OSPF Convergence

Resiliency and redundancy to circuit failure is provided by the convergence capabilities of OSPF at layer 3. There are two components to OSPF routing convergence: **detection** of topology changes and **recalculation** of routes. Detection of topology changes is supported in two ways by OSPF. The first, and quickest, is a failure or change of status on the physical interface, such as Loss of Carrier. The second is a timeout of the OSPF hello timer. An OSPF neighbor is deemed to have failed if the time to wait for a hello packet exceeds the dead timer, which defaults to four times the value of the hello timer. On a Serial, Fast Ethernet or Gigabit Ethernet interface, the default hello timer is set to 10 seconds; therefore the dead timer is 40 seconds. Recalculation of routes is done by each router after a failure has been detected. A link-state advertisement (LSA) is sent to all routers in the OSPF area to signal a change in topology. This causes all routers to recalculate all of their routes using the Djikstra (SPF) algorithm. This is a CPU intensive task, and a large network, with unreliable links, could cause a CPU overload. When link goes down and if layer 2 is not able to detect the failure, convergence can be improved by decreasing the value of the hello timer. The timer should not be set too low as this may cause phantom failures, hence unnecessary topology recalculations. Remember that these timers are used to detect failures that are not at the physical level. For example, carrier still exists but there is some sort of failure in the intermediate network. Once a topology change has been detected, a LSA is generated and flooded to rest of the devices in the network. Recalculation of the routes will not occur until the spf timer has expired. The default value of this timer is 5 seconds. SPF hold time is also used to delay consecutive SPF calculations (give the router some breathing space). The default for his value is 10 seconds. As a result, the min time for the routes to converge in case of failure is always going to be more than 5 secs unless the SPF timers are tuned using OSPF throttle timers. It is also possible to schedule SPF to run right after flooding the LSA information but this can potentially cause the instabilities in the network e.g. even a flash congestion in the network for a very short duration could declare the link down and trigger the SPF run. The following is the Cisco Default Timers: **ip ospf hello-interval 10 sec ip ospf dead-interval 40 sec (4 x hello) ip ospf retransmit-interval 5 sec ip ospf transmit-delay 1 sec timers spf spf-delay 5 sec timers spf spf-holdtime 10 sec LSA Generation Interval 0.5 sec** Each of these timers will affect the performance of OSPF and can be tuned to effect both convergence time and network resource utilization, but care should be taken in changing these values.