Contents

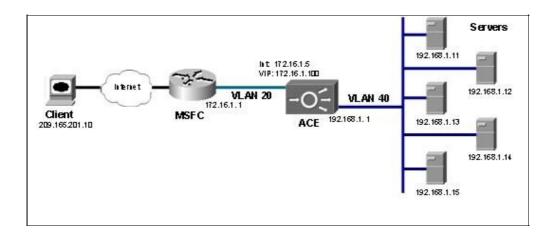
- 1 Goal
- 2 Design
- 3 Configuration
- <u>4 Related show</u> Commands
- 5 Comments
- <u>6 show</u> running-config
- <u>7 Related</u> Information

Goal

Configure basic load balancing (Layer 3) where client traffic enters on one VLAN and Network Address Translation (NAT) is used when sending the client request out the same VLAN to the servers. The servers will respond to the Cisco® Application Control Engine (ACE), where the server?s IP is replaced with the VIP and the response message is sent to the client via the multilayer switch feature card (MSFC). The serverfarm for each request will be chosen based on the URL being requested.

Design

Clients will send application requests through the MFSC, which routes them to a virtual IP address (VIP) within ACE. The VIP used in this example resides in an ACE context, which is configured with a single VLAN to handle client and server communication (Figure 1.). Client requests will arrive at the VIP and the Cisco ACE will pick the appropriate server to handle the request. ACE will rewrite the destination IP to that of the rserver and rewrite the source IP with one from a nat-pool. Once the client request is fully NAT?d it will be sent to the server over the same VLAN which it was originally received. The server will respond to the Cisco ACE, based on the source IP of the request. The Cisco ACE will receive the response, change the source IP to be the VIP, and send it to the MSFC. The MSFC will forward the response to the client.



Configuration

The Cisco ACE needs to be configured via access control lists (ACLs) to allow traffic into the Cisco 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 Cisco ACE how to handle traffic that has been

Contents 1

oad_Balancing_Using_One_Arm_Mode_with_Source_NAT_on_the_Cisco_Application_Control_Engine_Configuration_E

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 servers. The Cisco 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 Cisco 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.
- 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/onearm(config) # access-list everyone extended permit ip any any ACE-1/onearm(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 Cisco 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 reserver command is used to define the IP address of the service. In addition, each reserver 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 reserver 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 reserver level. This is very beneficial for users needing to upgrade or patch an reserver, because they no longer have to go to each application and remove each instance of the reserver.

```
ACE-1/onearm(config) # rserver lnx1
ACE-1/onearm(config-rserver-host) # ip add 192.168.5.11
ACE-1/onearm(config-rserver-host) # inservice
ACE-1/onearm(config-rserver-host) # rserver lnx2
ACE-1/onearm(config-rserver-host) # ip add 192.168.5.12
ACE-1/onearm(config-rserver-host) # inservice
ACE-1/onearm(config-rserver-host) # rserver lnx3
ACE-1/onearm(config-rserver-host) # ip add 192.168.5.13
ACE-1/onearm(config-rserver-host) # inservice
ACE-1/onearm(config-rserver-host) # rserver lnx4
ACE-1/onearm(config-rserver-host) # ip add 192.168.5.14
ACE-1/onearm(config-rserver-host) # inservice
ACE-1/onearm(config-rserver-host) # rserver lnx5
ACE-1/onearm(config-rserver-host) # ip add 192.168.5.15
ACE-1/onearm(config-rserver-host) # ip add 192.168.5.15
```

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

```
ACE-1/routed(config) # serverfarm webfarm
ACE-1/routed(config-sfarm-host) # rserver lnx1
ACE-1/routed(config-sfarm-host-rs) # inservice
ACE-1/routed(config-sfarm-host-rs) # rserver lnx2
```

Configuration 2

_oad_Balancing_Using_One_Arm_Mode_with_Source_NAT_on_the_Cisco_Application_Control_Engine_Configuration_E

```
ACE-1/routed(config-sfarm-host-rs) # inservice

ACE-1/routed(config) # serverfarm imagefarm

ACE-1/routed(config-sfarm-host) # rserver lnx3

ACE-1/routed(config-sfarm-host-rs) # inservice

ACE-1/routed(config-sfarm-host-rs) # rserver lnx4

ACE-1/routed(config-sfarm-host-rs) # inservice

ACE-1/routed(config-sfarm-host-rs) # exit
```

Use a class-map to define the VIP where clients will be sending their requests.

```
ACE-1/onearm(config) # class-map slb-vip
ACE-1/onearm(config-cmap) # match virtual-address 172.16.5.101 any
```

Next define an http class to match URLs requesting images. This is done using a simple http url match on the /image/ directory name. The wildcard allows for any image filename to be matched.

```
ACE-1/onearm(config) # class-map type http loadbalance match-all images ACE-1/onearm(config-cmap-http-lb) # match http url /images/.*
```

Next define the action to take when a new client request arrives. In this case we have two possible actions, based on the http matching class defined above requests for images will be sent to the imagefarm, while all other requests will be sent to the web servers in the webfarm.

```
ACE-1/onearm(config) # policy-map type loadbalance http first-match slb-logic ACE-1/onearm(config-pmap-lb) # class images
ACE-1/onearm(config-pmap-lb-c) # serverfarm imagefarm
ACE-1/onearm(config-pmap-lb-c) # class class-default
ACE-1/onearm(config-pmap-lb-c) # serverfarm webfarm
```

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. ?nat dynamic? is configured to make the Cisco ACE source NAT all client requests. The nat-pool will be defined in a later step.

```
ACE-1/onearm(config) # policy-map multi-match client-vips ACE-1/onearm(config-pmap) # class slb-vip ACE-1/onearm(config-pmap-c) # loadbalance policy slb ACE-1/onearm(config-pmap-c) # loadbalance vip inservice ACE-1/onearm(config-pmap-c) # nat dynamic 5 vlan 50
```

At this point the interface VLAN can be created to interconnect the Cisco ACE to the network.

```
ACE-1/onearm(config)# interface vlan 50
ACE-1/onearm(config-if)# description ?Client-Sever VLAN?
ACE-1/onearm(config-if)# ip address 172.16.5.5 255.255.255.0
ACE-1/onearm(config-if)# no shutdown
```

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. The nat-pool is also created, for use in the multi-match policy.

```
ACE-1/onearm(config) # interface vlan 50
ACE-1/onearm(config-if) # access-group input everyone
ACE-1/onearm(config-if) # service-policy input client-vips
ACE-1/onearm(config-if) # nat-pool 5 172.16.5.200 172.16.5.209 netmask 255.255.255.0 pat
```

Configuration 3

oad_Balancing_Using_One_Arm_Mode_with_Source_NAT_on_the_Cisco_Application_Control_Engine_Configuration_E

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.

When a client connects to the VIP they download the index.html page from the webfarm servers and all the embedded images from the imagefarm servers. This can be verified using the detailed output of the show service-policy command:

```
ACE-1/onearm# show service-policy client-vips detail
Status : ACTIVE
Description: -
Interface: vlan 50
  service-policy: client-vips
   class: slb-vip
     nat:
       nat dynamic 5 vlan 50
        curr conns : 0 dropped conns : 2
                                     , hit count : 83
                         : 2
       client pkt count : 524 , client byte count: 45320 server pkt count : 558 , server byte count: 146197
        bandwidth-rate-limit: 0 , drop-count: 0
     VIP Address: Protocol: Port:
     172.16.5.101
                     tcp eq 80
      loadbalance:
        L7 loadbalance policy: slb-logic
        VIP Route Metric : 77
        VIP Route Advertise : DISABLED
        VIP ICMP Reply : DISABLED
        VIP State: INSERVICE
        dropped conns : 0
client pkt count : 0 , client byte count: 0
server pkt count : 0 , server byte count: 0
conn-rate-limit : 0 , drop-count : 0
bandwidth-rate-limit : 0 , drop-count : 0
        L7 Loadbalance policy : slb-logic
          class/match : images
            LB action :
              primary serverfarm: imagefarm
                    state: UP
                backup serverfarm : -
            hit count : 20
            dropped conns
          class/match : class-default
            LB action :
               primary serverfarm: webfarm
                    state: UP
                backup serverfarm : -
            hit count : 5
            dropped conns
                              : 0
```

Related show Commands

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

Certain show commands are supported by the <u>Output Interpreter Tool (registered customers only)</u>, which allows you to view an analysis of show command output.

```
ACE-1/onearm #show arp
ACE-1/onearm #show acl
ACE-1/onearm #show service-policy client-vips
ACE-1/routed# show service-policy client-vips detail
ACE-1/onearm #show serverfarm
ACE-1/onearm #show rserver
ACE-1/onearm #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 route 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/onearm# sho 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 192.168.5.11
 inservice
rserver host lnx2
 ip address 192.168.5.12
 inservice
rserver host lnx3
 ip address 192.168.5.13
 inservice
rserver host lnx4
 ip address 192.168.5.14
 inservice
rserver host lnx5
  ip address 192.168.5.15
  inservice
serverfarm host web
 rserver lnx1
   inservice
 rserver lnx2
   inservice
  rserver lnx3
   inservice
 rserver lnx4
   inservice
 rserver lnx5
   inservice
class-map match-all slb-vip
  2 match virtual-address 172.16.5.100 any
policy-map type management first-match remote-access
 class class-default
   permit
```

Comments 5

```
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
   nat dynamic 5 vlan 50
interface vlan 50
  description "Client-Server VLAN"
  ip address 172.16.5.5 255.255.25.0
  access-group input everyone
  service-policy input client-vips
  service-policy input remote-access
  nat-pool 5 172.16.5.200 172.16.5.209 netmask 255.255.255.0 pat
  no shutdown
ip route 0.0.0.0 0.0.0.0 172.16.5.1
```

Related Information

Technical Support & Documentation - Cisco Systems

show running-config 6