

Traffic Zones

You can assign multiple interfaces to a *traffic zone*, which lets traffic from an existing flow exit or enter the ASA on any interface within the zone. This capability allows Equal-Cost Multi-Path (ECMP) routing on the ASA as well as external load balancing of traffic to the ASA across multiple interfaces.

- About Traffic Zones, on page 1
- Prerequisites for Traffic Zones, on page 7
- Guidelines for Traffic Zones, on page 9
- Configure a Traffic Zone, on page 10
- Monitoring Traffic Zones, on page 10
- Example for Traffic Zones, on page 13
- History for Traffic Zones, on page 15

About Traffic Zones

This section describes how you should use traffic zones in your network.

Non-Zoned Behavior

The Adaptive Security Algorithm takes into consideration the state of a packet when deciding to permit or deny the traffic. One of the enforced parameters for the flow is that traffic enters and exits the same interface. Any traffic for an existing flow that enters a different interface is dropped by the ASA.

Traffic zones let you group multiple interfaces together so that traffic entering or exiting *any* interface in the zone fulfills the Adaptive Security Algorithm security checks.

Related Topics

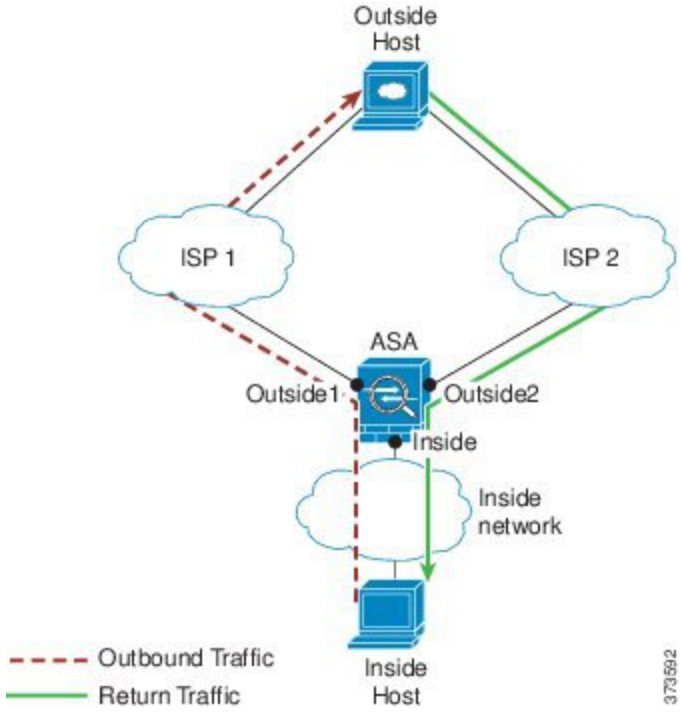
Stateful Inspection Overview

Why Use Zones?

You can use zones to accommodate several routing scenarios.

Asymmetric Routing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to asymmetric routing on the destination network, return traffic arrived from ISP 2 on the Outside2 interface.

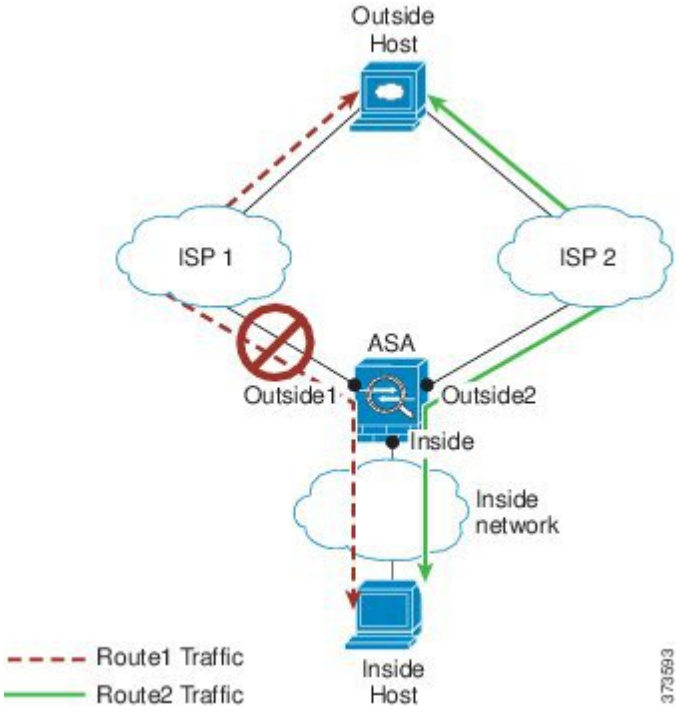


Non-Zoned Problem: The ASA maintains the connection tables on a per-interface basis. When the returning traffic arrives at Outside2, it will not match the connection table and will be dropped. For an ASA cluster, asymmetric routing when the cluster has multiple adjacencies to the same router can lead to unacceptible traffic loss.

Zoned Solution: The ASA maintains connection tables on a per-zone basis. If you group Outside1 and Outside2 into a zone, then when the returning traffic arrives at Outside2, it will match the per-zone connection table, and the connection will be allowed.

Lost Route

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to a lost or moved route between Outside1 and ISP 1, traffic needs to take a different route through ISP 2.

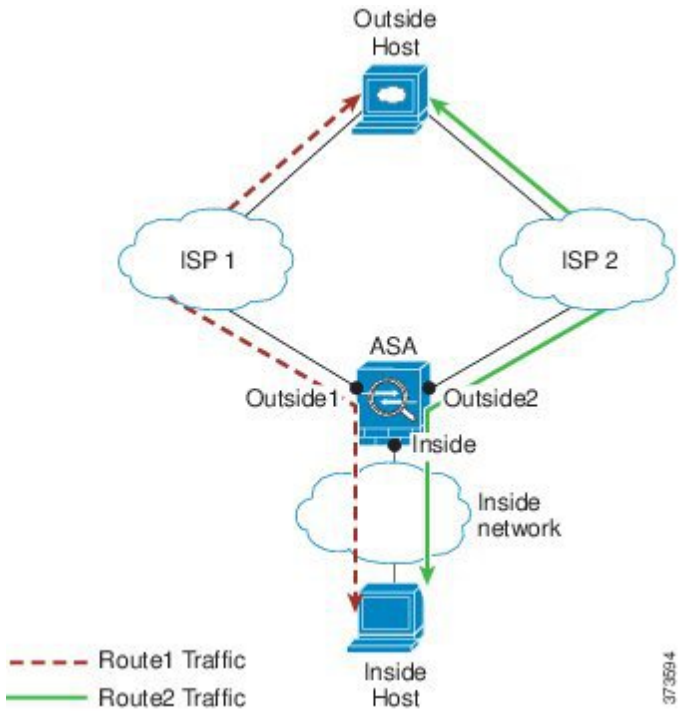


Non-Zoned Problem: The connection between the inside and outside host will be deleted; a new connection must be established using a new next-best route. For UDP, the new route will be used after a single packet drop, but for TCP, a new connection has to be reestablished.

Zoned Solution: The ASA detects the lost route and switches the flow to the new path through ISP 2. Traffic will be seamlessly forwarded without any packet drops.

Load Balancing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. A second connection was established through an equal cost route through ISP 2 on Outside2.



Non-Zoned Problem: Load-balancing across interfaces is not possible; you can only load-balance with equal cost routes on one interface.

Zoned Solution: The ASA load-balances connections across up to eight equal cost routes on all the interfaces in the zone.

Per-Zone Connection and Routing Tables

The ASA maintains a per-zone connection table so that traffic can arrive on any of the zone interfaces. The ASA also maintains a per-zone routing table for ECMP support.

ECMP Routing

The ASA supports Equal-Cost Multi-Path (ECMP) routing.

Non-Zoned ECMP Support

Without zones, you can have up to 8 equal cost static or dynamic routes per interface. For example, you can configure three default routes on the outside interface that specify different gateways:

```
route outside 0 0 10.1.1.2
route outside 0 0 10.1.1.3
route outside 0 0 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses, incoming interface, protocol, source and destination ports.

ECMP is not supported across multiple interfaces, so you cannot define a route to the same destination on a different interface. The following route is disallowed when configured with any of the routes above:

```
route outside2 0 0 10.2.1.1
```

Zoned ECMP Support

With zones, you can have up to 8 equal cost static or dynamic routes across interface within a zone.

When ECMP is determining how to distribute traffic, it sees each virtual access interface as a single participant or path option, even if that virtual access interface is associated with an interface that has several IPv6 addresses.

For example, you can configure three default routes across three interfaces in the zone:

```
route outside1 0 0 10.1.1.2
route outside2 0 0 10.2.1.2
route outside3 0 0 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The ASA load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the ASA seamlessly moves the flow to a different route.

How Connections Are Load-Balanced

The ASA load balances connections across equal cost routes using a hash made from the packet 6-tuple (source and destination IP address, source and destination port, protocol, and ingress interface). Unless the route is lost, a connection will stay on the chosen interface for its duration.

Packets within a connection are not load-balanced across routes; a connection uses a single route unless that route is lost.

The ASA does not consider the interface bandwidth or other parameters when load balancing. You should make sure all interfaces within the same zone have the same characteristics such as MTU, bandwidth, and so on.

The load-balancing algorithm is not user configurable.

Falling Back to a Route in Another Zone

When a route is lost on an interface, if there are no other routes available within the zone, then the ASA will use a route from a different interface/zone. If this backup route is used, then you may experience packet drops as with non-zoned routing support.

Interface-Based Security Policy

Zones allow traffic to and from any interface in the zone, but the security policy itself (access rules, NAT, and so on) is still applied per interface, not per zone. If you configure the same security policy for all interfaces within the zone, then you can successfully implement ECMP and load balancing for that traffic. For more information about required parallel interface configuration, see Prerequisites for Traffic Zones, on page 7.

Supported Services for Traffic Zones

The following services are supported with zones:

- Access Rules
- NAT
- Service Rules, except for QoS traffic policing.
- Routing

You can also configure to- and from-the-box services listed in To- and From-the-Box Traffic, on page 7, although full zoned support is not available.

Do not configure other services (such as VPN or Botnet Traffic Filter) for interfaces in a traffic zone; they may not function or scale as expected.

Note For detailed information about how to configure the security policy, see Prerequisites for Traffic Zones, on page 7.

Security Levels

The first interface that you add to a zone determines the security level of the zone. All additional interfaces must have the same security level. To change the security level for interfaces in a zone, you must remove all but one interface, and then change the security levels, and re-add the interfaces.

Primary and Current Interface for the Flow

Each connection flow is built based on the initial ingress and egress interfaces. These interfaces are the *primary* interfaces.

If a new egress interface is used because of route changes or asymmetric routing, then the new interfaces are the *current* interfaces.

Joining or Leaving a Zone

When you assign an interface to a zone, any connections on that interface are deleted. The connections must be reestablished.

If you remove an interface from a zone, any connections that have the interface as the primary interface are deleted. The connections must be reestablished. If the interface is the current interface, the ASA moves the connections back to the primary interface. The zone route table is also refreshed.

Intra-Zone Traffic

To allow traffic to *enter* one interface and *exit* another in the same zone, enable **Configuration > Device Setup > Interface Settings > Interfaces > Enable traffic between two or more hosts connected to the same interface**, which allows traffic to enter and exit the same interface, as well as **Configuration > Device**

Setup > Interface Settings > Interfaces > Enable traffic between two or more interfaces which are configured with same security level, which allows traffic between same-security interfaces. Otherwise, a flow cannot be routed between two interfaces in the same zone.

To- and From-the-Box Traffic

- You cannot add management-only or management-access interfaces to a zone.
- For management traffic on regular interfaces in a zone, only asymmetric routing on existing flows is supported; there is no ECMP support.
- You can configure a management service on only one zone interface, but to take advantage of asymmetric routing support, you need to configure it on all interfaces. Even when the configurations are parallel on all interfaces, ECMP is not supported.
- The ASA supports the following to- and from-the-box services in a zone:
  - Telnet
  - SSH
  - HTTPS
  - SNMP
  - Syslog

Overlapping IP Addresses Within a Zone

For non-zoned interfaces, the ASA supports overlapping IP address networks on interfaces so long as you configure NAT properly. However, overlapping networks are not supported on interfaces in the same zone.

Prerequisites for Traffic Zones

- Configure all interface parameters including the name, IP address, and security level. Note that the security level must match for all interfaces in the zone. You should plan to group together like interfaces in terms of bandwidth and other Layer 2 properties.
- Configure the following services to match on all zone interfaces:
  - Access Rules—Apply the same access rule to all zone member interfaces, or use a global access rule.

    For example:

    ```
    access-list ZONE1 extended permit tcp any host WEBSERVER1 eq 80
    access-group ZONE1 in interface outside1
    access-group ZONE1 in interface outside2
    access-group ZONE1 in interface outside3
    ```

  - NAT—Configure the same NAT policy on all member interfaces of the zone or use a global NAT rule (in other words, use “any” to represent the zone interfaces in the NAT rule).

Interface PAT is not supported.

For example:

```
object network WEBSERVER1
  host 10.9.9.9 255.255.255.255
  nat (inside,any) static 209.165.201.9
```



Note When you use interface-specific NAT and PAT pools, the ASA cannot switch connections over in case of the original interface failure.

If you use interface-specific PAT pools, multiple connections from the same host might load-balance to different interfaces and use different mapped IP addresses. Internet services that use multiple concurrent connections may not work correctly in this case.

- Service Rules—Use the global service policy, or assign the same policy to each interface in a zone.

  QoS traffic policing is not supported.

  For example:

```
service-policy outside_policy interface outside1
service-policy outside_policy interface outside2
service-policy outside_policy interface outside3
```



Note For VoIP inspections, zone load balancing can cause increased out-of-order packets. This situation can occur because later packets might reach the ASA before earlier packets that take a different path. Symptoms of out-of-order packets include:

- Higher memory utilization at intermediate nodes (firewall and IDS) and the receiving end nodes if queuing is used.
- Poor video or voice quality.

To mitigate these effects, we recommend that you use IP addresses only for load distribution for VoIP traffic.

- Configure routing with ECMP zone capabilities in mind.

Guidelines for Traffic Zones

Firewall Mode

Supported in routed firewall mode only. Does not support transparent firewall mode or bridge group interfaces in routed mode.

Failover

- You cannot add the failover or state link to a zone.
- In Active/Active failover mode, you can assign an interface in each context to an asymmetrical routing (ASR) group. This service allows traffic returning on a similar interface on the peer unit to be restored to the original unit. You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group. See Configure Support for Asymmetrically Routed Packets (Active/Active Mode) for more information about ASR groups.
- Only the primary interfaces for each connection are replicated to the standby unit; current interfaces are not replicated. If the standby unit becomes active, it will assign a new current interface if necessary.

Clustering

- You cannot add the cluster control link to a zone.

Model Guidelines

You cannot add Firepower 1010 switch ports and VLAN interfaces to a zone.

Additional Guidelines

- You can create a maximum of 256 zones.
- You can add the following types of interfaces to a zone:
  - Physical
  - VLAN
  - EtherChannel
- You cannot add the following types of interfaces:
  - Management-only
  - Management-access
  - Failover or state link
  - Cluster control link
  - Member interfaces in an EtherChannel
  - VNI; also, if a regular data interface is marked as nve-only, it cannot be a member of a zone.

  - BVI, or bridge group member interfaces.

- An interface can be a member of only one zone.
- You can include up to 8 interfaces per zone.
- For ECMP, you can add up to 8 equal cost routes per zone, across all zone interfaces. You can also configure multiple routes on a single interface as part of the 8 route limit.
- When you add an interface to a zone, all static routes for those interfaces are removed.
- You cannot enable DHCP Relay on an interface in a traffic zone.
- The ASA does not support fragmented packet reassembly for fragments that are load-balanced to separate interfaces; those fragments will be dropped.
- PIM/IGMP Multicast routing is not supported on interfaces in a zone.

Configure a Traffic Zone

Configure a named zone, and assign interfaces to the zone.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Zones**, and click **Add**.

You can alternately assign an interface to a zone from the **Configuration > Device Setup > Interface Settings > Interfaces > Add Interface** dialog box.

Step 2 Name the zone with a name up to 48 characters in length.

Step 3 Add one or more interfaces to the **Member** area. Ensure all interfaces have the same security level.

Step 4 Click **Apply**.

Monitoring Traffic Zones

This section describes how to monitor traffic zones.

Zone Information

- **show zone** [*name*]

  Shows zone ID, context, security level, and members.

  See the following output for the **show zone** command:

```
ciscoasa# show zone outside-zone

Zone: zone-outside id: 2
```

```
Security-level: 0
Context: test-ctx
 Zone Member(s) : 2
  outside1       GigabitEthernet0/0
  outside2       GigabitEthernet0/1
```

- **show nameif zone**

  Shows the interface names and zone names.

  See the following output for the **show nameif zone** command:

```
ciscoasa# show nameif zone
Interface                Name                    zone-name      Security
GigabitEthernet0/0       inside-1                inside-zone     100
GigabitEthernet0/1.21    inside                  inside-zone     100
GigabitEthernet0/1.31    4                                         0
GigabitEthernet0/2       outside                 outside-zone   0
Management0/0            lan                                       0
```

Zone Connections

- **show conn** [**long** | **detail**] [**zone** *zone_name* [**zone** *zone_name*] [...]]

  The **show conn zone** command displays connections for a zone. The **long** and **detail** keywords show the primary interface on which the connection was built and the one in the brackets is the current interface used to forward the traffic or the interface the last packet came from. Thus, the current interface in case of a connection coming from multiple interfaces can show different interfaces at different times depending on when the show conn command was issued.

  See the following output for the **show conn long zone** command:

```
ciscoasa# show conn long zone zone-inside zone zone-outside

TCP outside-zone:outside1(outside2): 10.122.122.1:1080
inside-zone:inside1(inside2): 10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

- **show asp table zone**

  Shows the accelerated security path tables for debugging purposes.

- **show local-host** [**zone** *zone_name* [**zone** *zone_name*] [...]]

  Shows the network states of local hosts within a zone.

  See the following output for the **show local-host zone** command. The primary interface is listed first, and the current interface is in parentheses.

```
ciscoasa# show local-host zone outside-zone

Zone:outside-zone: 4 active, 5 maximum active, 0 denied
local host: <10.122.122.1>,
    TCP flow count/limit = 3/unlimited
    TCP embryonic count to host = 0
    TCP intercept watermark = unlimited
    UDP flow count/limit = 0/unlimited
```

```
  Conn:
 TCP outside-zone:outside1(outside2): 10.122.122.1:1080
inside-zone:inside1(inside2): 10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

Zone Routing

- **show route zone**

  Shows the routes for zone interfaces.

  See the following output for the **show route zone** command:

```
ciscoasa# show route zone

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

S    192.168.105.1 255.255.255.255 [1/0] via 172.16.1.1, outside-zone:outside1
C    192.168.212.0 255.255.255.0 is directly connected, lan-zone:inside,
C    172.16.1.0 255.255.255.0 is directly connected, wan-zone:outside2
S    10.5.5.0 255.255.255.0 [1/0] via 172.16.1.1, wan-zone:outside2
O    10.2.2.1 255.255.255.255 [110/11] via 192.168.212.3, 2:09:24, lan-zone:inside
O    10.1.1.1 255.255.255.255 [110/11] via 192.168.212.2, 2:09:24, lan-zone:inside
```

- **show asp table routing**

  Shows the accelerated security path tables for debugging purposes, and shows the zone associated with each route.

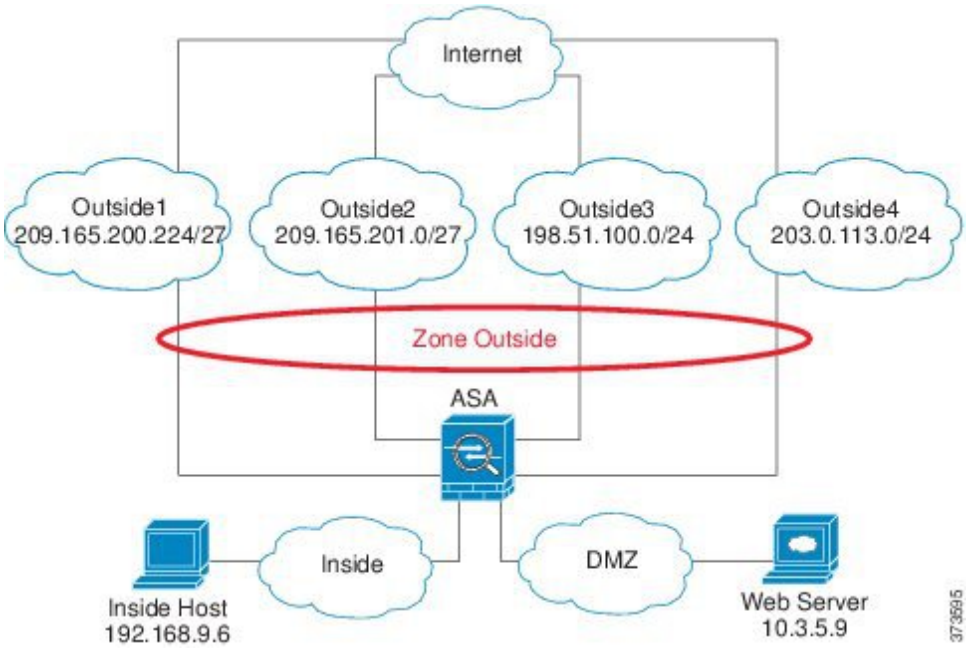
  See the following output for the **show asp table routing** command:

```
ciscoasa# show asp table routing
route table timestamp: 60
in   255.255.255.255 255.255.255.255 identity
in   10.1.0.1       255.255.255.255 identity
in   10.2.0.1       255.255.255.255 identity
in   10.6.6.4         255.255.255.255 identity
in   10.4.4.4         255.255.255.255 via 10.4.0.10 (unresolved, timestamp: 49)
in   172.0.0.67       255.255.255.255 identity
in   172.0.0.0       255.255.255.0   wan-zone:outside2
in   10.85.43.0      255.255.255.0   via 10.4.0.3 (unresolved, timestamp: 50)
in   10.85.45.0      255.255.255.0   via 10.4.0.20 (unresolved, timestamp: 51)
in   192.168.0.0        255.255.255.0   mgmt
in   192.168.1.0       255.255.0.0     lan-zone:inside
out  255.255.255.255 255.255.255.255 mgmt
out  172.0.0.67       255.255.255.255 mgmt
out  172.0.0.0        255.255.255.0   mgmt
out  10.4.0.0       240.0.0.0       mgmt
out  255.255.255.255 255.255.255.255 lan-zone:inside
out  10.1.0.1       255.255.255.255 lan-zone:inside
out  10.2.0.0       255.255.0.0     lan-zone:inside
```

```
out  10.4.0.0       240.0.0.0       lan-zone:inside
```

Example for Traffic Zones

The following example assigns 4 VLAN interfaces to the outside zone, and configures 4 equal cost default routes. PAT is configured for the inside interface, and a web server is available on a DMZ interface using static NAT.




```
interface gigabitethernet0/0
  no shutdown
  description outside switch 1
interface gigabitethernet0/1
  no shutdown
  description outside switch 2

interface gigabitethernet0/2
  no shutdown
  description inside switch

zone outside

interface gigabitethernet0/0.101
  vlan 101
  nameif outside1
  security-level 0
  ip address 209.165.200.225 255.255.255.224
  zone-member outside
  no shutdown

interface gigabitethernet0/0.102
  vlan 102
  nameif outside2
```

```
  security-level 0
  ip address 209.165.201.1 255.255.255.224
  zone-member outside
  no shutdown

interface gigabitethernet0/1.201
  vlan 201
  nameif outside3
  security-level 0
  ip address 198.51.100.1 255.255.255.0
  zone-member outside
  no shutdown

interface gigabitethernet0/1.202
  vlan 202
  nameif outside4
  security-level 0
  ip address 203.0.113.1 255.255.255.0
  zone-member outside
  no shutdown

interface gigabitethernet0/2.301
  vlan 301
  nameif inside
  security-level 100
  ip address 192.168.9.1 255.255.255.0
  no shutdown

interface gigabitethernet0/2.302
  vlan 302
  nameif dmz
  security-level 50
  ip address 10.3.5.1 255.255.255.0
  no shutdown

# Static NAT for DMZ web server on any destination interface
object network WEBSERVER
  host 10.3.5.9 255.255.255.255
  nat (dmz,any) static 209.165.202.129 dns

# Dynamic PAT for inside network on any destination interface
object network INSIDE
  subnet 192.168.9.0 255.255.255.0
  nat (inside,any) dynamic 209.165.202.130

# Global access rule for DMZ web server
access-list WEB-SERVER extended permit tcp any host WEBSERVER eq 80
access-group WEB-SERVER global

# 4 equal cost default routes for outside interfaces
route outside1 0 0 209.165.200.230
route outside2 0 0 209.165.201.10
route outside3 0 0 198.51.100.99
route outside4 0 0 203.0.113.87
# Static routes for NAT addresses - see redistribute static command
route dmz 209.165.202.129 255.255.255.255 10.3.5.99
route inside 209.165.202.130 255.255.255.255 192.168.9.99

# The global service policy
class-map inspection_default
  match default-inspection-traffic
policy-map type inspect dns preset_dns_map
  parameters
    message-length maximum client auto
```

```
    message-length maximum 512
    dns-guard
    protocol-enforcement
    nat-rewrite
policy-map global_policy
  class inspection_default
    inspect dns preset_dns_map
    inspect ftp
    inspect h323 h225 _default_h323_map
    inspect h323 ras _default_h323_map
    inspect ip-options _default_ip_options_map
    inspect netbios
    inspect rsh
    inspect rtsp
    inspect skinny
    inspect esmtp _default_esmtp_map
    inspect sqlnet
    inspect sunrpc
    inspect tftp
    inspect sip
    inspect xdmcp
service-policy global_policy global
```

History for Traffic Zones

Feature Name	Platform Releases	Description
Traffic Zones	9.3(2)	You can group interfaces together into a traffic zone to accomplish traffic load balancing (using Equal Cost Multi-Path (ECMP) routing), route redundancy, and asymmetric routing across multiple interfaces. **Note** You cannot apply a security policy to a named zone; the security policy is interface-based. When interfaces in a zone are configured with the same access rule, NAT, and service policy, then load-balancing and asymmetric routing operate correctly. We introduced or modified the following screens: **Configuration > Device Setup > Interface Parameters > Zones** **Configuration > Device Setup > Interface Parameters > Interfaces**.
clear local-host command	9.14(1)	The **clear local-host** command and all of its attributes and keywords were deprecated. They will be removed in a future release.