

## OSPF network types

OSPF addresses three classes of network (as listed in section 1.2 of [RFC 2328](#)): point-to-point, broadcast, and non-broadcast.

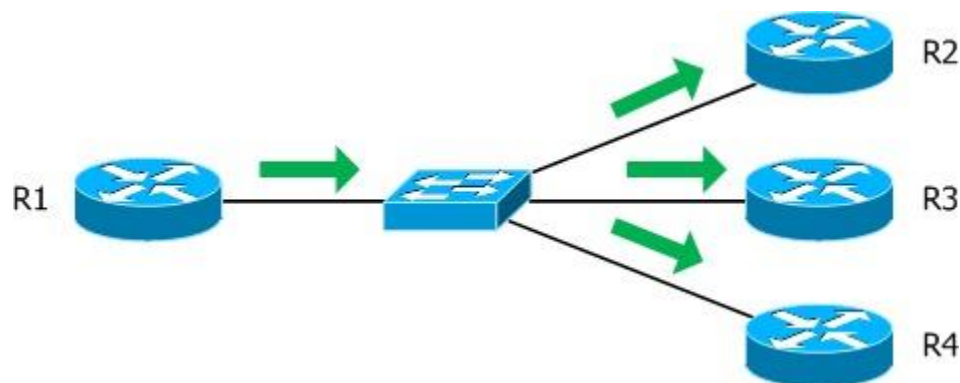
### Point-to-Point



This is by far the simplest network type, and serves as a convenient anchor from which to advance the discussion. A point-to-point network is, as its name aptly describes, a link between exactly two points (or routers). A packet sent from one of the routers will always have exactly one recipient on the local link.

### Broadcast

Obviously, point-to-point links don't scale well. A much more efficient manner of connecting a large number of devices is to implement a *multiaccess* segment; that is, a segment which can be accessed by multiple end points. An Ethernet segment is an example of such a network.



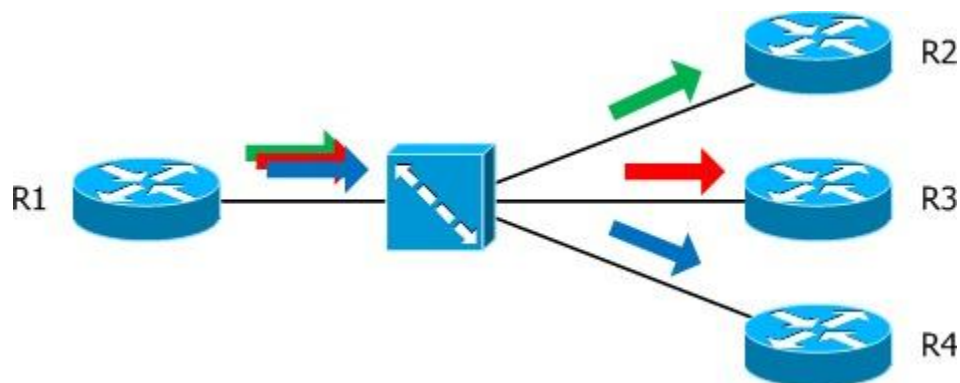
Ethernet networks support *broadcasts*; a single packet transmitted by a device can be multiplied by the medium (in this case an Ethernet switch) so that every other end point receives a copy. This is advantageous not only in bandwidth savings, but also in facilitating automatic neighbor discovery.

In the example pictured above, R1 can multicast (a broadcast intended only for certain recipients) an OSPF hello message to the link, knowing that all other OSPF routers connected to the link will receive it and reply with their own multicast message. Consequently, neighbors can quickly identify each other and form adjacencies without knowing addresses beforehand. Isn't that convenient?

OSPF routers on a multiaccess segment will elect a designated router (DR) and backup designated router (BDR) with which all non-designated routers will form an adjacency. This is to ensure that the number of adjacencies maintained does not grow too large; a segment containing ten routers would require 45 adjacencies to form a mesh, but only 17 when a DR and BDR are in place.

### Non-Broadcast

Unfortunately, not all multiaccess technologies support broadcast transmissions. Frame relay and ATM are probably the most common examples of non-broadcast transport, requiring individual permanent virtual circuits (PVCs) to be configured between end points.



Notice in the frame relay topology pictured above, R1 must craft and transmit an individual packet for every destination he wants to reach. Aside from being horribly inefficient with regard to bandwidth, this limitation requires the router to know the addresses of his neighbors before he can communicate to them.

OSPF can operate in one of two modes across a non-broadcast network: non-broadcast multi-access (NBMA) or point-to-multipoint. Each of these topologies tackles the absence of broadcast capability from a different direction.

#### ✚ **Non-Broadcast Multi-Access (NBMA)**

An NBMA segment emulates the function of a broadcast network. Every router on the segment must be configured with the IP address of each of its neighbors. OSPF hello packets are then individually transmitted as unicast packets to each adjacent neighbor. As in a true broadcast network, a DR and BDR are elected to limit the number of adjacencies formed.

#### ✚ **Point-to-Multipoint**

A point-to-multipoint configuration approaches the non-broadcast limitation in a different manner. Rather than trying to emulate broadcast capability, it seeks to organize the PVCs into a collection of point-to-point networks. Hello packets must still be replicated and transmitted individually to each neighbor, but the multipoint approach offers two distinct advantages: no DR/BDR is needed, and the emulated point-to-point links can occupy a common subnet.

All routers attached to a non-broadcast network must be manually configured to recognize it as a point-to-multipoint segment:

```
Router(config-if)# ip ospf network point-to-multipoint [non-broadcast]
```

The non-broadcast parameter can be appended to the OSPF network type to force unicasting of packets rather than relying on multicast. This might be necessary when dynamic circuits are in use.