

Implementing Router Interface Redundancy

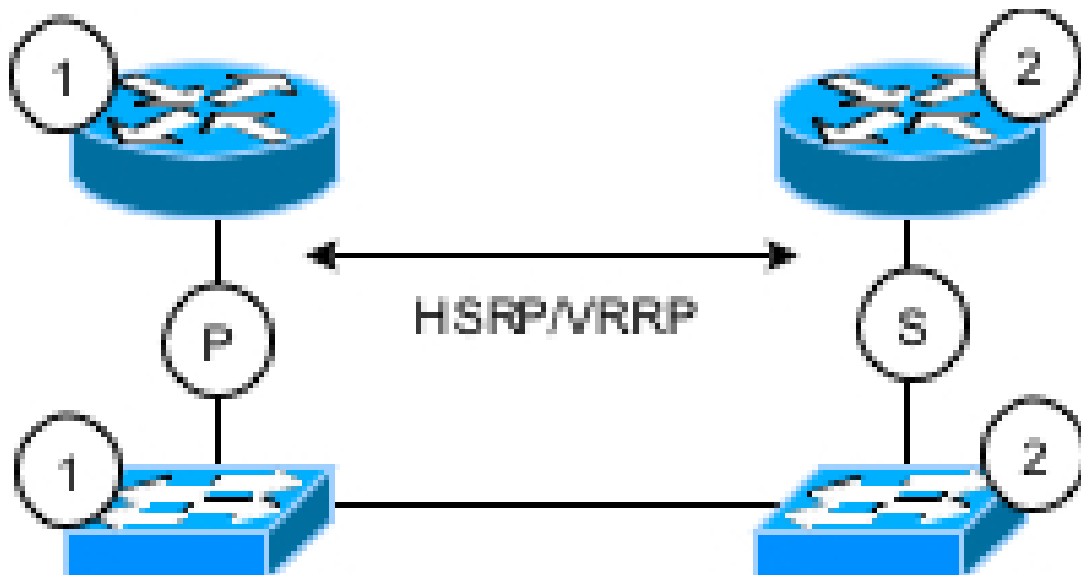
Michael J. Martin

Redundancy is the foundation of any high-availability network design. The concept of redundancy is quite straightforward. "A redundant system design is one that accounts for component failure and can continue to provide service. In the event that a failure occurs, the system will continue to perform at the normal or marginally diminished operational service level." While this is simple as a concept, redundancy becomes quite problematic in practice.

To start, in a data network there are three potential areas of failure:

- OSI Layer 1 – Physical (cabling)
- OSI Layer 2 – Data link (data transmission protocol, i.e., Ethernet, PPP, HDLC, etc.)
- OSI Layer 3 – Network (data delivery protocol, i.e., Internet Protocol, IPX, etc.)

Arguably, the three areas are interdependent. So using this assumption, many network administrators build for redundancy using what I like to call the "two of everything" approach. The premise is simple; if you need one to operate, then use two. That way the second can take over in the event of failure. Here is a typical implementation of this approach:

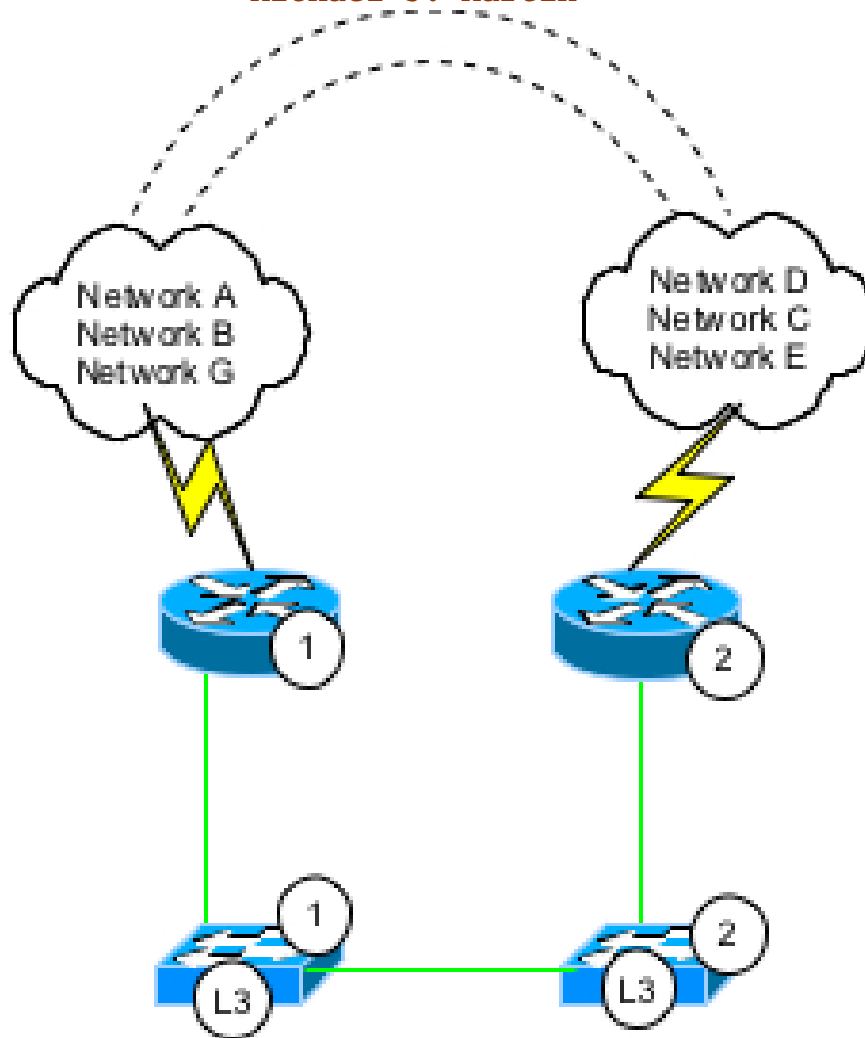


Two routers, two switches, L2 or L3 -- it does not matter. Group 1 makes up the primary path. Group 2 makes up the secondary. The primary path utilizes a high-speed WAN connection, while the secondary path utilizes a low-bandwidth or DDR type of connection. HSRP/VRRP can be run on the switches if they are L3. In a L2 switched environment, HSRP/VRRP can be run on the routers.

On its face, this is a redundant design. However, it is not really effective, and for several reasons. When using high- and low-speed links to provide primary and backup transport, this design can only tolerate a single component failure. Furthermore, this design is slanted toward recovering a total router hardware failure. If there is bad cable in the primary path, you're on backup. If it's a switch failure, you've lost half your users and you're on backup. If it's a router interface failure, you're on backup. In the event of a bad cable in Group 1 and a switch failure in Group 2, your network is down and out. Even more complex topologies that are designed to both fully utilize resources and accommodate for failure make similar mistakes.

Implementing Router Interface Redundancy

Michael J. Martin



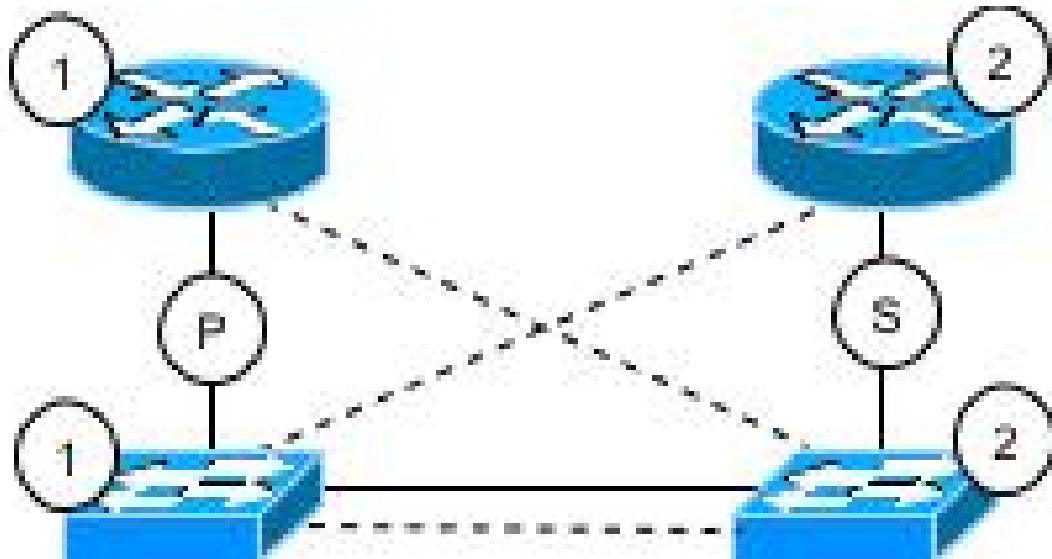
The example topology above is one commonly found in data center or server farm scenarios. Group 1 and Group 2 both provide primary path connectivity for a discrete group of networks. These remote networks are also interconnected at the remote end. So Group 1 can provide access to Group 2's network in the event of a failure and vice versa through the remote interconnections (granted, with much higher latency). However, just like our first design, this one is also only capable of sustaining a single component failure. If you lose a router interface, cable, or switch in one group, then half the network is running on backup. If you lose a component in each Group, then the network has bought the farm. To address these kinds of design issues, a shift in perspective is needed.

At first, it may seem you should add additional routers and switches, particularly in situations like Example 2, where multiple active WAN or MAN connections exist. The "spread risk" approach works, but it is costly and does not address the issue of cable or interface failure. Effective high-availability designs focus beyond component failure and instead design for link failure.

Of course, to design for link failure, you need to increase the number of links or paths that data can flow through when a failure arises. Let's take our first example network and increase the router and switch links to two.

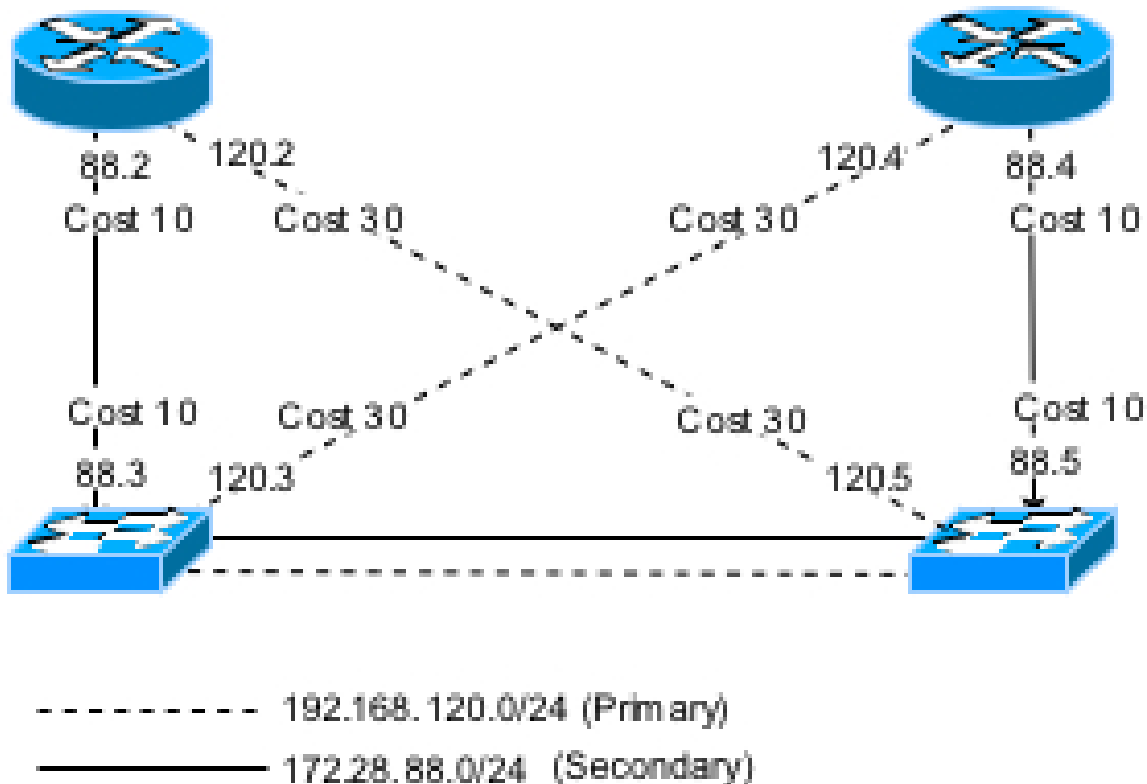
Implementing Router Interface Redundancy

Michael J. Martin



Think of this as the "two of everything" approach on steroids. By doubling the links, multiple failures can be sustained and the network can still maintain "normal" operation. This kind of topology can be accomplished at either the OSI-RM Layer 3 or the OSI-RM Layer 2 level.

The Layer 3 approach requires the use of L3 switches and routers. The interconnection network between the WAN/MAN routers and the "office" or data center LAN uses two different IP subnets with different dynamic routing protocol metric values, weighting one path to be preferred over the other. Here is an example of this using the Open Shortest Path First (OSPF) protocol and its only metric cost.

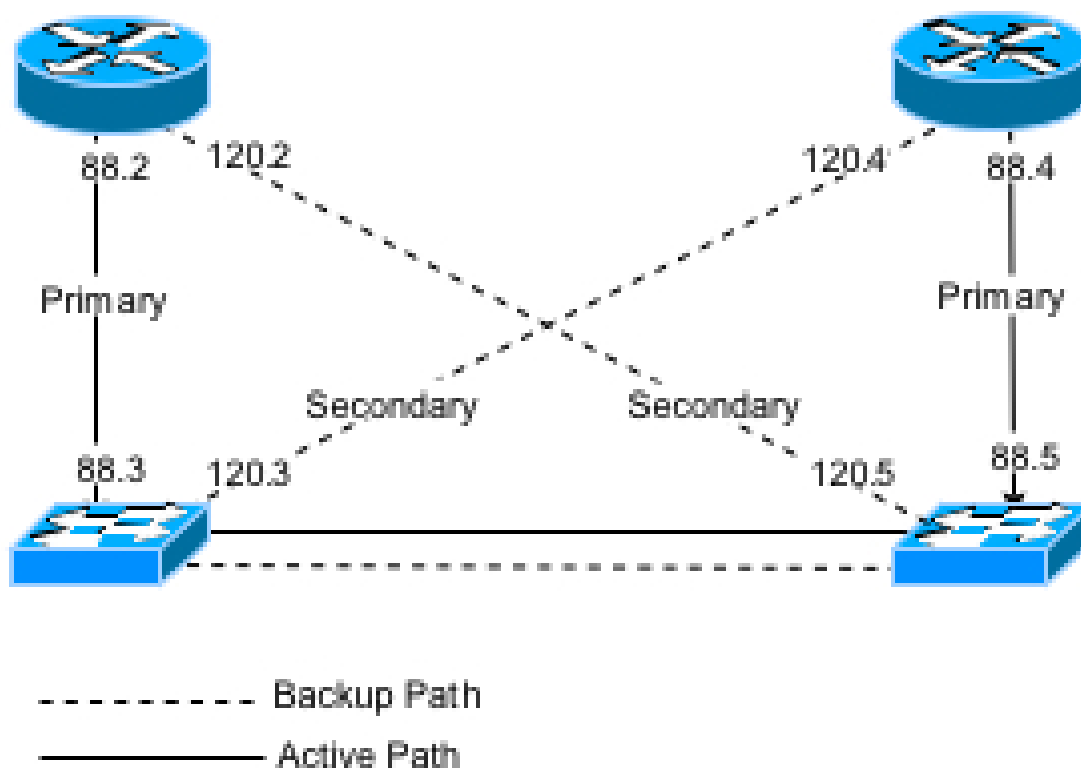


Implementing Router Interface Redundancy

Michael J. Martin

The primary path is subnet 192.168.120.0/24. OSPF assigns default costs to all router interfaces, according to type. In our example, all of the router and switch interfaces are Fast Ethernet, so we must modify the OSPF costs to set up the preferred path. This is done using the OSPF interface command `<ip ospf cost {1-65535}>`. If the network uses Enhanced Interior Gateway Routing Protocol (EIGRP), we can modify the router interface bandwidth or delay metrics to affect the path preference. These values are set using the EIGRP interface commands `<bandwidth {1-10000000 kbit}>` and `<delay {1-16777215 usec}>`. In our example, the defaults for Fast Ethernet are 100,000 Kbps and 100 ms. Cisco recommends using the delay statement for modifying route preference. So to prefer the 172.28.88.0/24 path, we could leave the delay as default and change the delay on the interfaces in the 192.168.120.0/24 to Ethernet's default of 1000 ms.

The Layer 2 approach can be implemented with or without L3 switches. Link recovery is achieved using the backup interface command.



The backup interface command "bonds" two interfaces on the router to behave as one. One is the primary, the other secondary. When the router detects a link failure on the primary, the secondary becomes active and remains active until the primary is restored. The backup interface command is defined on the primary interface:

- Router (config)#int fa 0/0
- Router (config-if)#backup interface Fast Ethernet 1/0

Both the primary interface and its backup should be configured identically. The primary interface dictates the operation state of the bond. If the primary interface is shut down, the backup is as well. Another little caveat of using interface backup is that the primary interface's slot and/or number must be lower than the backup interface. Otherwise, when the router reboots, the backup interface comes up first at boot. The router detects the IP address duplication and wipes out the IP address

Implementing Router Interface Redundancy

Michael J. Martin

configuration of the primary. When the router is up and running, you should check the operation status of the bond using the <show backup> command:

```
Router#show backup
Primary Interface Secondary Interface Status
-----
FastEthernet0/0 FastEthernet2/0 Normal
FastEthernet1/0 FastEthernet2/1 backup mode
```

Router#

One last thing when implementing interface backup -- be sure to either disable spanning tree or enable Cisco's STP portfast option on the switch ports to which the router's interfaces are connected. Otherwise, when the interfaces fail over, there will be a 60-second latency before the backup comes on line. This is because the switch port must go through the spanning tree process.

Which approach is better? Whenever there is more than one way to accomplish something, there will be opinions. Both approaches have merits. The L2 method is simpler. It works with both L2 and L3 switches. The L3 method is slightly more complicated. It has the advantage of being configured to prefer a single path or to load balance across both paths (if both router interfaces are the same type). The downside is that when a failure exists, routing will become asymmetric. In some cases, this is not an issue. However, in situations where load balancing or application caching is needed, this must be accounted for in the design, making sure the ingress and egress path flow for traffic is maintained.