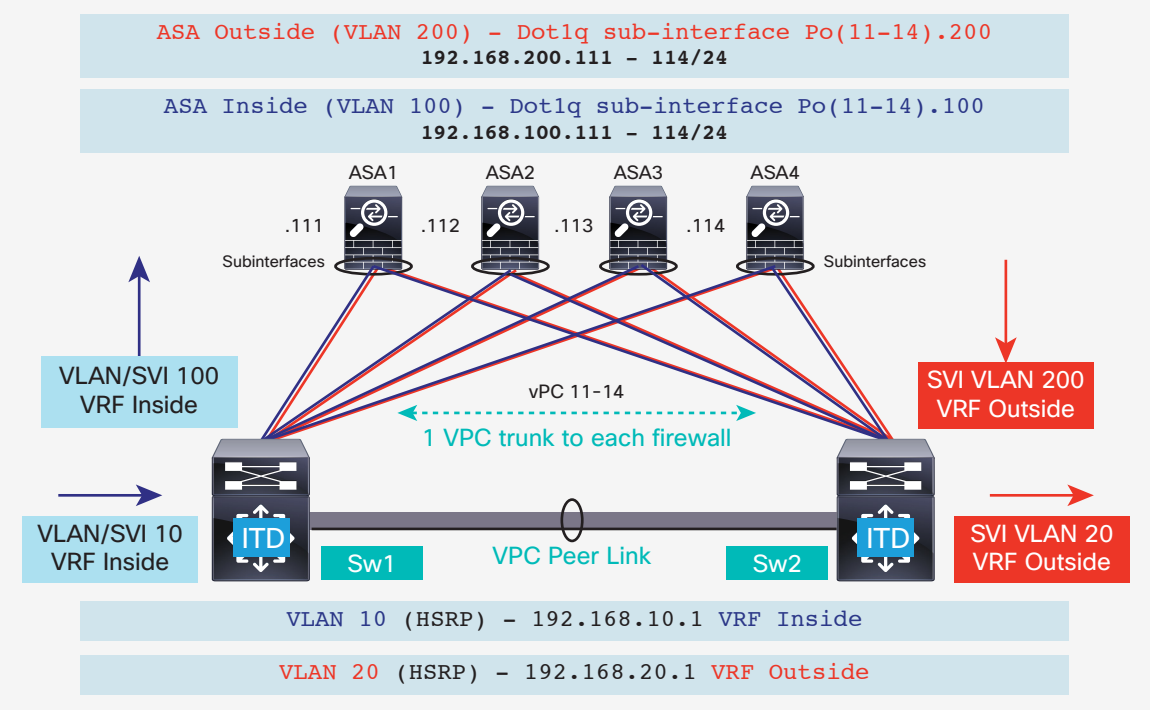

Configuration Snippets: Cisco Nexus 7000 Series

Sample configuration snippets for the Cisco Nexus 7000 Series are shown here. The example shows relevant configurations from one switch (sw1) and ASA (asa1). The configurations need to be extended similarly to other devices. Other necessary features are assumed to be already configured.

```
interface Vlan10
  description Inside_Vlan_to_Network
  vrf member INSIDE
  ip address 192.168.10.10/24
  hsrp 10
    ip 192.168.10.1

interface Vlan 20
  description Outside_Vlan_to_Network
  vrf member OUTSIDE
  ip address 192.168.20.10/24
  hsrp 20
    ip 192.168.20.1

interface Vlan100
  description Inside_Vlan_to_ASA
  vrf member INSIDE
  ip address 192.168.100.10/24
  hsrp 100
    ip 192.168.100.1

interface Vlan200
  description Outside_Vlan_to_ASA
  vrf member OUTSIDE
  ip address 192.168.200.10/24
  hsrp 200
    ip 192.168.200.1
```

```
interface Port-Channel11
  description VPC_TO_ASA1
  switchport mode trunk
  switchport trunk allowed vlan 100,200
  vpc 11
  no shutdown

interface Ethernet4/25
  description Link_To_ITD-ASA-1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 100,200
  channel-group 11 mode active
  no shutdown

interface port-channel41
  description Downstream_vPC_to_network
  switchport mode trunk
  switchport trunk allowed vlan 10,20
  vpc 41
  no shutdown

interface Ethernet 5/1-4
  description Downstream_vPC_member
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20
  channel-group 41
  no shutdown
```

```
itd device-group FW_INSIDE
  #Config Firewall Inside interfaces as nodes
 node ip 192.168.100.111
 node ip 192.168.100.112
 node ip 192.168.100.113
 node ip 192.168.100.114
probe icmp frequency 5 timeout 5 retry-count 1

itd device-group FW_OUTSIDE
  #Config Firewall Outside interfaces as nodes
 node ip 192.168.200.111
 node ip 192.168.200.112
 node ip 192.168.200.113
 node ip 192.168.200.114
probe icmp frequency 5 timeout 5 retry-count 1
```

```
itd INSIDE
 vrf INSIDE
 device-group FW_INSIDE
 ingress interface Vlan10
 failaction node reassign
 load-balance method src ip buckets 16
no shut


itd OUTSIDE
 vrf OUTSIDE
device-group FW_OUTSIDE
 ingress interface Vlan20
 failaction node reassign
 load-balance method dst ip buckets 16
no shut
```

Configuration Snippets: Cisco ASA

The ASA side of the configuration is show here for ASA1. A similar configuration must be extended to all other ASA devices.

```
interface Port-channel11
 nameif aggregate
 security-level 100
 no ip address
!
interface Port-channel11.100
 description INSIDE
 vlan 100
 nameif inside
 security-level 100
 ip address 192.168.100.111 255.255.255.0
!
interface Port-channel11.200
 description OUTSIDE
 vlan 200
 nameif outside
 security-level 100
 ip address 192.168.200.111 255.255.255.0
!
same-security-traffic permit inter-
interface
```

```
interface TenGigabitEthernet0/6
 description CONNECTED_TO_SWITCH-A-VPC
 channel-group 11 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/7
 description CONNECTED_TO_SWITCH-B-VPC
 channel-group 11 mode active
 no nameif
 no security-level
```

Note the following points in the preceding configurations and topology:

- Verify VLANs and SVIs and their mappings to appropriate VRF instances.
- Verify both ITD device-group IP addresses :.ASA inside and outside
- Verify ITD load-balancing configuration to achieve flow symmetry.
- In a vPC scenario, as long as one member of the vPC is active, there will be no change to ITD. Redirected traffic on the switch with the failed vPC segment will now traverse the peer switch through the peer link as in a typical vPC case.
- This deployment method does not black-hole traffic upon failure of an ASA inside or outside interface, because these interfaces are tied to the same physical or virtual interfaces on the ASA (IEEE 802.1q subinterfaces).
- To support routing protocol neighborship over vPC (NX-OS 7.2(0)D1(1) or later), you need to configure the command **layer3 peer-router** in the vPC domain.
- The VRF instances are needed because Layer 3 interfaces are used to connect to both inside and outside firewall interfaces. VRF instances are used to prevent traffic (inter-VLAN) from being routed around the firewall in certain cases.
- Traffic is redirected to ASA devices through policy-based routing (PBR). Thus, routes are not needed.

Single VDC: Cisco Intelligent Traffic Director with Single VDC (Two-Arm Mode)

In a two-arm single VDC ITD deployment, vPC (or single physical) access ports are typically used to connect the ASA devices to the switches. In this configuration, the inside and outside interfaces are separate port channels, and the switches have two VLANs and SVIs each going to the ASA in the inside and outside contexts with physical port separation between them (Figure 5).

Figure 5. Two-Arm Mode Deployment with vPC

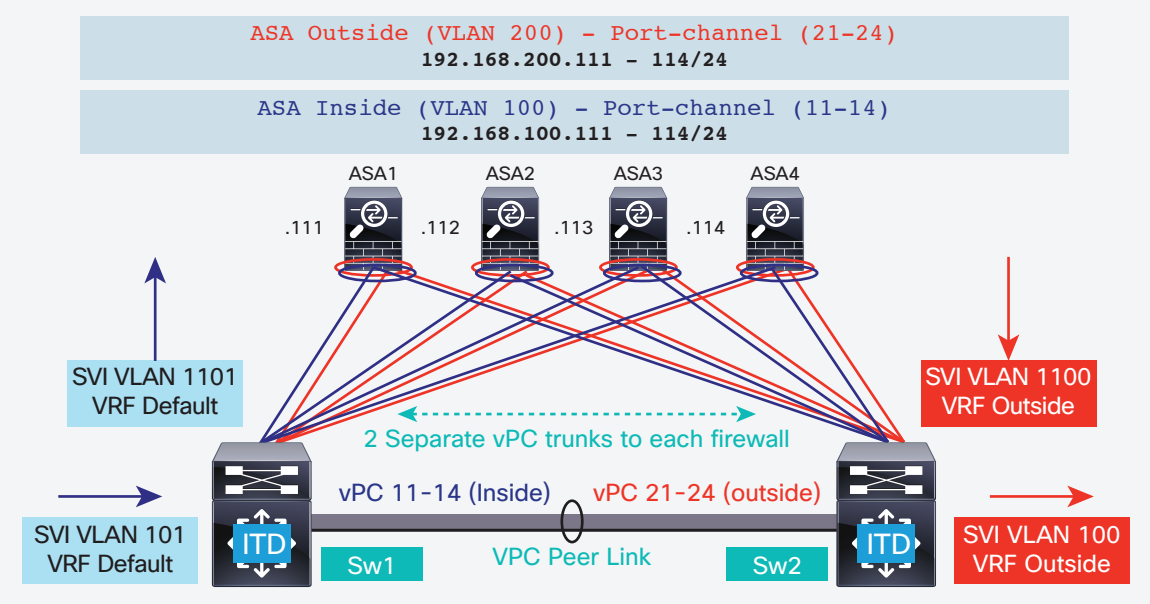

Configuration Snippets: Cisco Nexus 7000 Series

Configuration in this case is similar to that for the preceding case, with the difference being that in this case two vPCs (or separate physical links in a non-vPC case) go to each firewall. These port channels can be configured as either access ports or trunks, with only the appropriate VLAN allowed, to maintain physical port separation of the inside and outside VLANs going to the firewall.

A snippet of the configuration for the two port channels to ASA1 is shown here switch1.

```
interface Port-channel 11
  description To_ITD-ASA-1_Inside
  switchport mode access
  switchport access vlan 100
  vpc 11

interface Ethernet4/1
  description To_ITD-ASA-1_Inside
  switchport mode access
  switchport access vlan 100
  channel-group 11 mode active
```

```
interface Port-channel 21
  description To_ITD-ASA-1_Outside
  switchport mode access
  switchport access vlan 200
  vpc 21

interface Ethernet4/25
  description To_ITD-ASA-1_Outside
  switchport mode access
  switchport access vlan 200
  channel group 21 mode active
```

Configuration Snippets: Cisco ASA

The configuration of the ASA differs from that in the previous scenario in that the inside and outside interfaces are now configured on two separate port-channel interfaces.

```
interface Port-channel11
 description INSIDE
 vlan 100
 nameif inside
 security-level 100
 ip address 192.168.100.111 255.255.255.0
!
interface Port-channel21
 description OUTSIDE
 vlan 100
 nameif outside
 security-level 100
 ip address 192.168.200.111 255.255.255.0
!
same-security-traffic permit inter-
interface
```

```
interface TenGigabitEthernet0/6
 description CONNECTED_TO_SWITCH-A-VPC
 channel-group 11 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/7
 description CONNECTED_TO_SWITCH-B-VPC
 channel-group 11 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/8
 description CONNECTED_TO_SWITCH-A-VPC
 channel-group 21 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/9
 description CONNECTED_TO_SWITCH-B-VPC
 channel-group 21 mode active
 no nameif
 no security-level
```

Note the following points for the preceding configurations and topology:

- Verify VLANs and SVIs and their mappings to appropriate VRF instances.
- Verify both ITD device-group IP addresses:.ASA inside and outside.
- Verify the traffic director load-balancing configuration to achieve flow symmetry.
- In a vPC scenario, as long as one member of the vPC is active, there will be no change to ITD. Redirected traffic on the switch with the failed vPC segment will now traverse the peer switch through the peer link as in a typical vPC case.
- In this topology and deployment method, traffic can be black-holed if one of the port channels on the ASA (or a single physical link in a non-vPC case) fails.
- To support routing protocol neighborship over vPC (NX-OS 7.2(0)D1(1) or later), you need to configure the command **layer3 peer-router** in the vPC domain.
- The VRF instances are needed because Layer 3 interfaces are used to connect to both inside and outside firewall interfaces. VRF instances are used to prevent traffic (inter-VLAN) from being routed around the firewall in certain cases.
- Traffic is directed to the ASA device through PBR. Thus, routes are not needed.

Cisco Intelligent Traffic Director with Dual-VDC Sandwich Mode Without vPC

In a dual-VDC sandwich mode deployment, access ports (or trunks with only corresponding VLANs allowed) typically is used to connect the ASA devices to the switches. On the ASA, the inside and outside interfaces are separate interfaces (physical or virtual), achieving physical port seperation. The main difference in this deployment method is the seperation of inside and outside traffic into different VDCs on the Cisco Nexus 7000 Series Switch. As a result, the inside and outside VLANs and SVIs are configured in the appropriate VDCs (Figure 6).

Figure 6. Dual-VDC Single-Switch Sandwich Mode

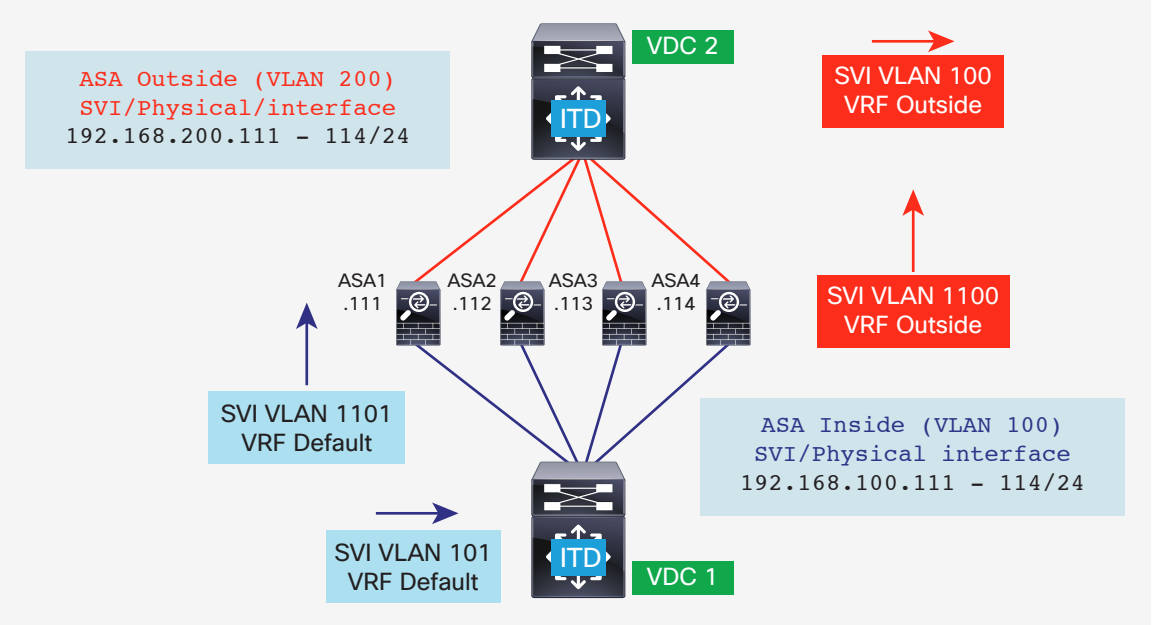

Configuration Snippets: Cisco Nexus 7000 Series

This deployment method (on the supported platforms/topologies) supports the synchronization of node states between the inside and outside ITD services. When configured this synchronization occurs internally on the Cisco Nexus 7000/9000 Series switches and helps ensure if an ASA's inside or outside interface fails(identified through traffic director probes), the ASA is removed from traffic distribution on both ITD services.

The following configuration snippets show the ITD service configured with this feature as well as the configurations on the two VDCs.

VDC1

```
interface Vlan10
  description INSIDE_VLAN
  ip address 192.168.10.10/24

interface Vlan100
  description FW_INSIDE_VLAN
  ip address 192.168.100.10/24


interface Ethernet4/1
  description To_ASA-1_INSIDE
  switchport mode access
  switchport access vlan 100
```

```
itd device-group FW_INSIDE
  #Config Firewall Inside interfaces as nodes
 node ip 192.168.100.111
 node ip 192.168.100.112
 node ip 192.168.100.113
 node ip 192.168.100.114
probe icmp frequency 5 timeout 5 retry-count 1


itd INSIDE
 vrf INSIDE
 device-group FW_INSIDE
 ingress interface Vlan10
 failaction node reassign
 load-balance method src ip buckets 16
 peer vdc ASA-OUTSIDE service OUTSIDE
   #Identifies the Peer-VDC, and the peer-ITD
   Service to sync node states with.
 no shut
```

VDC2

```
interface Vlan20
  description OUTSIDE_VLAN
  ip address 192.168.20.10/24

interface Vlan200
  description FW_OUTSIDE_VLAN
  ip address 192.168.200.10/24


interface Ethernet4/25
  description To_ASA-1_OUTSIDE
  switchport mode access
  switchport access vlan 200
```

```
itd device-group FW_OUTSIDE
  #Config Firewall Outside interfaces as nodes
 node ip 192.168.200.111
 node ip 192.168.200.112
 node ip 192.168.200.113
 node ip 192.168.200.114
probe icmp frequency 5 timeout 5 retry-count 1


itd OUTSIDE
 vrf OUTSIDE
device-group FW_OUTSIDE
 ingress interface Vlan20
 failaction node reassign
 load-balance method dst ip buckets 16
 peer vdc ASA-INSIDE service INSIDE
   #Identifies the Peer-VDC, and the peer-ITD
   Service to sync node states with.
 no shut
```

Configuration Snippets: Cisco ASA

The configuration on the ASA in this scenario is similar to that in the previous scenarios. The configuration snippet here shows two separate physical interfaces configured as the inside and outside interfaces.

```
interface TenGigabitEthernet0/6
 description INSIDE
 nameif inside
 security-level 100
 ip address 192.168.100.111 255.255.255.0
!
interface TenGigabitEthernet0/8
description OUTSIDE
 nameif outside
 security-level 100
 ip address 192.168.200.111 255.255.255.0
!
same-security-traffic permit inter-interface
```

Note the following points:

- To support routing protocol neighborship over vPC (NX-OS 7.2(0)D1(1) or later), you need to configure the command **layer3 peer-router** in the vPC domain.
- The VRF instances are needed because Layer 3 interfaces are used to connect to both inside and outside firewall interfaces. VRF interfaces are used to prevent traffic (inter-VLAN) from being routed around the firewall in certain cases.
- Traffic is directed to the ASA devices through PBR. Thus, routes are not needed.

Cisco Intelligent Traffic Director with Dual-VDC Sandwich Mode with vPC

Sandwich mode with vPC is similar to the sandwich mode without vPC; however, the inside and outside ASA interfaces are each assigned to separate port-channel bundles. This sandwich topology with vPC currently does not support the node-state synchronization feature across the vPC pair. As a consequence of vPC, a single link failure does not impede traffic flow, and ITD will continue to forward through the peer switch's link to the ASA, as in other scenarios with vPC (Figure 7).

Figure 7. Dual-VDC Two-Switch Sandwich Mode with vPC

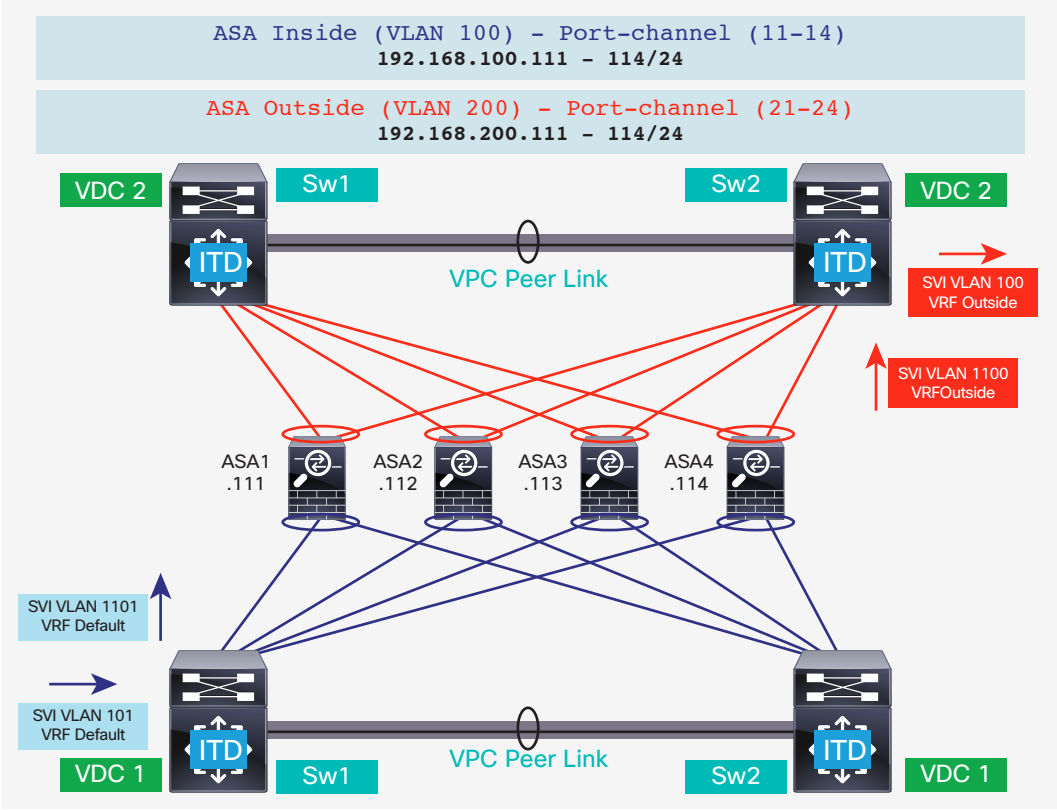

Configuration Snippets: Cisco Nexus 7000 Series

The main difference between this topology and the single-switch topology is that this topology uses vPCs instead of single links between the Cisco Nexus switch and the ASA. As in the previous scenario, the inside and outside interfaces on the switches are configured in different VDCs, as shown here.

#VDC1

```
interface Vlan10
  description INSIDE_VLAN
  ip address 192.168.10.10/24

interface Vlan100
  description FW_INSIDE_VLAN
  ip address 192.168.100.10/24
```

```
interface Port-channel11
  description To_ASA-1_INSIDE
  switchport mode access
  switchport access vlan 100
  vpc 11

interface Ethernet4/1
  description To_ASA-1_INSIDE
  switchport mode access
  switchport access vlan 100
  channel-group 11 mode active
```

#VDC2

```
interface Vlan20
  description OUTSIDE_VLAN
  ip address 192.168.20.10/24

interface Vlan200
  description FW_OUTSIDE_VLAN
  ip address 192.168.200.10/24
```

```
interface Port-channel21
  description To_ASA-1_OUTSIDE
  switchport mode access
  switchport access vlan 200
  vpc 21

interface Ethernet4/25
  description To_ASA-1_OUTSIDE
  switchport mode access
  switchport access vlan 200
  channel-group 21 mode active
```

The ITD configuration in the switch remains similar to that in the previous method except that the node-state synchronization command was removed because it is not supported.

Configuration Snippets: Cisco ASA

The ASA configuration is identical to that for the single VDC two-arm mode deployment with vPC. The following configuration snippet from the ASA is provided for reference.

```
interface Port-channel11
 description INSIDE
 vlan 100
 nameif inside
 security-level 100
 ip address 192.168.100.111 255.255.255.0
!
interface Port-channel21
 description OUTSIDE
 vlan 100
 nameif outside
 security-level 100
 ip address 192.168.200.111 255.255.255.0
!
same-security-traffic permit inter-
interface
```

```
interface TenGigabitEthernet0/6
 description CONNECTED_TO_SWITCH-A-VPC
 channel-group 11 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/7
 description CONNECTED_TO_SWITCH-B-VPC
 channel-group 11 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/8
 description CONNECTED_TO_SWITCH-A-VPC
 channel-group 21 mode active
 no nameif
 no security-level

interface TenGigabitEthernet0/9
 description CONNECTED_TO_SWITCH-B-VPC
 channel-group 21 mode active
 no nameif
 no security-level
```

Note the following points about the preceding configuration and topology:

- Verify the VLANs and SVIs and their mappings to appropriate VRF instances.
- Verify both ITD device-group IP addresses :.ASA inside and outside.
- Verify ITD load-balancing configuration to achieve flow symmetry.
- In a vPC scenario, as long as one member of the vPC is active, there will be no change to ITD. Redirected traffic on the switch with the failed vPC segment will now traverse the peer switch through the peer link as in a typical vPC case.
- In this topology and deployment method, traffic can be black-holed if one of the port channels on the ASA fails.
- To support routing protocol neighborship over vPC (NX-OS 7.2(0)D1(1) or later), you need to configure the command **layer3 peer-router** in the vPC domain.
- Traffic is directed to the ASA devices through PBR. Thus, routes are not needed.

Cisco Intelligent Traffic Director with Layer 3 Cisco ASA Clustering (with Any of the Preceding Topologies)

An ASA cluster consists of multiple ASA devices acting as a single unit. Grouping multiple ASAs together as a single logical device provides the convenience of a single device (management integration in a network) while achieving the increased throughput and redundancy of multiple devices.

ITD can load-balance to individual Layer 3 mode ASA clusters. ITD complements clustering by providing predictability so that you know which flows are handled by each firewall. ITD buckets determine which firewall handles a flow instead of relying on Open Shortest Path First (OSPF) equal-cost multipath (ECMP) or port-channel hashing algorithms.

With Layer 3 clusters, the flow owner can be predetermined based on the bucket allocation. Without the Intelligent Traffic Director and Layer 3 clustering, the initial choice of owner is typically unpredictable; however, with ITD, this choice can be predetermined.

ASA clustering also provides a backup flow owner. For every flow traversing any particular firewall in the cluster, another firewall stores the state of that flow and the ASA that owns the flow. If the active flow owner fails, ITD fail action reassignment, when configured, causes all flows (the buckets) from the failed ASA to shift to the next active ASA. If the new firewall to receive this traffic is not the backup owner for the flows it receives, it receives the flow state information from the backup owner and processes traffic transparently (Figure 8).

For more information, see http://www.cisco.com/c/en/us/td/docs/security/asa/asa90/configuration/guide/asa_90_cli_config/ha_cluster.html - pgfId-1556494.

A potential drawback of using ASA clustering with Intelligent Traffic Director is that backup flows and other cluster table operations consume memory and CPU resources that nonclustered firewalls do not. Therefore, you may get better firewall performance with nonclustered firewalls. However, the assurance of knowing that existing connections will not time out if an ASA cluster member fails may be of greater value to customers.

Figure 8. Cisco ASA Cluster in Dual-VDC Sandwich Mode with vPC

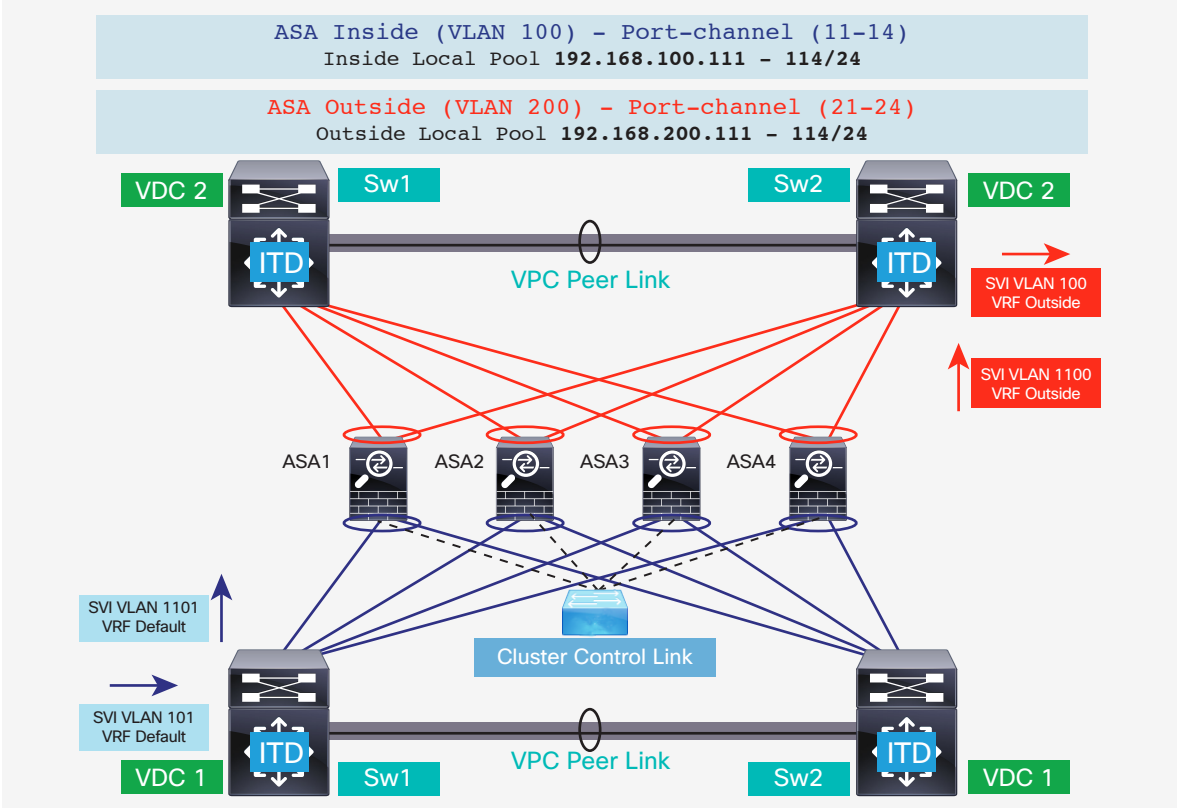

Configuration Snippets: Cisco Nexus 7000 Series

Introduction of clustering does not change ITD configuration. Hence, only the ASA configuration is discussed here. The ITD and Cisco Nexus switch configuration depends on the type of deployment used among those discussed previously. This example shows clustering deployed in dual-VDC sandwich mode with vPC.

Configuration Snippets: Cisco ASA

The ASA devices are configured in a Layer 3 cluster (similar to the PBR deployment scenario). Detailed information about the ASA clustering configuration can be found at http://www.cisco.com/c/en/us/td/docs/security/asa/asa90/configuration/guide/asa_90_cli_config/ha_cluster.html.

A sample configuration for the ASA is shown here for the topology in Figure 8.

```
cluster group ASA-CLUSTER-L3
 local-unit ASA1
 cluster-interface Port-channel31 ip 10.2.0.1 255.255.255.0
 priority 1
 health-check holdtime 1.5
 clacp system-mac auto system-priority 1
 enable


mac-address pool MAC-INSIDE aaaa.0101.0001 - aaaa.0101.0008
mac-address pool MAC-OUTSIDE aaaa.0100.0001 - aaaa.0100.0008
ip local pool IP-INSIDE 10.1.0.111-10.1.0.114
ip local pool IP-OUTSIDE 10.0.0.111-10.0.0.114


interface Port-channel11
 description INSIDE
 lacp max-bundle 8
 mac-address cluster-pool MAC-INSIDE
 nameif inside
 security-level 100
 ip address 10.1.0.11 255.255.255.0 cluster-pool IP-INSIDE
!
interface Port-channel21
 description OUTSIDE
 lacp max-bundle 8
 mac-address cluster-pool MAC-OUTSIDE
 nameif outside
 security-level 100
 ip address 10.0.0.11 255.255.255.0 cluster-pool IP-OUTSIDE

interface Port-channel31
 description Clustering Interface
 lacp max-bundle 8
!
```

```
interface TenGigabitEthernet0/6
 channel-group 11 mode active
 no nameif
 no security-level
 no ip address
!
interface TenGigabitEthernet0/7
 channel-group 11 mode active
 no nameif
 no security-level
 no ip address
!
interface TenGigabitEthernet0/8
 channel-group 21 mode active
 no nameif
 no security-level
 no ip address
!
interface TenGigabitEthernet0/9
channel-group 21 mode active
 no nameif
 no security-level
 no ip address
!
interface TenGigabitEthernet1/0
 channel-group 31 mode on
 no nameif
 no security-level
 no ip address
!
interface TenGigabitEthernet1/1
 channel-group 31 mode on
 no nameif
 no security-level
 no ip address
```

As seen in the configuration, port-channels 11 and 21 are used for the inside and outside interfaces as in previous cases. However, this configuration has an additional port channel, port-channel 31, for the clustering interface. Individual interfaces are normal routed interfaces, each with its own IP address taken from a pool of IP addresses. The main cluster IP address is a fixed address for the cluster that always belongs to the current master unit. Similarly, a MAC address pool is configured and addresses assigned to corresponding port channels (inside and outside).

Conclusion

Intelligent Traffic director offers traffic-distribution towards firewalls in different topologies that have been discussed in this document. Choice of a particular type of deployment is generally driven by the extent of requirement for segmentation between the Inside and Outside links and traffic flowing on them. By using ITD's inherently deterministic nature and chosing correct load-balancing options, flow symmetry can be achieved, which is a key requirement for Firewalls. Leveraging ITD's high scaling capacity, many Firewalls can be operated together (clustered or non-clustered) to cater to high network traffic requirements as well.

For More Information

- Nexus 7000 ITD configuration guide
- Nexus 9000 ITD configuration guide
- Nexus 5500 ITD configuration guide
- Cisco ASA configuration guide

C07-736769-00 04/16