



ECMP Load Balancing

Equal-cost multi-path routing (ECMP) is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple "best paths" which tie for top place in routing metric calculations. Multipath routing can be used in conjunction with most routing protocols, since it is a per-hop decision that is limited to a single router. It potentially offers substantial increases in bandwidth by load-balancing traffic over multiple paths.

Various routing protocols, including Open Shortest Path First (OSPF), Intermediate System to Intermediate System (ISIS), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP) allow ECMP routing.

Load balancing between ECMP paths is performed on IOS-XE based CEF object called loadbalance.


ECMP Per-Flow Load Balancing

Load balancing is a forwarding mechanism that distributes traffic over multiple links based on certain parameters. ECMP Per-Flow Load Balancing distributes packets across multiple links based on Layer 3 routing information. If the router discovers multiple paths to a destination, the routing table is updated with multiple entries for that destination. Per-flow load balancing allows the router to use multiple paths to achieve load sharing across multiple source-destination host pairs. Packets for a given source-destination host pair are guaranteed to take the same path, even if multiple paths are available. Traffic streams destined for different pairs tend to take different paths.

Benefits of Per-Flow Load Balancing

- Incoming data traffic is evenly distributed over multiple equal-cost connections.
- Incoming data traffic is evenly distributed over multiple equal-cost connections member links within a bundle interface.

Figure 1: ECMP Load Balancing with MPLS Enabled



- [Restrictions for ECMP Load Balancing, on page 2](#)
- [Configuring ECMP Load Balancing, on page 2](#)

Restrictions for ECMP Load Balancing

- Both 4 ECMP and 8 ECMP paths are supported.
- Load balancing is supported on global IPv4 and IPv6 traffic. For global IPv4 and IPv6 traffic, the traffic distribution can be equal among the available 8 links.
- Per packet load balancing is not supported.
- Label load balancing is supported.
- BGP multi-path is *not* supported with ECMP.
- BGP multi-path with PIC Edge is *not* supported
- When BGP PIC is configured, the L3VPN prefixes scale reduces by 1/4th of the supported value (Supported scale value/4), for better convergence value at the PIC core. For example, for RSP1A the supported L3VPN scale is 2000, if 4 ECMP path with PIC is enabled, then the maximum supported scale value is reduced to 5000 (20000/4).

Configuring ECMP Load Balancing

Perform the following steps to configure ECMP load balancing.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **platform loadbalance max-paths 8**

4. exit

DETAILED STEPS

| | Command or Action | Purpose |
|--------|--|--|
| Step 1 | enable Example: Device> enable | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | configure terminal Example: Device# configure terminal | Enters global configuration mode. |
| Step 3 | platform loadbalance max-paths 8 Example: Device(config)# platform loadbalance max-paths 8 | Configures the loadbalance maximum paths. Select Yes to save the configuration and reload the router. Note ISIS by default supports only 4 paths. To increase ISIS max-paths, use the command config-maximum-paths 8 under router ISIS. IGP by default supports only 4 paths. To increase IGP max-paths, use the command config-maximum-paths 8 under respective IGP (OSPF and ISIS) process. |
| Step 4 | exit Example: Device (config) #exit | Returns to privileged EXEC mode. |

Configuration Examples for ECMP Load Balancing

This section shows sample configurations for ECMP load balancing.

Example: Configuring ECMP Load balancing

The following is a sample configuration for ECMP load balancing.

Verifying ECMP Load Balancing

Use the following commands to verify ECMP load balancing.

```
Building configuration...

Current configuration : 10710 bytes
!
! Last configuration change at 00:29:01 IST Sat Jan 17 2015
!
version 15.5
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
platform loadbalance max-paths 8
```

```

no platform punt-keepalive disable-kernel-core
platform bfd-debug-trace 1
platform tcam-parity-error enable
platform tcam-threshold alarm-frequency 1
platform shell
!
hostname RM-PE1
!
boot-start-marker
boot-end-marker
!
!
vrf definition Mgmt-intf
!
address-family ipv4
exit-address-family
!
address-family ipv6
exit-address-family
!
!
no aaa new-model
clock timezone IST 5 30
facility-alarm critical exceed-action shutdown
no ip routing protocol purge interface
!
ip vrf test
rd 100:100
route-target export 1000:1000
route-target import 1000:1000
!
no ip domain lookup

!
!
!
!
!
!
!
!
!
!
!
mpls label protocol ldp
mpls ldp explicit-null
mpls ldp session protection
mpls ldp discovery targeted-hello accept
multilink bundle-name authenticated
!
!
license udi pid ASR-903 sn FOX1551P04E
license accept end user agreement
license boot level metroaggrservices
sdm prefer default
!
!
redundancy
mode sso
!
!
!
!
!
transceiver type all

```



```
ip address 65.1.116.1 255.255.255.255
!
interface Loopback117
ip address 65.1.117.1 255.255.255.255
!
interface Loopback118
ip address 65.1.118.1 255.255.255.255
!
interface Loopback119
ip address 65.1.119.1 255.255.255.255
!
interface Loopback120
ip address 65.1.120.1 255.255.255.255
!
interface Loopback121
ip address 65.1.121.1 255.255.255.255
!
interface Loopback122
ip address 65.1.122.1 255.255.255.255
!
interface Loopback123
ip address 65.1.123.1 255.255.255.255
!
interface Loopback124
ip address 65.1.124.1 255.255.255.255
!
interface Loopback125
ip address 65.1.125.1 255.255.255.255
!
interface Loopback126
ip address 65.1.126.1 255.255.255.255
!
interface Loopback127
ip address 65.1.127.1 255.255.255.255
!
interface Loopback128
ip address 65.1.128.1 255.255.255.255
!
interface Loopback129
ip address 65.1.129.1 255.255.255.255
!
interface Loopback130
ip address 65.1.130.1 255.255.255.255
!
interface Loopback131
ip address 65.1.131.1 255.255.255.255
!
interface Loopback132
ip address 65.1.132.1 255.255.255.255
!
interface Loopback133
ip address 65.1.133.1 255.255.255.255
!
interface Loopback134
ip address 65.1.134.1 255.255.255.255
!
interface Loopback135
ip address 65.1.135.1 255.255.255.255
!
interface Loopback136
ip address 65.1.136.1 255.255.255.255
!
interface Loopback137
ip address 65.1.137.1 255.255.255.255
```

```
!  
interface Loopback138  
ip address 65.1.138.1 255.255.255.255  
!  
interface Loopback139  
ip address 65.1.139.1 255.255.255.255  
!  
interface Loopback140  
ip address 65.1.140.1 255.255.255.255  
!  
interface Loopback141  
ip address 65.1.141.1 255.255.255.255  
!  
interface Loopback142  
ip address 65.1.142.1 255.255.255.255  
!  
interface Loopback143  
ip address 65.1.143.1 255.255.255.255  
!  
interface Loopback144  
ip address 65.1.144.1 255.255.255.255  
!  
interface Loopback145  
ip address 65.1.145.1 255.255.255.255  
!  
interface Loopback146  
ip address 65.1.146.1 255.255.255.255  
!  
interface Loopback147  
ip address 65.1.147.1 255.255.255.255  
!  
interface Loopback148  
ip address 65.1.148.1 255.255.255.255  
!  
interface Loopback149  
ip address 65.1.149.1 255.255.255.255  
!  
interface Loopback150  
ip address 65.1.150.1 255.255.255.255  
!  
interface Loopback151  
ip address 65.1.151.1 255.255.255.255  
!  
interface Loopback152  
ip address 65.1.152.1 255.255.255.255  
!  
interface Loopback153  
ip address 65.1.153.1 255.255.255.255  
!  
interface Loopback154  
ip address 65.1.154.1 255.255.255.255  
!  
interface Loopback155  
ip address 65.1.155.1 255.255.255.255  
!  
interface Loopback156  
ip address 65.1.156.1 255.255.255.255  
!  
interface Loopback157  
ip address 65.1.157.1 255.255.255.255  
!  
interface Loopback158  
ip address 65.1.158.1 255.255.255.255  
!
```

```
interface Loopback159
ip address 65.1.159.1 255.255.255.255
!
interface Loopback160
ip address 65.1.160.1 255.255.255.255
!
interface GigabitEthernet0/0/0
ip address 21.1.6.2 255.255.255.0
ip router isis core-agg
negotiation auto
mpls ip
!
interface GigabitEthernet0/0/1
ip address 21.2.6.2 255.255.255.0
ip router isis core-agg
negotiation auto
mpls ip
!
interface GigabitEthernet0/0/2
ip address 21.3.6.2 255.255.255.0
ip router isis core-agg
negotiation auto
mpls ip
!
interface GigabitEthernet0/0/3
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0/0/4
ip address 21.4.1.1 255.255.255.0
ip router isis core-agg
shutdown
negotiation auto
mpls ip
!
interface GigabitEthernet0/0/5
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0/0/6
ip address 51.1.0.1 255.255.255.0
negotiation auto
!
interface GigabitEthernet0/0/7
ip address 21.4.6.2 255.255.255.0
ip router isis core-agg
shutdown
negotiation auto
mpls ip
!
interface GigabitEthernet0/1/0
ip address 21.2.1.1 255.255.255.0
ip router isis core-agg
negotiation auto
mpls ip
!
interface GigabitEthernet0/1/1
ip address 21.3.1.1 255.255.255.0
ip router isis core-agg
negotiation auto
mpls ip
!
```



```
interface GigabitEthernet0/1/2
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0/1/3
ip address 21.1.1.1 255.255.255.0
ip router isis core-agg
negotiation auto
mpls ip
!
interface GigabitEthernet0/1/4
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0/1/5
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0/1/6
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0/1/7
no ip address
shutdown
negotiation auto
!
interface GigabitEthernet0
vrf forwarding Mgmt-intf
ip address 7.43.21.101 255.255.0.0
shutdown
negotiation auto
!
router isis core-agg
net 49.0000.0000.1111.00
is-type level-1
metric-style wide
fast-flood 10
ip route priority high tag 10000
set-overload-bit on-startup 360
max-lsp-lifetime 65535
lsp-refresh-interval 65000
spf-interval 5 50 200
prc-interval 5 50 200
lsp-gen-interval 5 50 200
log-adjacency-changes
passive-interface Loopback0
passive-interface Loopback101
passive-interface Loopback102
passive-interface Loopback103
passive-interface Loopback104
passive-interface Loopback105
passive-interface Loopback106
passive-interface Loopback107
passive-interface Loopback108
passive-interface Loopback109
passive-interface Loopback110
passive-interface Loopback111
passive-interface Loopback112
passive-interface Loopback113
```

```
passive-interface Loopback114
passive-interface Loopback115
passive-interface Loopback116
passive-interface Loopback117
passive-interface Loopback118
passive-interface Loopback119
passive-interface Loopback120
passive-interface Loopback121
passive-interface Loopback122
passive-interface Loopback123
passive-interface Loopback124
passive-interface Loopback125
passive-interface Loopback126
passive-interface Loopback127
passive-interface Loopback128
passive-interface Loopback129
passive-interface Loopback130
passive-interface Loopback131
passive-interface Loopback132
passive-interface Loopback133
passive-interface Loopback134
passive-interface Loopback135
passive-interface Loopback136
passive-interface Loopback137
passive-interface Loopback138
passive-interface Loopback139
passive-interface Loopback140
passive-interface Loopback141
passive-interface Loopback142
passive-interface Loopback143
passive-interface Loopback144
passive-interface Loopback145
passive-interface Loopback146
passive-interface Loopback147
passive-interface Loopback148
passive-interface Loopback149
passive-interface Loopback150
passive-interface Loopback151
passive-interface Loopback152
passive-interface Loopback153
passive-interface Loopback154
passive-interface Loopback155
passive-interface Loopback156
passive-interface Loopback157
passive-interface Loopback158
passive-interface Loopback159
passive-interface Loopback160
maximum-paths 8
mpls ldp sync
!
router bgp 100
  bgp router-id 100.111.14.1
  bgp log-neighbor-changes
  neighbor ABR peer-group
  neighbor ABR remote-as 100
  neighbor ABR update-source Loopback0
  neighbor 100.111.10.1 peer-group ABR
  neighbor 100.111.10.1 shutdown
  neighbor 100.111.10.2 peer-group ABR
  neighbor 100.111.10.2 shutdown
!
address-family ipv4
  bgp additional-paths install
  network 100.111.14.1 mask 255.255.255.255 route-map set-PAN-comm
```

```
neighbor ABR send-community both
neighbor ABR next-hop-self all
neighbor ABR route-map deny-PAN-loopbacks in
neighbor ABR send-label
neighbor 100.111.10.1 activate
neighbor 100.111.10.2 activate
exit-address-family
!
address-family vpnv4
neighbor ABR send-community both
neighbor 100.111.10.1 activate
neighbor 100.111.10.2 activate
exit-address-family
!
address-family ipv4 vrf test
redistribute connected
exit-address-family
!
ip forward-protocol nd
!
ip bgp-community new-format
ip community-list 1 permit 100:100
no ip http server
no ip http secure-server
ip route vrf Mgmt-intf 10.0.0.0 255.0.0.0 7.43.0.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 7.43.0.1
!
!
route-map set-service-nh permit 10
!
route-map deny-PAN-loopbacks deny 10
match community 1
!
route-map deny-PAN-loopbacks permit 20
!
route-map set-PAN-comm permit 10
set community 100:100
!
mpls ldp router-id Loopback0
!
!
control-plane
!
alias exec psh reques plat soft sys shell
alias exec shpp6 sh platform hard pp act fea cef da ipv6
alias exec shpp sh platform hard pp act fea cef da ipv4
!
line con 0
exec-timeout 0 0
logging synchronous
stopbits 1
line aux 0
stopbits 1
line vty 0
exec-timeout 0 0
password lab
login
line vty 1 4
login
!
!
!
end
```

