

Contents

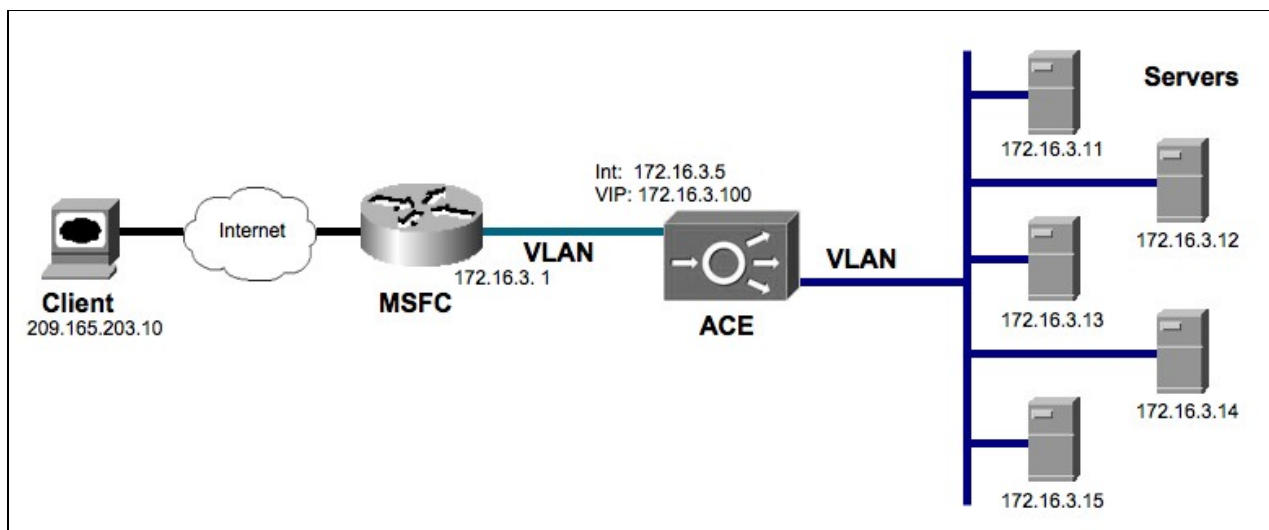
- [1 Goal](#)
- [2 Design](#)
- [3 Configuration](#)
- [4 Related show Commands](#)
- [5 Comments](#)
- [6 show running-config](#)
- [7 Related Information](#)

Goal

Configure basic load balancing (Layer 3) where client traffic enters on one VLAN and is bridged to servers residing on a second VLAN, while the VIP and server IPs reside in the same network.

Design

Clients will send application requests through the multilayer switch feature card (MSFC), which routes them to a virtual IP address (VIP) within the Cisco® Application Control Engine (ACE). The VIP used in this example resides in an ACE context, which is configured with a client VLAN and a server VLAN (Figure 1), and a Bridge Group Virtual Interface (BVI) to tie the VLANs together in a bridge-group. Client requests will arrive at the VIP, and the ACE will pick the appropriate server and then use the destination Network Address Translation (NAT) to send the client request to the server. The server will respond using the interface VLAN of the MSFC as its default gateway to the client. The ACE will then change the source IP to be the VIP and bridge the traffic to the MSFC. The MSFC will then forward the response to the client.



Configuration

The Cisco ACE needs to be configured via access control lists (ACLs) to allow traffic into the ACE data plane. After the ACL checks are made, a service policy, which is applied to the interface, is used to classify traffic destined for the VIP. The VIP is associated with a load-balancing action within the multimatch policy. The load-balancing action tells the ACE how to handle traffic that has been directed to a VIP. In this example, all traffic is sent to a server farm, where it is distributed in round-robin fashion to one of five real


Basic_Load_Balancing_Using_Bridged_Mode_on_the_Cisco_Application_Control_Engine_Configuration_Example

servers. The ACE configuration occurs in layers, such that it builds from the real IPs to applying the VIP on an interface. Due to this layered structure, it is optimal to create the configuration by working backward from the way the flow is processed. Thus, to enable server load balancing you need to do the following:

- Enable ACLs to allow data traffic through the ACE device, as it is denied by default.
- Configure the IPs of the servers (define rservers).
- Group the real servers (create a server farm).
- Define the virtual IP address (VIP).
- Define how traffic is to be handled as it is received (create a policy map for load-balancing).
- Associate a VIP to a handling action (create a multimatch policy map [a service policy])
- Create client- and server-facing interfaces and a BVI to join them.
- Apply the VIP and ACL permitting client connections to the interface (apply access group and service policy to interface).

To begin the configuration, create an access list for permitting client connections.

```
ACE-1/bridged(config)# access-list everyone extended permit ip any any
ACE-1/bridged(config)# access-list everyone extended permit icmp any any
```

 **Note:** Although this example shows a ?permit any any,? it is recommended that ACLs be used to permit only the traffic you want allow through the Cisco ACE. In the past, server load-balancing (SLB) devices have used the VIP and port alone to protect servers. Within the ACE, ACLs are processed first, and thus dropping traffic using an ACL requires fewer resources than dropping it once it passes the ACLs and reaches the VIP.

The Cisco ACE needs to know the IP address of the servers available to handle client connections. The rserver command is used to define the IP address of the service. In addition, each rserver must be place in service for it to be used. The benefit of this design is that no matter how many applications or services an rserver hosts, the entire real server can be completely removed from the load-balancing rotation by issuing a single ?no inservice? or ?no inservice-standby? command at the rserver level. This is very beneficial for users needing to upgrade or patch an rserver, because they no longer have to go to each application and remove each instance of the rserver.

```
ACE-1/bridged(config)# rserver lnx1
ACE-1/bridged(config-rserver-host)# ip add 172.16.3.11
ACE-1/bridged(config-rserver-host)# inservice
ACE-1/bridged(config-rserver-host)# rserver lnx2
ACE-1/bridged(config-rserver-host)# ip add 172.16.3.12
ACE-1/bridged(config-rserver-host)# inservice
ACE-1/bridged(config-rserver-host)# rserver lnx3
ACE-1/bridged(config-rserver-host)# ip add 172.16.3.13
ACE-1/bridged(config-rserver-host)# inservice
ACE-1/bridged(config-rserver-host)# rserver lnx4
ACE-1/bridged(config-rserver-host)# ip add 172.16.3.14
ACE-1/bridged(config-rserver-host)# inservice
ACE-1/bridged(config-rserver-host)# rserver lnx5
ACE-1/bridged(config-rserver-host)# ip add 172.16.3.15
ACE-1/bridged(config-rserver-host)# inservice
```

Now group the rservers to be used to handle client connections into a server farm. Again, the rserver must be placed in service. This allows a single instance of an rserver to be manually removed from rotation.

```
ACE-1/bridged(config-cmap)# serverfarm web
ACE-1/bridged(config-sfarm-host)# rserver lnx1
ACE-1/bridged(config-sfarm-host-rs)# inservice
ACE-1/bridged(config-sfarm-host-rs)# rserver lnx2
ACE-1/bridged(config-sfarm-host-rs)# inservice
ACE-1/bridged(config-sfarm-host-rs)# rserver lnx3
```

Basic_Load_Balancing_Using_Bridged_Mode_on_the_Cisco_Application_Control_Engine_Configuration_Example

```
ACE-1/bridged(config-sfarm-host-rs) # inservice
ACE-1/bridged(config-sfarm-host-rs) # rserver lnx4
ACE-1/bridged(config-sfarm-host-rs) # inservice
ACE-1/bridged(config-sfarm-host-rs) # rserver lnx5
ACE-1/bridged(config-sfarm-host-rs) # inservice
```

Use a class map to define the VIP to which clients will send their requests. In this example, the VIP is considered L3 (Layer 3) because there is a match on any port. If the VIP were to match only HTTP traffic, the match would be bound to port 80 and considered an L4 (Layer 4) VIP. (For example, ?match virtual-address 172.16.1.100 tcp eq 80?).

```
ACE-1/bridged(config) # class-map slb-vip
ACE-1/bridged(config-cmap) # match virtual-address 172.16.1.100 any
```

Next define the action to take when a new client request arrives. In this case, all traffic will be sent to the ?web? serverfarm. This type of load balancing is considered L4 since only class-default is used.

```
ACE-1/bridged(config) # policy-map type loadbalance http first-match slb
ACE-1/bridged(config-pmap-lb) # class class-default
ACE-1/bridged(config-pmap-lb-c) # serverfarm web
```

Since the VIPs and load-balancing actions are defined independently, they must be associated so that the Cisco ACE knows how to handle traffic destined for a VIP. The association is made using a multimatch policy map. Keep in mind that multimatch policy maps are applied to interfaces as service policies.

```
ACE-1/bridged(config) # policy-map multi-match client-vips
ACE-1/bridged(config-pmap) # class slb-vip
ACE-1/bridged(config-pmap-c) # loadbalance policy slb
ACE-1/bridged(config-pmap-c) # loadbalance vip inservice
```

At this point the interface VLANs and BVI can be created to interconnect the Cisco ACE to the client side of the network and to the servers.


```
ACE-1/bridged(config) # interface vlan 30
ACE-1/bridged(config-if) # description ?Client Side?
ACE-1/bridged(config-if) # bridge-group 3
ACE-1/bridged(config-if) # no shutdown

ACE-1/bridged(config-if) # interface vlan 31
ACE-1/bridged(config-if) # description ?Server Side?
ACE-1/bridged(config-if) # bridge-group 3
ACE-1/bridged(config-if) # no shutdown

ACE-1/bridged(config-if) # interface bvi 3
ACE-1/bridged(config-if) # description ?client ? server bridge group?
ACE-1/bridged(config-if) # ip address 172.16.3.5 255.255.255.0
```

The last step is to apply the ACL and service policy (?policy-map multi-match?) to the client-side interface. Both the access group and service policy are applied on the input side of the interface.

```
ACE-1/bridged(config) # interface vlan 30
ACE-1/bridged(config-if) # access-group input everyone
ACE-1/bridged(config-if) # service-policy input client-vips
```

 **Note:** There is no need to add an access group to the server side, as the Cisco ACE automatically creates pinholes to allow server response traffic to pass back to the client.

Related show Commands

This section provides information you can use to confirm your configuration is working properly.

Certain show commands are supported by the [Output Interpreter Tool \(registered customers only\)](#), which allows you to view an analysis of show command output.

```
ACE-1/bridged #show arp
ACE-1/bridged #show acl
ACE-1/bridged #show service-policy client-vips
ACE-1/bridged #show serverfarm
ACE-1/bridged #show rserver
ACE-1/bridged #show stats
```

Comments

Once you've completed the configuration, verify that the Cisco ACE has an Address Resolution Protocol (ARP) response for each rserver and the default Bridge to the client. Check the ACL hits to ensure that client connections are being accepted. Check the service policy output to see the client connection hits, and verify that the server is responding with response packets. The ?show? command for serverfarm and rserver can be used to display the exact rserver handling the connection and the amount of work the entire server farm has handled. The ?show stats? command provides a higher level of monitoring of ACE load balancing, inspection, probes, and other important metrics.

show running-config

```
ACE-1/bridged# show run
Generating configuration....

access-list everyone line 8 extended permit ip any any
access-list everyone line 16 extended permit icmp any any

rserver host lnx1
  ip address 172.16.3.11
  inservice
rserver host lnx2
  ip address 172.16.3.12
  inservice
rserver host lnx3
  ip address 172.16.3.13
  inservice
rserver host lnx4
  ip address 172.16.3.14
  inservice
rserver host lnx5
  ip address 172.16.3.15
  inservice

serverfarm host web
  rserver lnx1
    inservice
  rserver lnx2
    inservice
  rserver lnx3
    inservice
  rserver lnx4
    inservice
  rserver lnx5
```

Basic_Load_Balancing_Using_Bridged_Mode_on_the_Cisco_Application_Control_Engine_Configuration_Example

```
inservice

class-map match-all slb-vip
  2 match virtual-address 172.16.3.100 any

policy-map type management first-match remote-access
  class class-default
    permit

policy-map type loadbalance http first-match slb
  class class-default
    serverfarm web

policy-map multi-match client-vips
  class slb-vip
    loadbalance vip inservice
    loadbalance policy slb

interface vlan 30
  description "Client Side"
  bridge-group 3
  access-group input everyone
  service-policy input client-vips
  no shutdown

interface vlan 31
  description "Server Side"
  bridge-group 3
  service-policy input remote-access
  no shutdown

interface bvi 3
  ip address 172.16.3.5 255.255.255.0
  description "client - server bridge group"
  no shutdown

ip route 0.0.0.0 0.0.0.0 172.16.3.1
```

Related Information

[Technical Support & Documentation - Cisco Systems](#)