

## Multi-Topology Routing in OSPF

Multi-topology routing (MTR) introduces the capability of essentially class-based forwarding defined using MQC. A topology is a set of routers and links in a network for which a separate set of routes is calculated. A base topology is the default topology that a router uses to calculate a global routing table. Any additional topologies are called class-specific topologies which carry a specific class of traffic.

Multi-topology OSPF (MT-OSPF) exploits unused TOS (Type of Service) field in OSPF LSAs. MT-OSPF redefines TOS fields to carry MTIDs and their associated metrics.

A Router LSA originated by an MT-OSPF router lists all OSPF links for that router and all the router's neighbors. The modified Router LSA also indicates for each link what topologies the link belongs to (indicated by MTID) and the metric for the link specific to that topology (indicated by MTID metric). LSA types 3, 4, 5 and 7 are also modified to carry MTIDs and their metrics. For each MTID listed in the LSDB, a separate SPF is run, and a separate Routing Information Base (RIB) is maintained. The default (base) topology is indicated by MTID of 0. It consists of all routers and links in the OSPF domain. The class-specific topologies lie within the range 32-127.

The following packet capture shows an OSPF Router LSA with 2 MTIDs (although it shows as TOS):

## Open Shortest Path First

- ⊕ OSPF Header
- ⊖ LS Update Packet
  - Number of LSAs: 1
  - ⊖ LS Type: Router-LSA
    - LS Age: 3600 seconds
    - Do Not Age: False
    - ⊕ Options: 0x22 (DC, E)
      - Link-State Advertisement Type: Router-LSA (1)
      - Link State ID: 1.1.1.1
      - Advertising Router: 1.1.1.1 (1.1.1.1)
      - LS Sequence Number: 0x80000005
      - LS Checksum: 0x2c09
      - Length: 44
    - ⊕ Flags: 0x00 ()
      - Number of Links: 1
    - ⊖ Type: Transit ID: 10.1.1.2      Data: 10.1.1.1      Metric: 10
      - IP address of Designated Router: 10.1.1.2
      - Link Data: 10.1.1.1
      - Link Type: 2 - Connection to a transit network
      - Number of TOS metrics: 2
      - TOS 0 metric: 10
      - TOS: 32, Metric: 55
      - TOS: 33, Metric: 75

### Sample Scenario

The following network in figure 1 shows a base topology for which OSPF is enabled on all interfaces using **network 0.0.0.0 255.255.255.255 area 0** OSPF router configuration command. The metric for base topology is set to 10 for all interfaces using **ip ospf metric 10** interface configuration command.




Figure 1. Base Topology

The following output shows that R1 has 2 equal-cost paths to reach destination 4.4.4.4 on R4 via **R1-R2-R4** and **R1-R3-R4**.

Routing to 4.4.4.4 using base topology

```

R1# show ip route 4.4.4.4
Routing entry for 4.4.4.4/32
  Known via "ospf 1", distance 110, metric 21, type intra area
  Last update from 10.13.1.2 on FastEthernet1/1, 00:03:44 ago
  Routing Descriptor Blocks:
    10.13.1.2, from 4.4.4.4, 00:03:44 ago, via FastEthernet1/1
      Route metric is 21, traffic share count is 1
    * 10.12.1.2, from 4.4.4.4, 00:03:54 ago, via FastEthernet1/0
      Route metric is 21, traffic share count is 1

```

```

R1# traceroute 4.4.4.4

```

```

Type escape sequence to abort.
Tracing the route to 4.4.4.4

```

```

 0 10.12.1.2 140 msec
 1 10.13.1.2 124 msec
 2 10.12.1.2 120 msec
 3 10.34.1.2 152 msec
 4 10.24.1.2 148 msec
 5 10.34.1.2 172 msec

```

## Enabling MT-OSPF for Critical (Red) and Non-critical (Blue) Topologies

These topologies are overlaid on top of the base topology in figure 1. It is required that for all Critical class, only links R1-R2 and R2-R4 be used with the metric defined in figure 2.




Figure 2. Critical (Red) Topology

For all non-critical class, only links R1-R2, R2-R3 and R3-R4 be used with metric defined in figure 3.




Figure 3. Non-Critical (Blue) Topology

## Step 1. Configuring the unicast topologies for MTR

### Configuring topologies

```

global-address-family ipv4
  topology CRITICAL
    forward-base
  !
  topology NON_CRITICAL
    forward-base
  !
  !

```

The **global-address-family** command allows to define the topology. Only IPv4 address-family is supported. The default sub-address family is unicast. The other is multicast. The **topology** command is entered to label a class-specific topology. An optional **all-interfaces** command under topology instance configures the instance to use all interfaces on the router.

The **forward-base** command under topology instance enables the router to look first in the class-specific topology FIB. If a forwarding route is not found, then the router will look in the base topology FIB. This is also called **Incremental mode**. The default is **Strict mode** which configures the router to look into class-specific topology FIB only. If the path is not found, the packet is dropped.

## Step 2. MTR Traffic Classification

Traffic classification is used to assign different classes of traffic to different topologies. Traffic classification is configured using MQC. Traffic classification is defined globally for each topology rather than on interface.

```

class-map match-any CRITICAL_TRAFFIC
  match ip dscp 46
!
class-map match-any NON_CRITICAL_TRAFFIC
  match ip dscp af31
  match ip dscp af32
  match ip dscp af33
!
policy-map type class-routing ipv4 unicast TOPOLOGY_POLICY
  class CRITICAL_TRAFFIC
    select-topology CRITICAL
  !
  class NON_CRITICAL_TRAFFIC
    select-topology NON_CRITICAL
  !
!
global-address-family ipv4
  topology CRITICAL
    forward-base
  !
  topology NON_CRITICAL
    forward-base
  !
service-policy type class-routing TOPOLOGY_POLICY
!
!

```

### Step 3. Enabling topologies to use interfaces

Since ***all-interfaces*** command was not configured under topology instances, topology instances must be enabled on per-interface basis. Topology specific OSPF metric (cost) can be configured here.

#### Enabling topology instances on interfaces on R1 and R2

##### R1 router

```

interface Loopback 0
  ip address 1.1.1.1 255.255.255.255
  !
  topology ipv4 unicast CRITICAL
  !
  topology ipv4 unicast NON_CRITICAL
  !
!
interface Fa 1/0
  ip address 10.12.1.1 255.255.255.0
  ip ospf cost 10

```

##### R2 router

```

interface Fa 1/0
  ip address 10.12.1.2 255.255.255.0
  ip ospf cost 10
  !
  topology ipv4 unicast CRITICAL
    ip ospf cost 20
  !
  topology ipv4 unicast NON_CRITICAL
    ip ospf cost 50
  !
!

```

```

!
topology ipv4 unicast CRITICAL
ip ospf cost 20
!
topology ipv4 unicast NON_CRITICAL
ip ospf cost 50
!
!
interface Fa 1/1
ip address 10.13.1.1 255.255.255.0
ip ospf cost 10
!
!
interface Fa 1/1
ip address 10.24.1.1 255.255.255.0
ip ospf cost 10
!
!
topology ipv4 unicast CRITICAL
ip ospf cost 20
!
!
interface Fa 2/0
ip address 10.23.1.1 255.255.255.0
ip ospf cost 10
!
!
topology ipv4 unicast NON_CRITICAL
ip ospf cost 30
!
!

```

## Enabling topology instance on interfaces on R3 and R4

### R3 router

```

interface Fa 1/0
ip address 10.13.1.2 255.255.255.0
ip ospf cost 10
!
interface Fa 1/1
ip address 10.23.1.2 255.255.255.0
!
topology ipv4 unicast NON_CRITICAL
ip ospf cost 30
!
!
interface Fa 2/0
ip address 10.34.1.1 255.255.255.0
!
topology ipv4 unicast NON_CRITICAL
ip ospf cost 30
!
!

```

### R4 router

```

interface Loopback 0
ip address 4.4.4.4 255.255.255.255
!
topology ipv4 unicast CRITICAL
!
topology ipv4 unicast NON_CRITICAL
!
!
interface Fa 1/0
ip address 10.24.1.2 255.255.255.0
ip ospf cost 10
!
topology ipv4 unicast CRITICAL
ip ospf cost 20
!
!
interface Fa 1/1
ip address 10.34.1.2 255.255.255.0
ip ospf cost 10
!
topology ipv4 unicast NON_CRITICAL
ip ospf cost 30
!
!

```

## Step 4. Activating MTR using OSPF

All the topologies defined in global-address-family must be activated in OSPF for it to send

the metrics for a specific topology in LSAs. Also, each topology must be assigned a topology ID (tid) in OSPF from the range 32 - 127. The following configuration must be applied to all routers:

### Activating MTR using OSPF

```
router ospf 1
network 0.0.0.0 255.255.255.255 area 0
!
address-family ipv4
topology CRITICAL tid 32
!
topology NON_CRITICAL tid 33
!
!
!
```

### Verification

The following output shows the LSDB on R1 router. Notice the Router LSAs received have metric for specific topologies along with MTIDs. Also note that from router LSAs, R1 figures out which advertising routers are reachable on which topology. Hence, R3 is not reachable on CRITICAL topology.

### LSDB on R1

```
R1# show ip ospf database router
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1)
```

#### Router Link States (Area 0)

```
LS age: 527
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 1.1.1.1
Advertising Router: 1.1.1.1
LS Seq Number: 80000008
Checksum: 0x1299
Length: 76
Number of Links: 3

Link connected to: a Stub Network
(Link ID) Network/subnet number: 1.1.1.1
(Link Data) Network Mask: 255.255.255.255
Number of MTID metrics: 2
```

TOS 0 Metrics: 1            *! TOS 0 is metric for base topology*  
MTID 32 Metrics: 1        *! metric for CRITICAL topology (MTID 32)*  
MTID 33 Metrics: 1        *! metric for NON\_CRITICAL topology (MTID 33)*

Link connected to: a Transit Network  
(Link ID) Designated Router address: 10.13.1.1  
(Link Data) Router Interface address: 10.13.1.1  
Number of MTID metrics: 0  
TOS 0 Metrics: 10

Link connected to: a Transit Network  
(Link ID) Designated Router address: 10.12.1.1  
(Link Data) Router Interface address: 10.12.1.1  
Number of MTID metrics: 2  
TOS 0 Metrics: 10  
MTID 32 Metrics: 20  
MTID 33 Metrics: 50

LS age: 987  
Options: (No TOS-capability, DC)  
LS Type: Router Links  
Link State ID: 2.2.2.2  
Advertising Router: 2.2.2.2  
LS Seq Number: 8000000F  
Checksum: 0x3CCE  
Length: 76  
Number of Links: 3

Link connected to: a Transit Network  
(Link ID) Designated Router address: 10.23.1.1  
(Link Data) Router Interface address: 10.23.1.1  
Number of MTID metrics: 1  
TOS 0 Metrics: 10  
MTID 33 Metrics: 30

Link connected to: a Transit Network  
(Link ID) Designated Router address: 10.24.1.1  
(Link Data) Router Interface address: 10.24.1.1  
Number of MTID metrics: 1  
TOS 0 Metrics: 10  
MTID 32 Metrics: 20

Link connected to: a Transit Network  
(Link ID) Designated Router address: 10.12.1.1  
(Link Data) Router Interface address: 10.12.1.2  
Number of MTID metrics: 2  
TOS 0 Metrics: 10  
MTID 32 Metrics: 20  
MTID 33 Metrics: 50

Adv Router is not-reachable in topology CRITICAL with MTID 32  
LS age: 985



Options: (No TOS-capability, DC)

LS Type: Router Links

Link State ID: 3.3.3.3

Advertising Router: 3.3.3.3

LS Seq Number: 80000005

Checksum: 0x2F4D

Length: 68

Number of Links: 3

Link connected to: a Transit Network

(Link ID) Designated Router address: 10.34.1.1

(Link Data) Router Interface address: 10.34.1.1

Number of MTID metrics: 1

TOS 0 Metrics: 10

MTID 33 Metrics: 30

Link connected to: a Transit Network

(Link ID) Designated Router address: 10.23.1.1

(Link Data) Router Interface address: 10.23.1.2

Number of MTID metrics: 1

TOS 0 Metrics: 10

MTID 33 Metrics: 30

Link connected to: a Transit Network

(Link ID) Designated Router address: 10.13.1.1

(Link Data) Router Interface address: 10.13.1.2

Number of MTID metrics: 0

TOS 0 Metrics: 10

LS age: 457

Options: (No TOS-capability, DC)

LS Type: Router Links

Link State ID: 4.4.4.4

Advertising Router: 4.4.4.4

LS Seq Number: 80000007

Checksum: 0x68EF

Length: 76

Number of Links: 3

Link connected to: a Stub Network

(Link ID) Network/subnet number: 4.4.4.4

(Link Data) Network Mask: 255.255.255.255

Number of MTID metrics: 2

TOS 0 Metrics: 1

MTID 32 Metrics: 1

MTID 33 Metrics: 1

Link connected to: a Transit Network

(Link ID) Designated Router address: 10.34.1.1

(Link Data) Router Interface address: 10.34.1.2

Number of MTID metrics: 1

TOS 0 Metrics: 10

MTID 33 Metrics: 30

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.24.1.1
(Link Data) Router Interface address: 10.24.1.2
Number of MTID metrics: 1
  TOS 0 Metrics: 10
  MTID 32 Metrics: 20
```

The following output shows the topology routing table entry for 4.4.4.4. Notice the metric to reach 4.4.4.4 from R1 for CRITICAL topology is 41 while the metric for NON\_CRITICAL topology is 111.

### Routing table entries for 4.4.4.4 for specific topologies

```
R1# show ip route topology CRITICAL 4.4.4.4
```

```
Routing Table: CRITICAL
Routing entry for 4.4.4.4/32
  Known via "ospf 1", distance 110, metric 41, type intra area
  Last update from 10.12.1.2 on FastEthernet1/0, 03:02:59 ago
  Routing Descriptor Blocks:
  * 10.12.1.2, from 4.4.4.4, 03:02:59 ago, via FastEthernet1/0
    Route metric is 41, traffic share count is 1
```

```
R1# show ip route topology NON_CRITICAL 4.4.4.4
```

```
Routing Table: NON_CRITICAL
Routing entry for 4.4.4.4/32
  Known via "ospf 1", distance 110, metric 111, type intra area
  Last update from 10.12.1.2 on FastEthernet1/0, 03:03:01 ago
  Routing Descriptor Blocks:
  * 10.12.1.2, from 4.4.4.4, 03:03:01 ago, via FastEthernet1/0
    Route metric is 111, traffic share count is 1
```

Now notice the path taken for traffic to destination 4.4.4.4 with **DSCP value 46** i.e. for traffic belonging to CRITICAL topology. It takes the optimal path **R1-R2-R4**.

### Traceroute to 4.4.4.4 using CRITICAL topology

```
R1# traceroute topology CRITICAL ip
```

```
Target IP address: 4.4.4.4
Source address: 1.1.1.1
DSCP Value [0]: 46
Numeric display [n]:
Resolve AS number in (G)lobal table, (V)RF or(N)one [G]:
Timeout in seconds [3]:
Probe count [3]:
```

```
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 4.4.4.4
```

```
 1 10.12.1.2 144 msec 136 msec 152 msec
 2 10.24.1.2 156 msec 168 msec *
```

Now notice the path taken for traffic to destination 4.4.4.4 with **DSCP value 26 (DSCP AF32)** i.e. for traffic belonging to NON\_CRITICAL topology. It takes the path **R1-R2-R3-R4**.

### Traceroute to 4.4.4.4 using NON\_CRITICAL topology

```
R1# traceroute topology NON_CRITICAL ip
```

```
Target IP address: 4.4.4.4
Source address: 1.1.1.1
DSCP Value [0]: 26
Numeric display [n]:
Resolve AS number in (G)lobal table, (V)RF or(N)one [G]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 4.4.4.4
```

```
 1 10.12.1.2 108 msec 120 msec 156 msec
 2 10.23.1.2 180 msec 140 msec 148 msec
 3 10.34.1.2 220 msec 200 msec *
```