White Papers

from the files of

Networking Unlimited, Inc.

http://www.networkingunlimited.com

Multi−homing−−Connecting to Two ISPs

by Dr. Vincent C. Jones, PE

Version 1.02a — 8 August 2001

Many organizations depend upon Internet connectivity to support critical applications. One popular approach for improving Internet connectivity is to connect to more than one Internet service provider (ISP), a technique called multi−homing.

Multi−homing can be very effective for ensuring continuous connectivity−− eliminating the ISP as a single point of failure−−and it can be cost effective as well. However, your multi−homing strategy must be carefully planned to ensure that you actually improve connectivity for your company.

Important Copyright and License Information

The Need for Physical Diversity

Taking advantage of redundant links requires three conditions to always be present. First, we must be able to detect when a link has failed. Second, we must have a mechanism for directing traffic which would normally flow across a failed link to take the path which is still functional. Third, meeting the first two conditions only helps us to the extent that whatever causes the failure of one link does not also result in failure of its backup.

Let's look at the last requirement first because no protocol design will save us if all links have failed. Because most network failures are due to problems in the WAN links, it does little good to connect to a second ISP if both ISP links are carried over the same communications circuit. Even if independent circuits are used, if they are not physically diverse they will still be subject to common failure events such as construction work inside your building or digging in the street outside. This independence ultimately needs to extend to the physical environment in the data center, where the routers and interfaces should have independent sources of power and be physically separated so that an accident affecting one link will not affect the devices supporting the other.

Providing complete physical diversity is difficult and expensive, and the requirement is not limited to ISP connections. All critical network links for internal communications should also be diversified. Assuming a well designed internal network, the easiest way to achieve physical diversity in your ISP connections is to connect from two different locations that are already well connected to each other but far enough apart to not share any common facilities or communications infrastructure.

Routing with the Border Gateway Protocol

Once we have physical connectivity with adequate diversity in place, we can address the first two requirements. In a multi−homing environment, we normally detect and react to link failures by running Border Gateway Protocol (BGP) between our routers and those of the ISP.

BGP is frequently assumed to mean complex configurations on expensive, high−end routers in order to handle the huge routing tables required to fully describe the Internet. However, depending upon the specific application requirements and the degree of load−balancing desired across all available links, it is frequently practical to implement multi−homing using the smallest routers available capable of handling the traffic load.

In other words, the decision to implement multi−homing does not have to be an all−or−nothing choice. There are choices you can make along the way based upon the equipment you have available and the level of connectivity you need to provide. The underlying need for multihoming should also be examined. For many users, adequate availability may be achievable at lower cost by using multiple links to a single, higher quality ISP. This is particularly true in environments where load balancing is important as well as availability.

Multi−homing without BGP

Even in a multihomed situation, BGP is not the only solution. If your goal is to simply provide internal users with access to the Internet, there is no need to run BGP at all. As long as the link layer protocol supports the exchange of keep−alive messages from router to router, link failure will be detected by the link layer protocol. Floating static routes can then reliably direct all outbound traffic to a working ISP link.

Network Address Translation (NAT) is then used to send outbound packets with a source IP address associated by the ISP with that outbound link. Return traffic will automatically come back via the same, working link because that link is the only link servicing that address range.

Of course this approach will not work if we are providing services to the outside world, as the addresses associated with the failed link will disappear from the Internet. Similarly, connections which were established over the link which failed will need to be reconnected. However, for many applications this impact is minor.

This approach is also sufficient to provide high availability virtual private networks (VPN) across the Internet if we use a routing protocol to detect and route around failed IPSec tunnels.

When BGP Is Mandatory

Only when we need to support a common IP address range using both ISPs do we need to run BGP. This will normally be the case any time our applications include providing services to Internet users. Before we turn on BGP, we must obtain a globally unique autonomous system number (ASN). The ASN is used by BGP to identify our organization as the source of routing advertisements for our network prefixes. We can then arrange with both ISPs to accept BGP advertisements of our IP address prefixes from us and to in turn advertise those address prefixes to the rest of the Internet.

Getting our address prefixes advertised is usually not a problem, although we do have to use care in our configuration to ensure that we do not inadvertently advertise any other address prefixes. In particular, we must ensure that we do not advertise ourselves as a path between the two ISPs or we could find our links consumed by transit traffic of no interest to us. More challenging is setting up our advertisements so that incoming traffic is reasonably balanced between the ISP links. This is extremely implementation dependent and in some situations may not even be possible.

BGP Implementation Choices

The final decision is determining what routes to accept from each ISP. This can range from merely accepting a default route (used to detect if the link is up or down) to accepting all routes (so called "running defaultless"). The former is usually insufficient, because it does not protect us from an ISP which has an internal failure cutting them off from the rest of the Internet. The latter requires using expensive "carrier−class" routers with lots of memory installed. Fortunately, there are some "in−between" choices.

Rather than using a simple default route, we can use a conditional default route to protect against ISP failure behind the ISP's router that serves our site. A conditional default route is a default route that is defined by a router only if a specific address is already in that router's routing table. Each ISP is only used for a default route if it is advertising one or more routes that indicate it is receiving advertisements from the rest of the Internet. That way, we will always use a default route which promises to be useful.

Another option is to have the ISP send us just its local routes. That way, we can optimize our outbound routing to avoid sending packets that could be locally delivered to the wrong ISP, adding to delivery delays. Care must be taken when using this option, however, because some ISPs have so many local routes that there is no cost benefit in the size of the routers required to handle them compared to running defaultless.

Options can also be combined. In some cases, taking local routes and a conditional default route can provide almost all the availability benefits of running defaultless, while still allowing the use of low

cost routers. As is always the case in networking, a good understanding of the requirements and the available capabilities is essential to maximizing cost−effectiveness.

For More Information

Detailed descriptions and example Cisco configurations for all the approaches discussed here can be found in Chapter 8 of my book, *High Availability Networking with Cisco* from Addison−Wesley, ISBN 0-201-70455-2. But be aware that my book concentrates on the support of redundancy in pursuit of higher availability. Load balancing techniques are covered in Sam Halabi's book *Internet Routing Architectures* from Cisco Press, ISBN 1-578-70233-X, considered by myself and many others to be the definitive guide to BGP theory and practice.
