Configuring ITD

This chapter describes how to configure Intelligent Traffic Director (ITD) on the Cisco NX-OS device.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table below.

Information About ITD

Intelligent Traffic Director (ITD) is an intelligent, scalable clustering and load-balancing engine that addresses the performance gap between a multi-terabit switch and gigabit servers and appliances. The ITD architecture

integrates Layer 2 and Layer 3 switching with Layer 4 to Layer 7 applications for scale and capacity expansion to serve high-bandwidth applications.

ITD provides adaptive load balancing to distribute traffic to an application cluster. With this feature on the Cisco Nexus 7000 Series switch, you can deploy servers and appliances from any vendor without a network or topology upgrade.

ITD Feature Overview

The ITD feature offers the following:

- Provides an ASIC-based ulti-terabit Layer 3 or Layer 4 solution to load balance traffic at line-rate.
- No service module or external Layer 3 or Layer 4 load-balancer is required.
- Every Cisco Nexus 7000 Series port can be used for load balancing.
- Can be used to redirect line-rate traffic to any device, such as web cache engines, Web Accelerator Engines (WAE), or video-caches, and so on.
- Can be used to load balance traffic to other software load balancers.
- Allows non-DSR Virtual IP (VIP) load-balancing deployments.
- Weighted load-balancing provides load-balancing to a large number of devices or servers.
- Allows ACL with simultaneous redirection and load balancing.
- Provides bi-directional flow-coherency; traffic from A to B and from B to A goes to same node.
- Provides the capability to create clusters of devices, such as firewalls, Intrusion Prevention System (IPS), Web Application Firewall (WAF), or Hadoop cluster IP-stickiness Resilient (like resilient ECMP).
- Supports the order of magnitude OPEX savings for a reduction in configuration and ease of deployment.
- Supports the order of magnitude CAPEX savings for wiring, power, rackspace and cost savings.
- The servers or appliances do not have to be directly connected to the Cisco Nexus 7000 Series switch.
- Monitors the health of servers and appliances.
- Provides N + M redundancy.
- Provides automatic failure handling of servers or appliances.
- Supports VRFs, vPCs, and VDCs.
- Supported on both the Cisco Nexus 7000 Series and Nexus 7700 Series switches.
- Supports both IPv4 and IPv6.

The following example use cases are supported by the Cisco ITD feature:

- Load-balance traffic to 256 servers of 10Gbps each.
- Load-balance to a cluster of Firewalls. ITD is much superior than policy-based routing (PBR).
- Scale up NG IPS and WAF by load-balancing to standalone devices.
- Scale the WAAS / WAE solution.
- Scale the VDS-TC (video-caching) solution.

- Replace ECMP/Port-channel to avoid re-hashing. ITD is resilient.

Benefits of ITD

ITD on the Cisco NX-OS switch enables the following:

- Horizontal scale—groups N servers for linear scaling and capacity expansion.
- Health monitoring of servers.
- Automatic failure detection of links and/or servers.
- Automatic traffic redistribution in case of a failure.
- Weight-based load balancing.
- Hot standby support of N+1 redundancy, with M standby.
- Node level standby support.
- Complete transparency to the end devices.
- No manual configuration or intervention required if a link or server fails.
- The use of heterogeneous types of servers and devices.
- Large number of servers supported.
- Simplified provisioning and ease of deployment.
- No certification, integration, or qualification needed between the devices and the Cisco NX-OS switch.
- The feature does not add any load to the supervisor CPU.
- ITD uses orders of magnitude less hardware TCAM resources than WCCP.
- Handles unlimited number of flows.

Deployment Modes

One-Arm Deployment Mode

You can connect servers to the Cisco NX-OS device in one-arm deployment mode. In this topology, the server is not in the direct path of client or server traffic, which enables you to plug in a server into the network with no changes to the existing topology or network.

Figure 1: One-Arm Deployment Mode

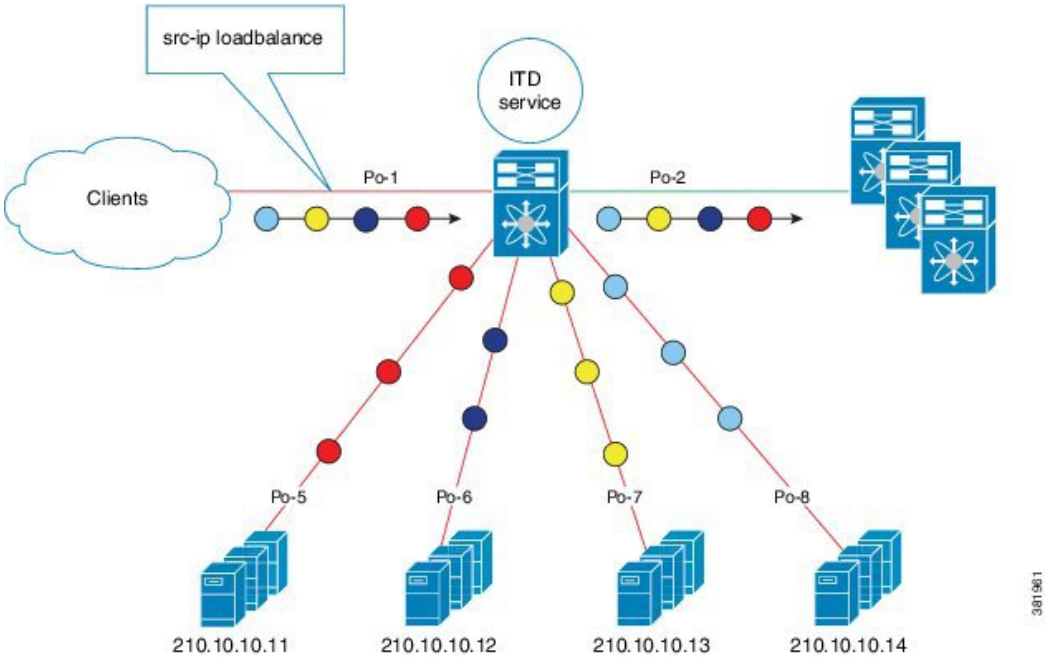

One-Arm Deployment Mode with VPC

The ITD feature supports an appliance cluster connected to a virtual port channel (vPC). The ITD service runs on each Cisco NX-OS switch and ITD programs each switch to provide flow coherent traffic passing through the nodes.

Figure 2: One-Arm Deployment Mode with VPC

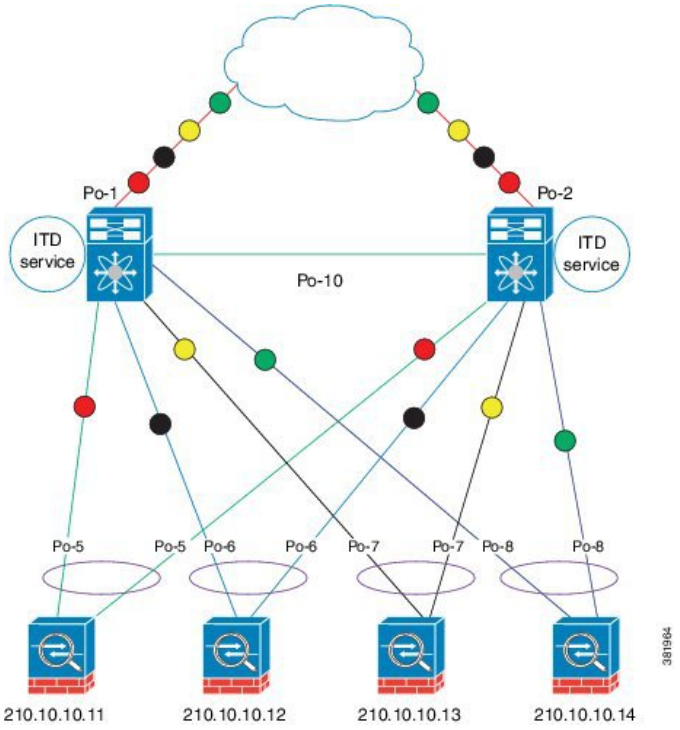

Sandwich Deployment Mode

The sandwich deployment mode uses two Cisco NX-OS 7000 Series switches to provide stateful handling of traffic.

The main requirement in this mode is that both forward and reverse traffic of a flow must go through the same appliance. Examples include firewalls and load balancer deployments, where traffic between client and server must flow through the same appliance.

The key features are:

- An ITD service for each network segment—one for outside network and another for inside network.
- A source-IP load balancing scheme where the ITD service operates on the interface that connects to the outside world in an ingress direction.

- A destination-IP load balancing scheme where the ITD service operates on the interface that connects to the servers in the ingress direction.

Figure 3: Sandwich Deployment Mode

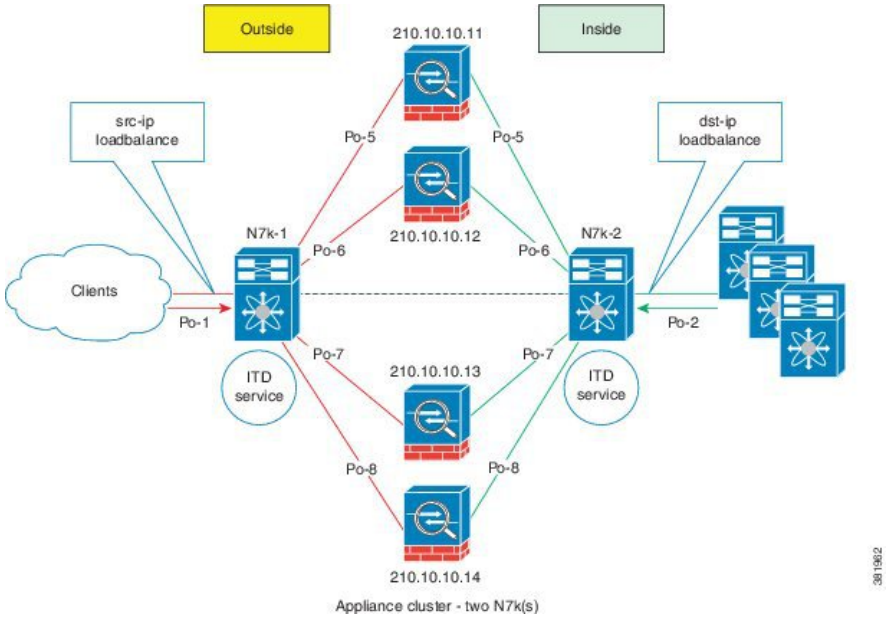

Server Load-Balancing Deployment Mode

The ITD service can be configured to host a virtual IP (VIP) on a Cisco NX-OS 7000 Series switch. Internet traffic destined for the VIP will be load balanced to the active nodes. Unlike traditional server load balancers, source NAT is not needed as the ITD service is not a stateful load balancer.

Note You need to configure ITD service similarly on each Cisco NX-OS 7000 Series switch. The ITD service configuration needs to be done manually on each switch.

Figure 4: ITD Load Distribution with VIP

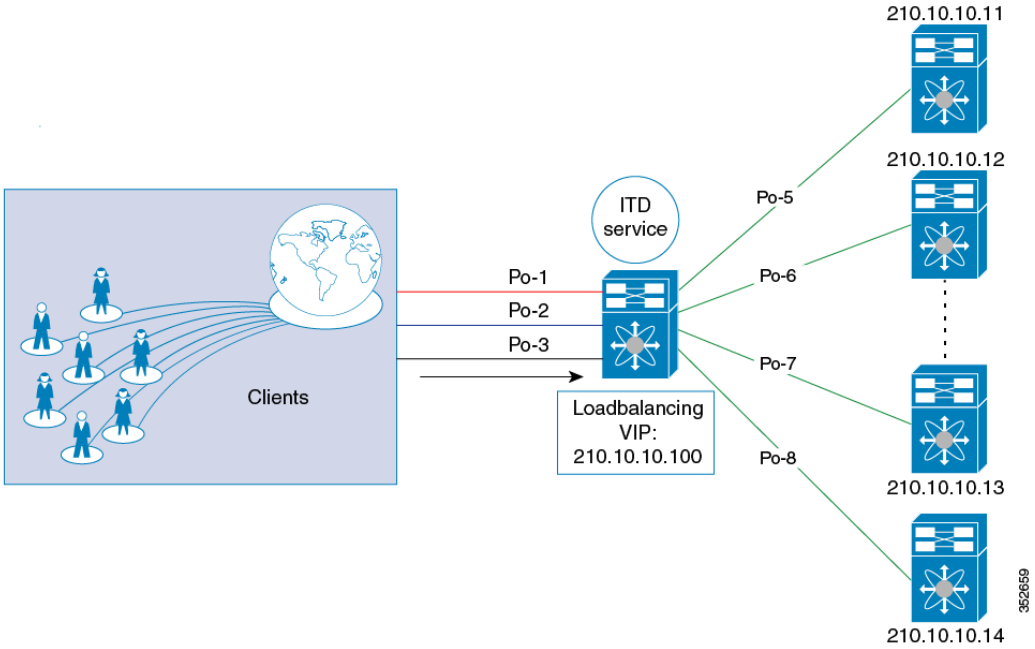

Device Groups

The ITD feature supports device groups. When you configure a device group you can specify the following:

- The device group's nodes
- The device group's probe

VRF Support

The ITD service can be configured in the default VRF as well as non-default VRFs.

Ingress interface(s) and device-group nodes must all belong to the same VRF for the ITD service to redirect traffic. You must ensure that all ingress interface(s) and node members of the associated device group are all reachable in the configured VRF.

Load Balancing

The ITD feature enables you to configure specific load-balancing options by using the **loadbalance** command.

The optional keywords for the **loadbalance** command are as follows:

- **buckets—**Specifies the number of buckets to create. Buckets must be configured in powers of two. One or more buckets are mapped to a node in the cluster. If you configure more buckets than the number of nodes, the buckets are applied in round robin fashion across all the nodes.
- **mask-position—** Specifies the mask position of the load balancing. This keyword is useful when a packet classification has to be made based on specific octets or bits of an IP addresses. By default the system uses the last octet or least significant bits (LSBs) for bucketing. If you prefer to use nondefault bits/octets, you can use the **mask-position** keyword to provide the starting point at which bits the traffic classification is to be made. For example, you can start at the 8th bit for the second octet and the 16th bit for the third octet of an IP address.
- **src** or **dst ip—** Specifies load balancing based on source or destination IP address.
- **src ip** or **src ip-l4port—** Specifies load balancing based on source IP address, or source IP address and source L4 port.
- **dst ip** or **dst ip-l4port—** Specifies load balancing based on destination IP address, or destination IP address and destination L4 port.

Hot Standby

ITD supports N+1 redundancy where M nodes can act as standby nodes for N active nodes.

When an active node fails, ITD looks for an operational standby node and selects the first available standby node to replace the failed node. ITD reconfigures the switch to redirect the traffic segment that was originally headed toward the failed node to the newly active node. The service does not impose any fixed mapping of standby nodes to active nodes.

When the failed node becomes operational again, it is reinstated as an active node and traffic from the acting standby node is redirected back to the original node and the standby node reverts to the pool of standby nodes.

When multiple nodes fail, traffic destined to all failed nodes gets redirected to the first available standby node.

A node can be configured as a standby at the node-level or device-group-level. A node-level standby receives traffic only if its associated active node fails. A device-group-level standby receives traffic if any of the active nodes fail.

Router Access Control Lists

Cisco Nexus 7000 Series devices support router access control lists (RACL) with ITD.

You can configure on the same ingress interface the ITD feature and an RACL together. The resulting RACL, which is downloaded to the TCAM, is a cross product of the ACL generated by ITD and the user-configured RACL. The permit and deny statements configured on the RACL are combined with the ACL permits and redirect entries created by ITD. This helps you to filter and load distribute selected traffic.

For more information on configuring an RACL with ITD, see Configuration Examples for ITD, on page 17.

Multiple Ingress Interfaces

You can configure the ITD service to apply traffic redirection policies on multiple ingress interfaces. This feature allows you to use a single ITD service to redirect traffic arriving on different interfaces to a group of nodes. The **ingress interface** command enables you to configure multiple ingress interfaces.

System Health Monitoring

ITD supports health monitoring functionality to do the following:

- Monitor the health of the node through the configured probe.
- Monitor the state of ingress interface(s).

With health monitoring, the following critical errors are detected and remedied:

- ITD service is shut/no shut or deleted.
- Switch reboot.
- Supervisor switchover.
- In-service software upgrade (ISSU).
- ITD service node failure.
- Ingress interface is down.

Monitor Node

The ITD health monitoring module periodically monitors nodes to detect any failure and to handle failure scenarios.

Note IPv6 probes are not supported.

Health of an Interface Connected to a Node

ITD leverages the IP service level agreement (IP SLA) feature to periodically probe each node. The probes are sent at a one second frequency and sent simultaneously to all nodes. You can configure the probe as part of the cluster group configuration. A probe is declared to have failed after retrying three times.

Node Failure Handling

Upon marking a node as down, the ITD performs the following tasks automatically to minimize traffic disruption and to redistribute the traffic to remaining operational nodes:

- Determines if a standby node is configured to take over from the failed node.
- Identifies the node as a candidate node for traffic handling, if the standby node is operational.
- Redefines the standby node as active for traffic handling, if an operational standby node is available.
- Programs automatically to reassign traffic from the failed node to the newly active standby node.

Failaction Reassignment

Failaction for ITD enables traffic on the failed nodes to be reassigned to the first available active node. Once the failed node comes back, it automatically resumes serving the connections. The **failaction** command enables this feature.

When the node is down, the traffic bucket associated with the node is reassigned to the first active node found in the configured set of nodes. If the newly reassigned node also fails, traffic is reassigned to the next available active node. Once the failed node becomes active again, traffic is diverted back to the new node and resumes serving connections.

Note You must configure probe under an ITD device group, before enabling the failaction feature.

Failaction Reassignment Without a Standby Node

When the node is down, the traffic bucket associated with the node is reassigned to the first active node found in the configured set of nodes. If the newly reassigned node also fails, the traffic bucket is reassigned to the next available active node. Once the failed node comes back and becomes active, the traffic is diverted back to the new node and starts serving the connections.

If all the nodes are down, the packets get routed automatically.

- When the node goes down (probe failed), the traffic is reassigned to the first available active node.
- When the node comes up (probe success) from the failed state, it starts handling the connections.
- If all the nodes are down, the packets get routed automatically.

Failaction Reassignment with a Standby Node

When the node is down and if the standby is active, the traffic serves the connections and there is no change in the bucket assignment. When both the active and standby nodes are down, the traffic bucket associated with the node is reassigned to the first active node found in the configured set of nodes. If the newly reassigned node also fails, the traffic bucket is reassigned to the next available active node. Once the failed node comes back up and becomes active, the traffic is diverted back to the new node and begins serving connections.

- When the node goes down (probe failed) and when there is a working standby node, traffic is directed to the first available standby node.
- When all nodes are down including the standby node, the traffic is reassigned to the first available active node.
- When the node comes up (probe success) from failed state, the node that came up starts handling the connections.
- If all the nodes are down, the packets are routed automatically.

No Failaction Reassignment

When failaction node reassignment is not configured, there are two possible scenarios:

- Scenario 1: Probe configured; and:
  - with standby configured; or
  - without standby configured.
- Scenario 2: No probe configured.

No Failaction Reassignment with a Probe Configured

The ITD probe can detect the node failure or the lack of service reachability.

- If the node fails and a standby is configured, the standby node takes over the connections.
- If the node fails and there is no standby configuration, the traffic gets routed and does not get reassigned, as failaction is not configured. Once the node recovers, the recovered node starts handling the traffic.

No Failaction Reassignment without a Probe Configured

Without a probe configuration, ITD cannot detect the node failure. When the node is down, ITD does not reassign or redirect the traffic to an active node.

Licensing Requirements for ITD

ITD requires an Enhanced Layer 2 Package license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

Prerequisites for ITD

ITD has the following prerequisites:

- You must enable the ITD feature with the **feature itd** command.
- The following commands must be configured prior to entering the **feature itd** command:
  - **feature pbr**
  - **feature sla sender**
  - **feature sla responder**
  - **ip sla responder**

Guidelines and Limitations for ITD

ITD has the following configuration guidelines and limitations:

- Virtual IP type and the ITD device group nodes type should be either IPv4 or IPv6, but not both.
- Configuration rollback is only supported when the ITD service is in shut mode in both target and source configurations.
- Probes are not supported for a device group with IPv6 nodes.
- The **failaction** command is supported only for IPv4.
- SNMP is not supported for ITD.
- An in-service software upgrade (ISSU) from Release 6.2(8) to Release 6.2(8a) or an in-service software downgrade (ISSD) from Release 6.2(8a) to Release 6.2(8) is not supported. Before performing an ISSU or ISSD, you must remove the ITD configuration by using the **no feature itd** command. After the upgrade or downgrade, you must manually reapply the configuration.

Configuring ITD

The server can be connected to the switch through a routed interface or port-channel, or via a switchport port with SVI configured.

Enabling ITD

Before You Begin

Before you configure the **feature itd** command you must enter the **feature pbr** and **feature ipsla** commands.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature itd**	Enables the ITD feature.

Configuring a Device Group

Before You Begin

Enable the ITD feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **itd device-group** *name*	Creates an ITD device group and enters into device group configuration mode.
Step 3	switch(config-device-group)# **node ip** *ipv4-address* [**mode hot-standby**] [**standby** *ipv4-address*] [**weight** *value*]	Specifies the nodes for ITD. Repeat this step to specify all nodes. To configure IPv6 nodes, use the **node ipv6** *ipv6-address* [**mode hot-standby**]. **Note** An ITD device group can have either IPv4 or IPv6 nodes, but not both. The **weight** *value* keyword specifies the proportionate weight for the node for weighted traffic distribution.
Step 4	switch(config-device-group)# **probe** {**icmp** | **tcp port** *port-number* | **udp port** *port-number* | **dns** {*hostname* |*target-address*} } [**frequency** *seconds*] [[**retry-down-count** | **retry-up-count**] *number*] [**timeout** *seconds*]	Configures the cluster group service probe. You can specify the following protocols as the probe for the ITD service: • ICMP • TCP • UDP • DNS The keywords are as follows: • **retry-down-count—**Specifies the consecutive number of times the probe must have failed prior to the node being marked DOWN. • **retry-up-count—**Specifies the consecutive number of times the probe must have succeeded prior to the node being marked UP. • **timeout—**Specifies the number of seconds to wait for the probe response. • **frequency—**Specifies the time interval in seconds between successive probes sent to the node. **Note** IPv6 probes are not supported.

Configuring an ITD Service

Before You Begin

- Enable the ITD feature.
- Configure the device-group to be added to the ITD service.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **itd** *service-name*	Configures an ITD service and enters into ITD configuration mode.
Step 3	switch(config-itd)# **device-group** *device-group-name*	Adds an existing device group to the ITD service. The *device-group-name* specifies the name of the device group. You can enter up to 32 alphanumeric characters.
Step 4	switch(config-itd)# **ingress interface** *interface*	Adds an ingress interface or multiple interfaces to an ITD service. • Use a comma (“,”) to separate multiple interfaces. • Use a hyphen (“-”) to separate a range of interfaces.
Step 5	switch(config-itd)# **load-balance** {**method** {**src** {**ip** | **ip-l4port** [**tcp** | **udp**] **range** *x y*} | **dst** {**ip** | **ip-l4port** [**tcp** | **udp**] **range** *x y*}} | **buckets** *bucket-number* | **mask-position** *position*}	Configures the load-balancing options for the ITD service. The keywords are as follows: • **buckets—**Specifies the number of buckets to create. Buckets must be configured in powers of two. • **mask-position—** Specifies the mask position of the load balance. • **method—**Specifies the source IP address or destination IP address, or source IP address and source port, or the destination IP address and destination port based load-balancing.
Step 6	switch(config-itd)# **virtual ip** *ipv4-address ipv4-network-mask* [**tcp** | **udp** {*port-number* | **any**}] [**advertise** {**enable** | **disable**}]	Configures the virtual IPv4 address of the ITD service. **Note** To configure an IPv6 virtual address, use the **virtual ipv6** *ipv6-address ipv6-network-mask* | *ipv6-prefix/length*} [**ip** | **tcp** {*port-number* | **any**} | **udp** {*port-number* | **any**}] [**advertise** {**enable** | **disable**}] The **advertise enable** keywords specify that the virtual IP route is advertised to neighboring devices.

	Command or Action	Purpose
		The **tcp**, **udp**, and **ip** keywords specify that the virtual IP address will accept flows from the specified protocol.
Step 7	switch(config-itd)# **failaction node reassign**	Enables traffic to be reassigned, following a node failure. The traffic to the failed node gets reassigned to the first available active node.
Step 8	switch(config-itd)# **vrf** *vrf-name*	Specifies the VRF for the ITD service.
Step 9	switch(config-itd)# **no shutdown**	Enables the ITD service.

Verifying the ITD Configuration

To display the ITD configuration, perform one of the following tasks:

Command	Purpose
show itd [*itd-name*] [**brief**]	Displays the status and configuration for all or specified ITD instances. • Use the *itd-name* argument to display the status and configuration for the specific instance. • Use the **brief** keyword to display summary status and configuration information.
show itd [*itd-name* | **all**] {**src** | **dst**} *ip-address*] **statistics** [**brief**]	Displays the statistics for ITD instances. • Use the *itd-name* argument to display statistics for the specific instance. • Use the **brief** keyword to display summary information. **Note** Before using the **show itd statistics** command, you need to enable ITD statistics by using the **itd statistics** command.
show running-config services	Displays the configured ITD device-group and services.

These examples show how to verify the ITD configuration:

```
switch# show itd

Name           Probe LB Scheme  Status   Buckets
-------------- ----- ---------- -------- -------
WEB            ICMP  src-ip     ACTIVE   2
```

```
Device Group                                          VRF-Name
----------------------------------------------- -------------
WEB-SERVERS


Pool                           Interface    Status Track_id
------------------------------ ------------ ------ ---------
WEB_itd_pool                   Po-1       UP       3

Virtual IP                                Netmask/Prefix Protocol     Port
----------------------------------------- -------------- ---------- ----------
10.10.10.100 / 255.255.255.255                          IP           0

  Node  IP                 Config-State Weight Status     Track_id  Sla_id
  ------------------------ ------------ ------ ---------- --------- ---------
  1     10.10.10.11        Active       1      OK           1         10001

        Bucket List
        ----------------------------------------------------------------------
        WEB_itd_vip_1_bucket_1


  Node  IP                 Config-State Weight Status     Track_id  Sla_id
  ------------------------ ------------ ------ ---------- --------- ---------
  2     10.10.10.12        Active       1      OK           2         10002

        Bucket List
        ----------------------------------------------------------------------
        WEB_itd_vip_1_bucket_2

switch# show itd brief

Name           Probe LB Scheme  Interface  Status   Buckets
-------------- ----- ---------- ---------- -------- --------
WEB            ICMP  src-ip     Eth3/3     ACTIVE   2

Device Group                                          VRF-Name
----------------------------------------------- -------------
WEB-SERVERS


Virtual IP                                Netmask/Prefix Protocol     Port
----------------------------------------- -------------- ---------- ----------
10.10.10.100 / 255.255.255.255                          IP           0

  Node  IP                 Config-State Weight Status     Track_id  Sla_id
  ------------------------ ------------ ------ ---------- --------- ---------
  1     10.10.10.11        Active       1      OK           1         10001
  2     10.10.10.12        Active       1      OK           2         10002

switch(config)# show itd statistics

Service        Device Group             VIP/mask               #Packets
-------------------------------------------------------------------------------
test            dev              9.9.9.10 / 255.255.255.0       114611  (100.00%)

Traffic Bucket        Assigned to      Mode       Original Node   #Packets
-------------------------------------------------------------------------------
test_itd_vip_0_acl_0  10.10.10.9       Redirect    10.10.10.9      57106  (49.83%)

Traffic Bucket        Assigned to      Mode       Original Node   #Packets
-------------------------------------------------------------------------------
test_itd_vip_0_acl_1  12.12.12.9       Redirect    12.12.12.9      57505  (50.17%)
switch (config)# show running-config services

version 6.2(10)
feature itd

itd device-group WEB-SERVERS
node ip 10.10.10.11
node ip 10.10.10.12
```

```
probe icmp

itd WEB
device-group WEB-SERVERS
virtual ip 10.10.10.100 255.255.255.255
ingress interface po-1
no shut
```

Warnings and Error Messages for ITD

The following warnings and error messages are displayed for ITD:

When you reach the maximum number of configurable nodes, this message is displayed:

```
Already reached maximum nodes per service
```

If you configure the same node IP when it is already configured part of an ITD service, this message is displayed:

```
This IP is already configured, please try another IP
```

When you try to change or remove a device group, probe, or ingress interface after the IDT service is enabled, one of these messages is displayed:

```
Probe configuration is not allowed, service is enabled
Ingress interface configuration is not allowed, service is enabled
Node configuration is not allowed, service is enabled
```

If the ITD service is already enabled or disabled, one of these messages is displayed:

```
In service already enabled case
In service already disabled case
```

When you try to change the failaction configuration after the ITD service is enabled, this message is displayed:

```
Failaction configuration is not allowed, service is enabled.
```

Configuration Examples for ITD

This example shows how to configure an ITD device group:

```
switch(config)# feature itd
switch(config)# itd device-group dg
switch(config-device-group)# node ip 210.10.10.11
switch(config-device-group)# node ip 210.10.10.12
switch(config-device-group)# node ip 210.10.10.13
switch(config-device-group)# node ip 210.10.10.14
switch(config-device-group)# probe icmp
```

This example shows how to configure a virtual IPv4 address:

```
switch(config)# feature itd
switch(config)# itd test
switch(config-itd)# device-group dg
switch(config-itd)# ingress interface Po-1
switch(config-itd)# virtual ip 210.10.10.100 255.255.255.255 advertise enable tcp any
```

This example shows how to configure a virtual IPv6 address:

```
switch(config)# feature itd
switch(config)# itd test
switch(config-itd)# device-group dg
switch(config-itd)# ingress interface Po-1
switch(config-itd)# virtual ipv6 ffff:eeee::cccc:eeee dddd:efef::fefe:dddd tcp 10 advertise
 enable
```

This example shows how to configure an RACL with ITD. The user-defined RACL, test, is displayed:

```
switch(config-itd)# show ip access-lists test

IP access list test
        10 permit ip 1.1.1.1/32 2.2.2.2/16
```

```
        20 permit ip 3.3.3.3/20 4.4.4.4/32
```

Below is the ITD configuration that has the ingress interface as Po-1

```
itd demo
  device-group dg
  virtual ip 11.22.33.44 255.255.255.255 tcp any
  virtual ip 11.22.33.55 255.255.0.0
  virtual ip 11.22.33.66 255.255.255.255 tcp any
  ingress interface Po-1
  no shut
```

Here we see both the route-map created by ITD and the RACL are both part of the same physical interface Po-1:

```
interface Po-1
  ip access-group test in
  ip policy route-map demo_itd_routemap
  no shutdown
```

This example shows how to configure device-group-level standby node. Node 210.10.10.15 is configured as standby for the entire device group. If any of the active nodes fail, the traffic going to the failed node will be redirected to 210.10.10.15:

```
switch(config)# feature itd
switch(config)# itd device-group dg
switch(config-device-group)# node ip 210.10.10.11
switch(config-device-group)# node ip 210.10.10.12
switch(config-device-group)# node ip 210.10.10.13
switch(config-device-group)# node ip 210.10.10.14
switch(config-device-group)# node ip 210.10.10.15 mode hot-standby
switch(config-device-group)# probe
```

This example shows how to configure node-level standby node. Node 210.10.10.15 is configured as standby for node 210.10.10.11 only. Only when node 210.10.10.11 fails, the traffic going to the failed node will be get redirected to 210.10.10.15:

```
switch(config)# feature itd
switch(config)# itd device-group dg
switch(config-device-group)# node ip 210.10.10.11 standby 210.10.10.15
switch(config-device-group)# node ip 210.10.10.12
switch(config-device-group)# node ip 210.10.10.13
switch(config-device-group)# node ip 210.10.10.14
switch(config-device-group)# probe
```

This example shows how to configure weight for proportionate distribution of traffic. Nodes 1 and 2 would get three times as much traffic as nodes 3 and 4:

```
switch(config)# feature itd
switch(config)# itd device-group dg
switch(config-device-group)# node ip 210.10.10.11 weight 3
switch(config-device-group)# node ip 210.10.10.12 weight 3
switch(config-device-group)# node ip 210.10.10.13
switch(config-device-group)# node ip 210.10.10.14
switch(config-device-group)# probe
```

This example shows how to configure VRF for ITD service:

```
switch(config)# feature itd
switch(config)# itd test
switch(config-device-group)# device-group dg
switch(config-device-group)# ingress interface Po-1
switch(config-device-group)# vrf RED
```

This example shows how to enable statistics collection for ITD service:

Note You must enable statistics collection for 'show itd statistics' to show the packet counters.

```
switch(config)# itd statistics test
```

This example shows how to disable statistics collection for ITD service:

```
switch(config)# no itd statistics test
```

Configuration Example: One-Arm Deployment Mode

The configuration below uses the topology in the following figure:

Figure 5: One-Arm Deployment Mode

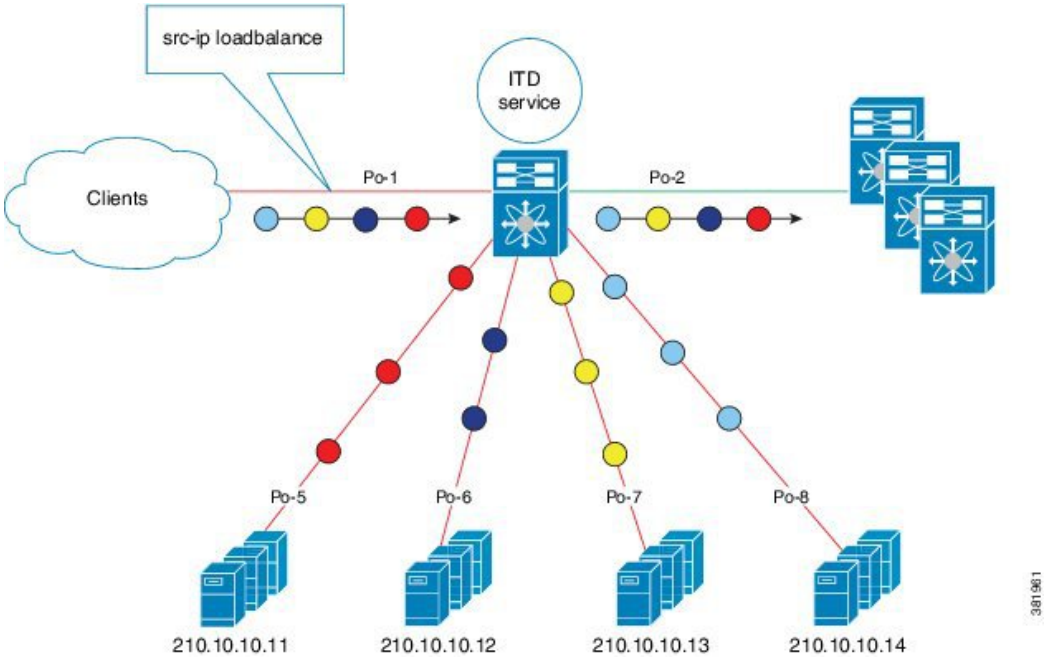

Step 1: Define device group

```
switch(config)# itd device-group DG
switch(config-device-group)# node ip 210.10.10.11
switch(config-device-group)# node ip 210.10.10.12
switch(config-device-group)# node ip 210.10.10.13
switch(config-device-group)# node ip 210.10.10.14
switch(config-device-group)# probe icmp
```

Step 2: Define ITD service

```
switch(config)# itd HTTP
switch(config-itd)# ingress interface port-channel 1
switch(config-itd)# device-group DG
switch(config-itd)# no shutdown
```

Configuration Example: One-Arm Deployment Mode with VPC

The configuration below uses the topology in the following figure:

Figure 6: One-Arm Deployment Mode with VPC

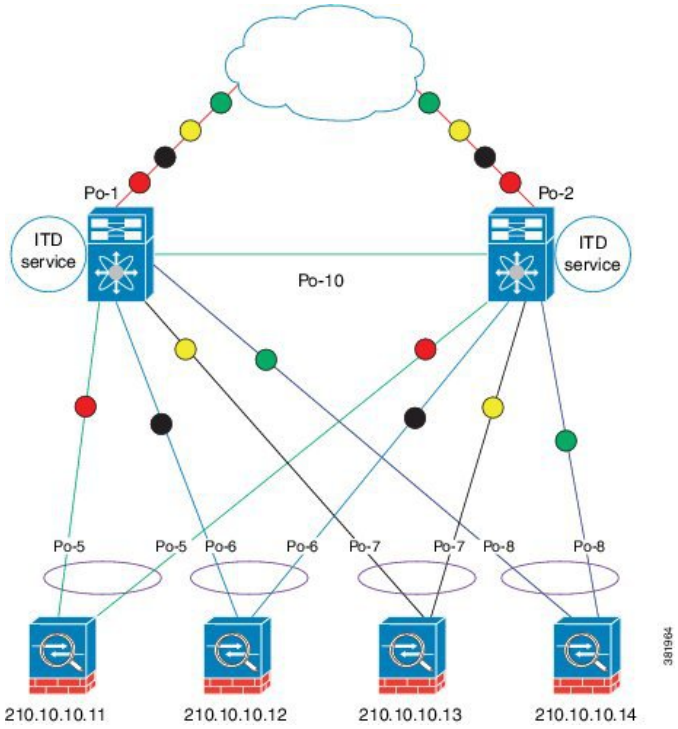

Device 1

Step 1: Define device group

```
N7k-1(config)# itd device-group DG
N7k-1(config-device-group)# node ip 210.10.10.11
N7k-1(config-device-group)# node ip 210.10.10.12
N7k-1(config-device-group)# node ip 210.10.10.13
N7k-1(config-device-group)# node ip 210.10.10.14
N7k-1s(config-device-group)# probe icmp
```

Step 2: Define ITD service

```
N7k-1(config)# itd HTTP
N7k-1(config-itd)# ingress interface port-channel 1
N7k-1(config-itd)# device-group DG
N7k-1(config-itd)# no shutdown
```

Device 2

Step 1: Define device group

```
N7k-2(config)# itd device-group DG
N7k-2(config-device-group)# node ip 210.10.10.11
N7k-2(config-device-group)# node ip 210.10.10.12
```

```
N7k-2(config-device-group)# node ip 210.10.10.13
N7k-2(config-device-group)# node ip 210.10.10.14
N7k-2(config-device-group)# probe icmp
```

Step 2: Define ITD service

```
N7k-2(config)# itd HTTP
N7k-2(config-itd)# ingress interface port-channel 2
N7k-2(config-itd)# device-group DG
N7k-2(config-itd)# no shutdown
```

Configuration Example: Sandwich Deployment Mode

The configuration below uses the topology in the following figure:

Figure 7: Sandwich Deployment Mode

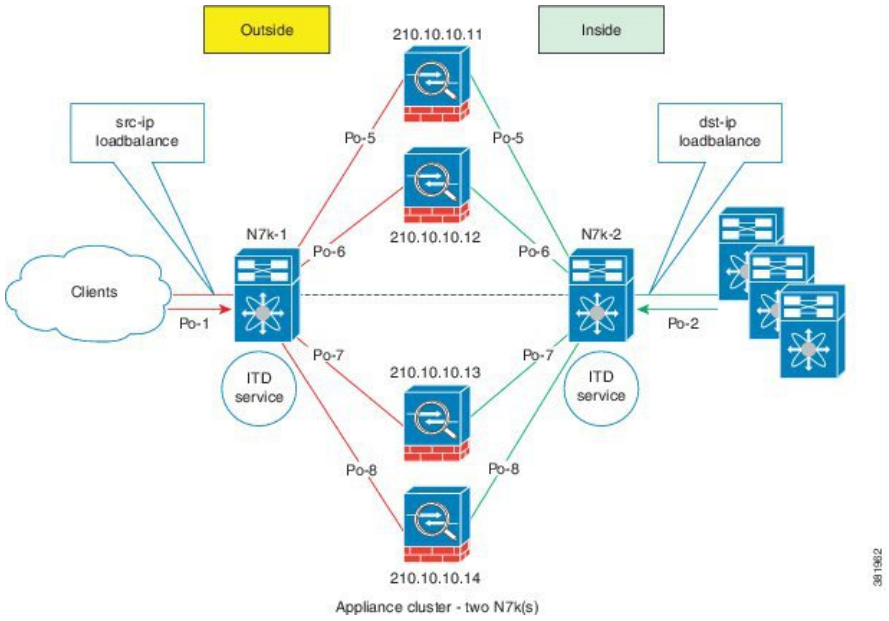

Device 1

Step 1: Define device group

```
N7k-1(config)# itd device-group DG
N7k-1(config-device-group)# node ip 210.10.10.11
N7k-1(config-device-group)# node ip 210.10.10.12
N7k-1(config-device-group)# node ip 210.10.10.13
N7k-1(config-device-group)# node ip 210.10.10.14
N7k-1s(config-device-group)# probe icmp
```

Step 2: Define ITD service

```
N7k-1(config)# itd HTTP
N7k-1(config-itd)# ingress interface port-channel 1
N7k-1(config-itd)# device-group DG
N7k-1(config-itd)# load-balance method src ip
N7k-1(config-itd)# no shutdown
```

Device 2

Step 1: Define device group

```
N7k-2(config)# itd device-group DG
N7k-2(config-device-group)# node ip 220.10.10.11
N7k-2(config-device-group)# node ip 220.10.10.12
N7k-2(config-device-group)# node ip 220.10.10.13
N7k-2(config-device-group)# node ip 220.10.10.14
N7k-2(config-device-group)# probe icmp
```

Step 2: Define ITD service

```
N7k-2(config)# itd HTTP
N7k-2(config-itd)# ingress interface port-channel 2
N7k-2(config-itd)# device-group DG
N7k-2(config-itd)# load-balance method dst ip
N7k-2(config-itd)# no shutdown
```

Configuration Example: Server Load-Balancing Deployment Mode

The configuration below uses the topology in the following figure:

Figure 8: ITD Load Distribution with VIP

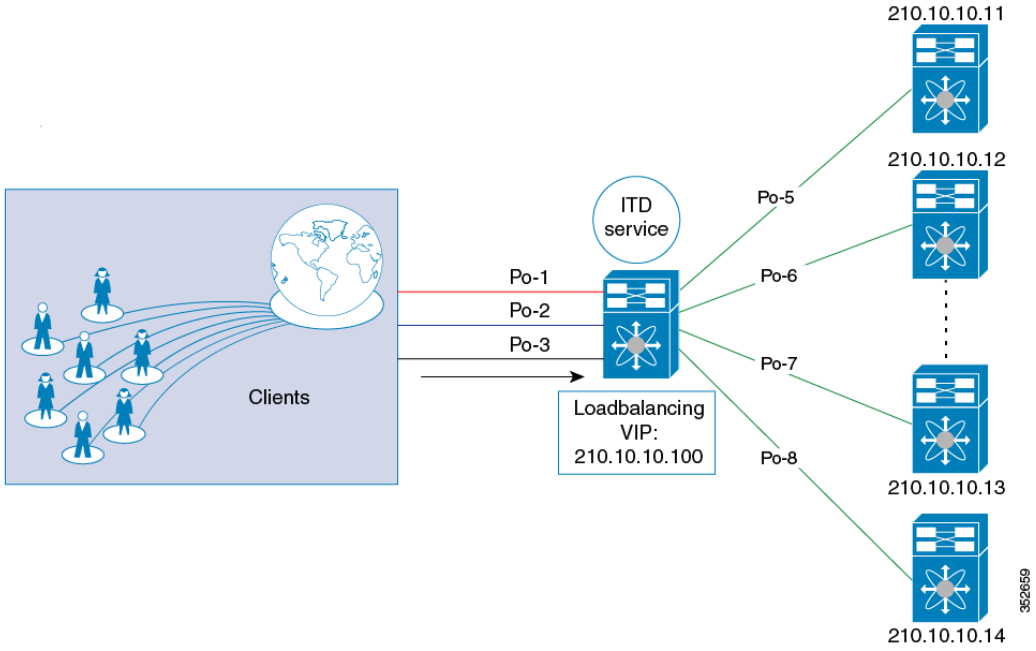

Step 1: Define device group

```
switch(config)# itd device-group DG
switch(config-device-group)# node ip 210.10.10.11
switch(config-device-group)# node ip 210.10.10.12
switch(config-device-group)# node ip 210.10.10.13
switch(config-device-group)# node ip 210.10.10.14
switch(config-device-group)# probe icmp
```

Step 2: Define ITD service

```
switch(config)# itd HTTP
switch(config-itd)# ingress interface port-channel 1
switch(config-itd)# ingress interface port-channel 2
switch(config-itd)# ingress interface port-channel 3
switch(config-itd)# device-group DG
Switch(config-itd)# virtual ip 210.10.10.100 255.255.255.255
switch(config-itd)# no shutdown
```

Related Documents for ITD

Related Topic	Document Title
Intelligent Traffic Director commands	*Cisco Nexus 7000 Series NX-OS Intelligent Traffic Director Command Reference*

Standards for ITD

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for ITD

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Feature Name	Release	Feature Information
ITD	6.2(10)	Added the following enhancements: • Weighted load-balancing. • Node-level standby. • Layer 4 port load-balancing. • Sandwich mode node-state synchronization across two VDCs on the same device. • DNS probe. • Start/stop/clear ITD statistics collection. • VRF support for the ITD service and probes.
Intelligent Traffic Director (ITD)	6.2(8)	This feature was introduced.