

Redundant Inbound Connectivity by Name without BGP

by: Brian G. Hughes © 2007

Over the past few years as infrastructure consultants, we have helped several companies install redundant connectivity, either intra-campus, inter-campus or backup internet connectivity. This paper examines maximizing multiple connections to the internet by diverse ISPs.

Most administrators agree that, with the growing importance of inter-networking, a secondary connection for any mission critical service is a good idea. At this time many medium and even some small sized businesses have secured, or are planning to, a secondary connection for backup. Unfortunately, after the initial enthusiasm many administrators don't maximize their secondary link.

It is neither difficult nor expensive to install an appliance to aggregate the bandwidth, or just allow for a failover for address translated clients on the inside of the network, but this does nothing for external inbound connections. One solution to this is to implement Border Gateway Protocol (BGP). Unfortunately, BGP is difficult and expensive to setup and maintain, and you may not have the support of your ISP, which is quite necessary. It may even be impossible for a small or medium sized business to make arrangements to implement BGP.

The premise for this easily implemented plan is that an authoritative DNS server on each of the connections can only be reached by external clients if the connection is alive, and it "points" to services on its own connection. The fail-over portion is handled at the top level internet DNS servers, and the critical settings are at your registrar, likely to be completely independent from you and either of your ISPs. Consequently they are not likely to be affected by any outage you or either of your ISPs may experience.

The following is a work around that can be implemented in house by any administrator. This solution takes a minimum of hardware and is completely out of the hands of your ISP. The only external requirement is that you have the ability to designate authoritative name servers for your domain at your registrar.

In the diagram below we have created a fictitious company using two ISP connections. In this example, your company would receive a small range of static public IP addresses from each ISP. (Note: when I originally sketched out this plan I used a few non-contiguous IP addresses which were in the public routable range. Although these are real IP addresses belonging to the Department of Defense and Qwest, they should be treated as fictitious examples.) Please use addresses assigned to you by your respective ISPs if you implement this plan.

A likely scenario for implementing this plan would be that you are the administrator at a small or medium company that needs to host some services, such as: mail, web, VPN,

terminal services, or any number of other services. You may also host your DNS locally, or currently have DNS hosting at your ISP, registrar, or other third party location.

The basic interconnection diagram is shown below (diagram 1) and can be implemented prior to actively “turning on” the system.

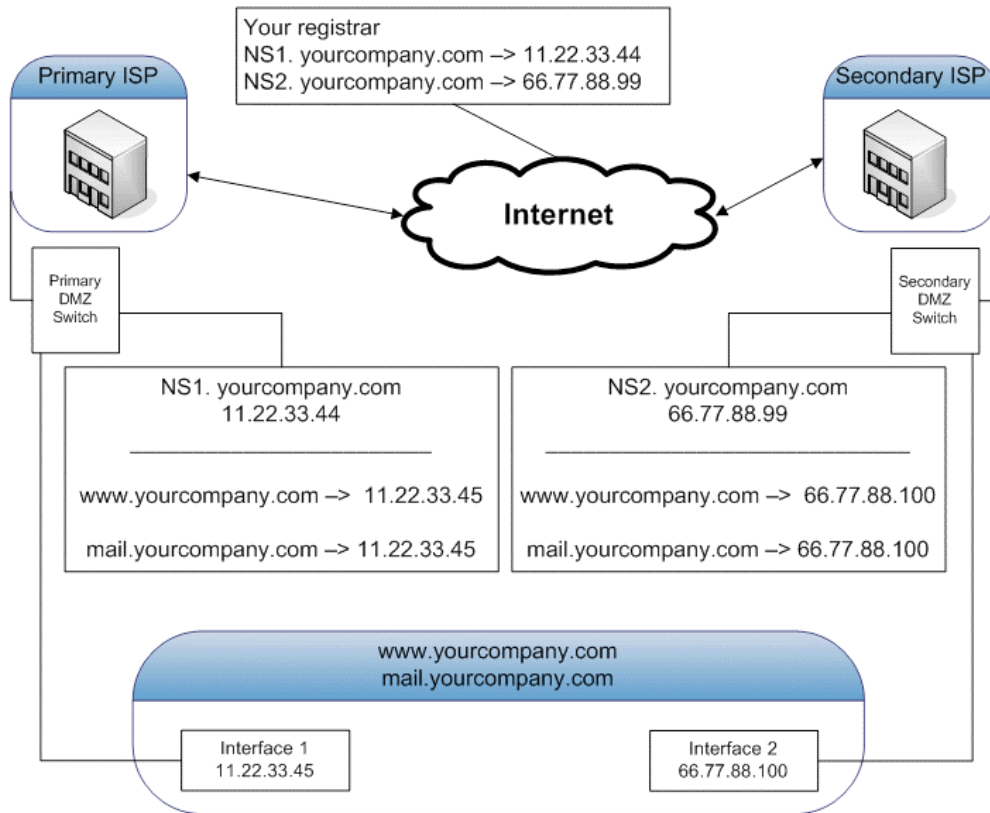


Diagram 1: Functional block diagram for Inbound Redundancy by Name

The interconnections can probably be implemented with little or no hardware, assuming that your company has a few small switches laying around to act as the DMZs if they are not already in place. The DNS servers are such light duty appliances that they can be built on retired desktops if necessary.

The first active step in this plan would be to build and test your own DNS servers. Assuming that you have a platform of choice, I will not advise on how to go about doing this.

Once the servers are built and tested you need only add the interface 2 IP information to your web, mail, or other server and contact your registrar with the change to the authoritative DNS pointers.

There is of course a downside to an implementation of this type, the hierarchal nature of DNS has real propagation delays. Fortunately many networks have TTLs set down to zero, and propagation delays that used to take days now usually takes minutes.

You will notice that the included diagram does not include any outbound or LAN connectivity pieces. These were left out for the sake of keeping the diagram legible. The LAN connectivity could be implemented with a properly configured three-interface router. A 1700 series router from that popular network hardware company, interconnected to the two DMZ switches and performing NAT between them and your LAN could provide outbound internet for your on-site people with consistently improved uptime.

Don't forget to research your ISPs to ensure that they are connected through different tier one providers. Preferably they are multi-homed through different tier one providers.

Appendix A: Uptime Calculations

For arguments sake, lets assume that your small business purchases one inexpensive ADSL circuit and an inexpensive metro wireless connection with the following uptime statistics:

Circuit	Up(%)	Down(%)
ADSL	99.6	0.4
Wireless	99.0	1.0

Your total Down-Time would be represented as the intersection of when both circuits were out of service, calculated as: $0.004 \times 0.01 = 0.00004$ or 0.004%

This equates to a 99.996% uptime, while this ratio is not the coveted "6 nines", it is significantly better than either circuit alone.

Link to original web document:

<http://www.worleyconsulting.com/publications/2007/redundant.html>