



Configuring ITD

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

- [Finding Feature Information, page 1](#)
- [Information About ITD, page 1](#)
- [Licensing Requirements for ITD, page 9](#)
- [Prerequisites for ITD, page 9](#)
- [Guidelines and Limitations for ITD, page 9](#)
- [Configuring ITD, page 9](#)
- [Verifying the ITD Configuration, page 12](#)
- [Warnings and Error Messages for ITD, page 13](#)
- [Configuration Examples for ITD, page 14](#)
- [Standards for ITD, page 18](#)
- [Feature History for ITD, page 18](#)

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, 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 multi-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 5500 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, etc.
- Can be used to load balance traffic to other software load balancers.
- Allows DSR load-balancing deployments.
- Weighted load-balancing provides load-balances to large number of devices or servers ACL along 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) and 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 switch.
- Supports VRFs and vPCs.
- Supports IPv4 only.

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.


- Weight-based load balancing.
- Complete transparency to the end devices.
- 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




381361

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 cluster nodes.

Figure 2: One-Arm Deployment Mode with VPC




Server Load-Balancing Deployment Mode

The ITD service can be configured to host a virtual IP (VIP) on a Cisco NX-OS 5500 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**

The ITD service configuration needs to be done manually on each switch.

Figure 3: ITD Load Distribution with VIP



Device Groups

The ITD feature supports device groups. When you configure a device group you can specify the the device group's cluster nodes.

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 layer 4 port.
- **dst ip** or **dst ip-l4port**— Specifies load balancing based on destination IP address or destination layer 4 port .

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 ITD channel and peer ITD service.
- Monitor the state of the interface connected to each node.
- 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.
- iSCM process crash.
- iSCM process restart.
- Switch reboot.
- Supervisor switchover.
- In-service software upgrade (ISSU).
- ITD service node failure.
- ITD service node port or interface down.
- Ingress interface down.

Monitor Node

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

ICMP probes are supported to probe each node periodically for health monitoring. A probe can be configured at the device-group level or at node-level. A probe configured at the device-group level is sent to each node member of the device-group. A probe configured at a node-level is sent only to the node it is associated with. If a node-specific probe is configured, only that probe is sent to the node. For all the nodes that do not have node-specific probe configuration, the device-group level probe (if configured) is sent.

IPv4 Control Probe for IPv6 Data Nodes

For an IPv6 node (in an IPv6 device-group), if the node is a dual-homed node (that is, it supports IPv4 and IPv6 network interfaces), an IPv4 probe can be configured to monitor the health. Since IPv6 probes are not supported, this provides a way to monitor health of IPv6 data nodes using a IPv4 probe.



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.

Monitor Peer ITD Service

For sandwich mode cluster deployments, the ITD service runs on each Cisco NX-OS 5500 series switch. The health of the ITD channel is crucial to ensure flow coherent traffic passing through cluster nodes in both directions.

Each ITD service probes its peer ITD service periodically to detect any failure. A ping is sent every second to the peer ITD service. If a reply is not received it is retried three times. The frequency and retry count are not configurable.



Note Since only a single instance of the ITD service is running on the switch in one-arm mode deployment, monitoring of the peer ITD is not applicable.

ITD channel failure handling

If the heartbeat signal is missed three times in a row, then the ITD channel is considered to be down.

While the ITD channel is down, traffic continues to flow through cluster nodes. However, since the ITD service on each switch is not able to exchange information about its view of the cluster group, this condition requires immediate attention. A down ITD channel can lead to traffic loss in the event of a node failure.

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.

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

Starting with NX-OS Release 7.2(0)N1(1), ITD requires the Services License N55-SERVICES1K9. Releases prior to NX-OS Release 7.2(0)N1(1) require the Enhanced Layer 2 Package license for ITD support.

Prerequisites for ITD

ITD has the following prerequisites:

- You must enable the ITD feature with the **feature itd** command.
- The **feature pbr** command must be configured prior to entering the **feature itd** command:

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 IPv4.
- Configuration rollback is only supported when the ITD service is in shut mode in both target and source configurations.
- SNMP is not supported for ITD.
- 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** command.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **feature itd**

DETAILED STEPS

	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.

DETAILED STEPS

	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	Specifies the nodes for ITD. Repeat this step to specify all nodes.
Step 4	switch(config-device-group)# probe icmp	Configures the cluster group service probe. 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.

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# itd <i>service-name</i>	Configures an ITD service and enters into ITD configuration mode.
Step 3	switch(config-itd)# device-group <i>device-group-name</i>	Adds an existing device group to the ITD service. The <i>device-group-name</i> specifies the name of the device group. You can enter up to 32 alphanumeric characters.
Step 4	switch(config-itd)# ingress interface <i>interface</i>	Adds an ingress interface or multiple interfaces to an ITD service. <ul style="list-style-type: none"> • 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 <i>x y</i> } dst { ip ip-l4port [tcp udp] range <i>x y</i> }} buckets <i>bucket-number</i> mask-position <i>position</i> }	Configures the load-balancing options for the ITD service. The keywords are as follows: <ul style="list-style-type: none"> • buckets—Specifies the number of buckets to create. Buckets must be configured in powers of two. • mask-position— Specifies the mask position of the loadbalance. • method—Specifies the source IP address or destination IP address based load/traffic distribution.
Step 6	switch(config-itd)# virtual ip <i>ipv4-address</i> <i>ipv4-network-mask</i> [tcp udp { <i>port-number</i> any }] [advertise { enable disable }]	Configures the virtual IPv4 address of the ITD service. The advertise enable keywords specify that the virtual IP route is advertised to neighboring devices. The tcp , udp , and ip keywords specify that the virtual IP address will accept flows from the specified protocol.
Step 7	switch(config-itd)# vrf <i>vrf-name</i>	Specifies the VRF for the ITD service.
Step 8	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
<code>show itd [itd-name] [brief]</code>	<p>Displays the status and configuration for all or specified ITD instances.</p> <ul style="list-style-type: none"> Use the <i>itd-name</i> argument to display the status and configuration for the specific instance. Use the brief keyword to display summary status and configuration information.
<code>show itd [itd-name all] {src dst} ip-address] statistics [brief]</code>	<p>Displays the statistics for ITD instances.</p> <ul style="list-style-type: none"> Use the <i>itd-name</i> argument to display statistics for the specific instance. Use the brief keyword to display summary information. <p>Note Before using the show itd statistics command, you need to enable ITD statistics by using the itd statistics command.</p>
<code>show running-config services</code>	Displays the configured ITD device-group and services.

These examples show how to verify the ITD configuration:

```
switch# show itd

Name          LB Scheme  Status  Buckets
-----
WEB           src-ip    ACTIVE  2

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

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

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

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

Bucket List
-----
WEB_itd_vip_1_bucket_1
```

```

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

Bucket List
-----
WEB_itd_vip_1_bucket_2

switch# show itd brief

Name LB Scheme Interface Status Buckets
-----
WEB 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
-----
1 10.10.10.11 Active 1 OK -
2 10.10.10.12 Active 1 OK -

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 7.1(1)N1(1)
feature itd

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

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 or ingress interface after the IDT service is enabled, one of these messages is displayed:

```
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
```

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 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
```

Configuration Example: One-Arm Deployment Mode

The configuration below uses the topology in the following figure:

Figure 4: 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
```


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 5: 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
```

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
```



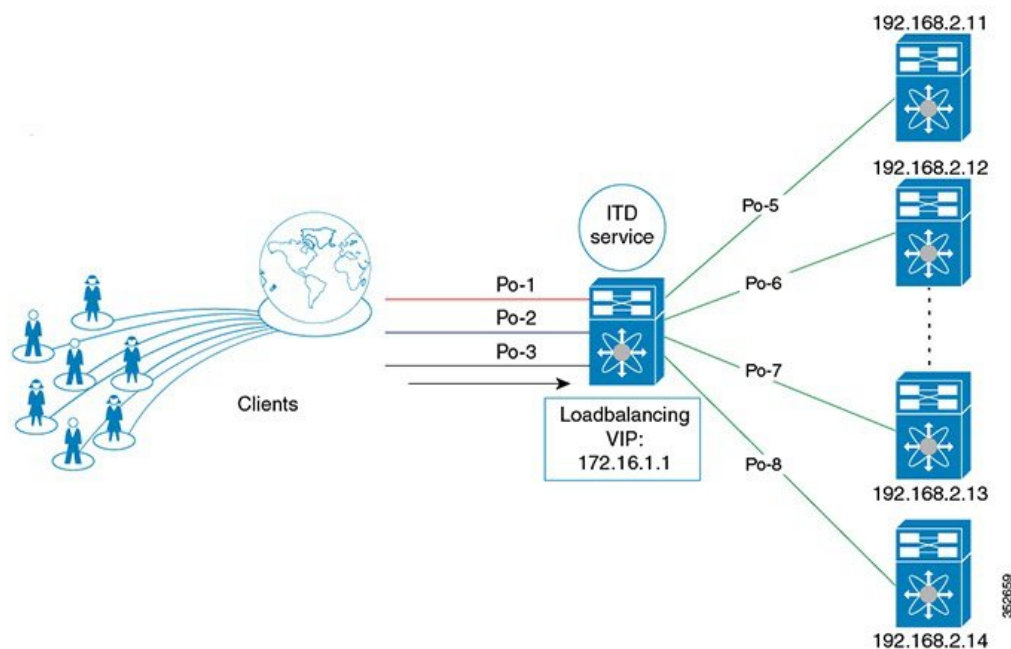
```
N7k-2(config-device-group)# node ip 210.10.10.13
N7k-2(config-device-group)# node ip 210.10.10.14
Step 2: Define ITD service
```

```
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: Server Load-Balancing Deployment Mode

The configuration below uses the topology in the following figure:

Figure 6: 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
```

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
```

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 Probe	7.2(1)N1(1)	The support for ITD Probe was added.
Intelligent Traffic Director (ITD)	7.1(1) N1(1)	This feature was introduced.