



UCMP Load Balancing

Table 1: Feature History

Feature Name	Release Information	Description
UCMP Load Balancing	Cisco IOS XE Dublin 17.12.1	<p>This feature provides the capability to load balance traffic proportionally across multiple paths, with different cost.</p> <p>Prior to this release, the higher bandwidth links used to carry the same traffic as the lower bandwidth links were underutilized.</p> <p>Use the following new command to configure local Unequal Cost Multi Path (UCMP):</p> <pre>ucmp local <i>prefix-list</i> <i>prefix-list-name</i></pre>

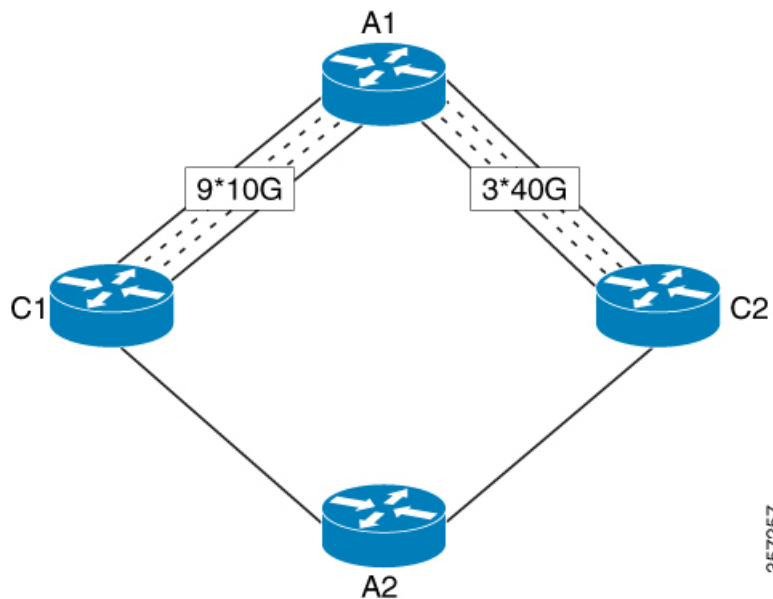
Load balancing is a forwarding mechanism that distributes traffic over multiple links based on certain parameters. Prior to Cisco IOS XE Cupertino Release 17.8.1, the traffic with different bandwidths could not be load balanced. Configuring same metric would make higher bandwidth links carry the same traffic as the lower bandwidth links. Hence, the higher bandwidth links were underutilized.

Generally, higher bandwidth paths have lower Interior Gateway Protocol (IGP) metrics configured, so that they form the shortest IGP paths. Starting with Cisco IOS XE Cupertino Release 17.8.1, with the Unequal Cost Multi Path (UCMP) load balancing enabled, protocols can use even lower bandwidth paths or higher cost paths for traffic, and can install these paths to the Forwarding Information Base (FIB). This feature is only supported on Intermediate System to Intermediate System (IS-IS) protocol.

Starting with Cisco IOS XE Dublin 17.12.1, UCMP is supported on ASR 920 Routers .

Starting with Cisco IOS XE Cupertino 17.8.1, UCMP is supported on .

Figure 1: Example 1: Topology for UCMP




Consider the above topology where there are nine 10G POS links from A1 to C1 and three 40G gigabit ethernet links from A1 to C2. If you want to load balance traffic from A1 to A2, across all the links from A1 to C1 and A1 to C2, you must configure equal metric on all the links. This will create Equal Cost Multipath (ECMP) paths. However, the higher bandwidth links will carry the same traffic as the lower bandwidth links and the higher bandwidth links are underutilized. More specifically, if 9X10G POS links together carry a load of 90 units, then the 3X40G links should also carry a “load” of 90 units, though they have a capability of carrying a total load of 120 units.

To avoid this problem, UCMP allows you configure all the links to distribute the traffic proportionately across the links based on bandwidth, even if the configured metrics on all links are the same. With the UCMP load-balancing enabled, the protocols still install multiple paths to the same destination in FIB, but each path will have a 'load metric/weight' associated with it. FIB uses this load metric/weight to decide the amount of traffic that needs to be sent on a higher bandwidth path and the amount of traffic that needs to be sent on a lower bandwidth path.

In the following example, there are three paths to reach Network X as follows:

Figure 2: Example 2: Topology for UCMP



Paths	Cost from Router E to Net -X
E-B-A	40
E-C-A	50
E-D-A	40

IGP selects the lowest path links, i.e E-B-A and E-D-A. The path E-C-A is not considered for load balancing because of higher cost. The lowest path link E-D (5) is not a tie breaker, as the end to end cost to the Network X is considered.

Restrictions

- UCMP is supported only by the IS-IS protocol.
- IPv6 UCMP is not supported by the IS-IS protocol.
- Native UCMP is not supported.
- UCMP is not supported in L2VPNs.
- In L3VPN, BGP Prefix-Independent Convergence (PIC) core is enabled by default. When the number of prefixes is more than six and UCMP load balancing is converted into Indirect load balancing, ECMP and UCMP do not work. We recommend that you disable BGP PIC Core for UCMP load balancing.
- The router supports only 6K prefixes when UCMP is enabled. This includes both IPv4 and VPNv4 prefixes.
- The packet processor has a maximum limit of 8 buckets to load-balance for each load balance ID. It is applicable for the Cisco IOS XE RSP2 module.

ECMP vs UCMP

Equal Cost Multi Path (ECMP) is a forwarding mechanism for routing packets along multiple paths of equal cost with the goal to achieve almost equally distributed link load sharing. This significantly impacts a router's next-hop (path) decision.

UCMP applies a *weight* to a path. The weight applied is *static* and is derived by the DMZ bandwidth extended community either assigned to a peer or as configured via the Route Policy Language (RPL) route manipulation functionality.

In local UCMP, the best paths have the same metric. Though the metrics are same, IGP calculates the 'load metric or weight' that is based on the bandwidth of each of the links. This information is passed on to FIB and FIB takes care of load balancing the traffic accordingly across the links.

- [Advantages of UCMP, on page 4](#)
- [Configure UCMP Load Balancing, on page 4](#)
- [Verification of UCMP Configuration, on page 4](#)

Advantages of UCMP

- This is a simple process.
- No major changes are required in IGP other than obtaining the bandwidth of the links. You can calculate local UCMP weights based on bandwidths and pass the information to RIB or FIB.

Configure UCMP Load Balancing

To enable UCMP load balancing for IS-IS protocol per address family:

```
router isis
ucmp local prefix-list prefix-list-name
```

To enable local UCMP to calculate local weights based on bandwidths for IPv4 MT-0 routes:

```
router isis
 address-family ipv4 unicast
  ucmp prefix-list list1
```

To enable local UCMP for IPv6 address family:

```
router isis
 address-family ipv6
  ucmp local prefix-list prefix-list-name
```

Verification of UCMP Configuration

Use **show ip route** to verify local UCMP configuration:

```
R1#show ip route 12.12.12.12
Routing entry for 12.12.12.12/32
  Known via "isis", distance 115, metric 60, type level-2
  Redistributing via isis 1
  Last update from 140.0.0.2 on BDI140, 00:00:09 ago
  Routing Descriptor Blocks:
  * 141.0.0.2, from 12.12.12.12, 00:00:09 ago, via BDI141, prefer-non-rib-labels, merge-labels

    Route metric is 60, traffic share count is 1
```

```
MPLS label: 17012
MPLS Flags: NSF
Repair Path: 140.0.0.2, via BDI140
140.0.0.2, from 12.12.12.12, 00:00:09 ago, via BDI140, prefer-non-rib-labels, merge-labels
```

Route metric is 60, traffic share count is 10---UCMP Enabled,7:1 load share Expected

```
MPLS label: 17012
MPLS Flags: NSF
Repair Path: 141.0.0.2, via BDI141
```

