

OSPF DR Election - Why Priority Doesn't Take Effect Right Away

Contributed by Dale Henninger
Tuesday, 15 November 2005
Last Updated Monday, 08 May 2006

I get asked quite often why the OSPF designated router will not change immediately when you set a high priority on a router. This is my stab at explaining why this is, and why you may see some strange behavior. The router priority and OSPF router ID are the factors for determining who will become the DR and BDR. If a loopback interface is present the lowest numbered loopback interface will become the OSPF router ID. But the router ID can be an address of another interface if the loopback does not exist. It is a best practice to always configure a loopback with OSPF. The problem begins with the re-election process not being pre-emptive. The first router that runs OSPF on the network will become the DR by default because a DR must exist in a multi-access network. The problem is that if the 2nd router on the network has a better priority or higher router ID, it will not pre-empt and de-throne the existing DR. If you clear the OSPF process on the DR (clear ip ospf process) you will force a DR/BDR election. This new election will take into account all of the priorities and router ID's of the routers on the network at that time and correct choose the best based on the above rules. So even if your 2nd router comes online just a second after the first router, it may sit as a BDR even though it has a higher priority. RFC 2328 (<http://www.faqs.org/rfcs/rfc2328.html>) in section 9.4 states that the DR election is triggered by the Interface State Machine in OSPF. I interpret this as the only time a router will trigger an election is when that router interface changes state (or you reset OSPF). So once a DR is elected, no other routers will start an election because they are happy with the DR that they see in the first hello packet that they receive from a neighbor. A new election process has to be initiated on the DR with a clear command or that router going down. Here is where the mystery is. Based on different versions of Cisco IOS router code, I have observed different behavior. On many 12.1 versions if you walk away for a while (20-30 minutes) you may come back to see that they have re-elected themselves. But I have also observed where it will go 8-12 hours without an election. This could be where the bug is. I have no idea why they would have re-elected. I don't see an interface go down, etc... Copyright 2006 - Dale R. Henninger