

## Providing High Performance Web Services

DOING BUSINESS ON A GLOBAL SCALE REQUIRES A GLOBAL INTERNET PRESENCE

LOCALIZED CONTENT IS A START BUT LOCALIZED CONTENT DELIVERY IS A REQUIREMENT IN TODAY'S WORLD

BY JAMES HOLLINGSHEAD

Back in the mid 90's, when the world wide web was relatively new, just having a website that had information about your company was sufficient. They were pretty expensive to maintain because of bandwidth and server costs involved, so most businesses only had one.

Things are a little different now. With the decrease in the cost of bandwidth and server grade computers, those websites are now used for everything from simply supplying information to e-commerce and the large-scale distribution of content (such as this magazine).

With the changes in the way the Internet is used, the way that websites were handed had to change as well. One set of servers in a single geographic location providing all the global services for your customers is no longer sufficient.

### WHY SHOULD I BE CONCERNED?

Having multiple servers worldwide which hosts your business services has several advantages. The most obvious is redundancy. If one server goes down, customers can still access the remaining servers. This is very important if you use your server for e-commerce or critical customer services because downtime has a direct effect on the bottom line.

There was one very dramatic example of what can happen to a server in one location when Hurricane Katrina made landfall at New Orleans. Many businesses located there had great difficulties in relocating their servers and experienced severe amounts of downtime. One data center there was the notable exception, because they staffed their floor of the building through the entire ordeal, keeping the servers up with a diesel generator and the looters away through other means.

They even kept a weblog during their ordeal to show people what the city was like during the disaster. In fact, they were one of the main, unbiased sources of news for the city of New Orleans for a significant portion of the disaster. However, while

their efforts were extremely impressive (and I admit that I read the blog daily), not every place which hosts servers will go to such lengths to ensure your uptime.

Having multiple servers in different geographic locations allows you to let customers in other parts of the world have faster access to your sites and allows you to better tailor your content to other cultures. This means that someone in Asia would be able to access your site on a server closer to where they are rather than having to deal with the latency of a connection between them and North America (for instance) and will be able to view a site with, to them, a more native feel than the one that your American customers would use. This is very important because we all want to feel comfortable when using on-line services, and most people will not wait more than a few seconds to begin seeing results when accessing your website.

Tied to these first two points is an important third point – sometimes it isn't the server that goes down. While it's a very good start, it's not always enough simply to have servers in different parts of the world. The unfortunate truth is that, on thankfully rare occasions, entire networks can lose connectivity to the outside world. The bad news is that it can sometimes take days to get service fully resorted if something catastrophic happens to the network.

The solution to this is to have at least one of your servers on a different network provider than the others. Don't worry too much about this point right now, because we'll discuss it later.

Having multiple servers also allows you to do something else that can be very useful – load balancing. If one of your servers is experiencing heavy traffic for whatever reason (like we tend to after releasing a new issue), traffic can be routed to one of your other servers which is less busy. This means that, instead of having to wait or make multiple attempts to connect to your site, a customer gets access to your site hosted on a different server

with little delay.

## HOW DO I DO THIS?

There are a couple of ways to obtain co-location for your servers. One way is to set up another office at a separate location and hire staff for it in order to manage your servers. As you might expect, this option is pretty expensive.

The other option is to pay another company for dedicated hosting and co-location services. This saves you the cost of having full time staff in multiple locations, though you still need to buy your own servers. It does however come with a great deal of choice and a few pitfalls to overcome.

Granted, you could also go with shared hosting (where you share the server that you use at the hosting company with some of their other customers) to save even more money, but that comes with a whole other set of problems. The first of these is that it isn't your server. You only have so much space on it and the rest is reserved for the other customers. You also have to hope that what they are using the server for isn't too processor or memory intensive. The upshot of this is, unless you're really strapped for finances, it's usually better to get your own dedicated server.

First, you need to find a company to host your server. When checking out prospective hosting companies, the first thing you want to do is make sure that they're actually up and running. While it may be funny to find hosting companies who can't even keep their own servers on-line, you really don't want to use one for obvious reasons.

Next, try digging around a little on-line and see if you can find any reviews of them. Google is really your friend in this, because seeing the praise and complaints that other customers have had for the company will give you a better idea of what to expect. Remember to take all reviews with a grain of salt, because not all of them are honest, but the chances are if you see nothing but bad reviews for the company, it's best to look elsewhere.

After you've looked around for reviews of your potential hosting company, get in touch with them. Ask them for a test IP address and get it in writing that the address is on the same network as the service that you want. Perform a traceroute on the IP that you've been given to make sure that it doesn't take too many hops to get to.

Now it's time to get in touch with the people that will actually be running the server. Before we even get to the physical requirements, ask them what your hardware options are. Do they host both towers and rack mounted servers? If they host rack mounted servers, do they host (or charge more for hosting) 1U, 2U and 3U servers? Do they have deals with server vendors which may save you money on support in the event of hardware failure?

What operating systems do they support? If they're a windows only shop and you have things which have to run on Linux or BSD, then they're really not an option. Make sure that they do a minimal install on the server. One of the main causes of security breaches is having a server that is running services which you aren't using.

Ask them what their power situation is. See if they have adequate battery backups and an on site generator for extended power outages.

Make sure that their network is in good order. Remember that a network can only transmit data as quickly as its slowest part. If they're running a firewall on a 386 box which all traffic has to go through, you're going to have trouble. Also make sure that the number of hops from your server to the outside world isn't too high. If the traffic just goes through a couple of switches, you're fine. If, on the other hand, to paraphrase a colleague of mine, it goes through half a dozen hubs, a hamster running on a wheel, the aforementioned 386, a hot air balloon, and then another hub or two before reaching your server, there's going to be a problem.

Make sure that the fiber they use to connect to the outside world has multiple access points so their network won't become inaccessible in the event of a backhoe accidentally cutting the cable. I know it sounds unlikely, but it can and does happen, and when it happens, it takes a considerable amount of time to fix.

Be absolutely positive that their hosting facility is climate controlled and has appropriate fire suppression systems. Also make sure that it's secured – preferably guarded and in a building with closed circuit cameras and the servers in cages. Nothing can ruin your day faster than a hard drive that's died from the heat unless it's the server being engulfed in flames or someone just completely walking off with it.

It's also important to ask when they have staff there

in case of emergency reboot and how often they make backups of the server. In addition, be sure to ask if they keep a copy of the backups off site in case of fire or flood.

Finally, check for hidden costs. For example, if they offer the server as 5Mbit/sec but have it on a 100Mbit/sec port with huge overage charges, it might be worth considering finding a different host.

## WHY WOULD I NEED TO BE ON DIFFERENT NETWORKS?

No matter how many precautions are taken, sometimes backbone networks do become inaccessible to other backbone networks. It can be because a backhoe cuts a vital cable or any number of other reasons, but the truth is that it does happen albeit rarely.

The way to prevent this from affecting your business is to have at least one of your co-located servers on a different backbone network than the others. For instance, if all of your servers are on networks supplied by Qwest, it would be a good idea to have one hosted on a network that is supplied by another backbone network like MCI. This means that, even if Qwest disappears off the face of the earth for a few days, your site will still be accessible. This can help prevent a lot of headaches.

## DO I NEED A SEPARATE WEBSITE FOR EVERY SERVER?

The short answer is no, you don't need a separate website for every server. You may, however, want to have separate websites for region or language specific content.

For example, commerce carried out in the United States has a whole different set of taxes, shipping requirements, etc compared to commerce that might be carried out in the European Union. In order to minimize the strain on your customers, it would make sense to set up a separate domain (say, .co.uk) for purchases made in that region.

## HOW DOES ALL THIS WORK?

I realize that, at the moment, how all of this works together is probably a little fuzzy. After all, we're talking about having servers in different parts of the world that, as far as anyone accessing them is concerned, are all the same machine. They don't want to have to worry about choosing a server closer to them, and you don't see how to get around that.

Don't worry. There's help with that too.

Companies such as Nortel ([www.nortel.com](http://www.nortel.com)), F5 ([www.f5.com](http://www.f5.com)), Radware ([www.radware.com](http://www.radware.com)) and Coyotepoint ([www.coyotepoint.com](http://www.coyotepoint.com)) sell networking devices that help take care of this problem for you. Basically, the devices they sell for this purpose are DNS servers that reply with a specific site based on the source of the client. This way, if someone from Asia accesses your site, they get sent to the server closest to them instead of the server in the United States or wherever it is you happen to be.

Their products are fairly simple and painless to get up and running and some of them can even handle the network load balancing that was discussed earlier.

If you want to save yourself the cost of buying an appliance to do the job of making this all work, you can do it yourself with the language and country data in web browsers and open source tools like `mod_rewrite`.

Any way you look at it, unless you're just running a personal site or one that isn't very important to your business, it's a good idea to have it set up in such a way that there is no single point of failure and that it can be accessed quickly from any part of the world in which you might do business. You can easily save more than the cost of the co-location due to increased uptime.

While this hasn't been an overly in-depth look at the problem and its solutions (that would take entire books), I hope it has been a useful overview and has helped point out why and how you should distribute your web presence around the world. If you're still in doubt as to whether or not you should distribute your web presence, honestly consider the loss your company would experience in the event of the catastrophic destruction of your web server verses the fairly small cost of a dedicated co-located setup. You'll probably find that the second option is a lot cheaper.

**James Hollingshead is Executive Editor of o3 magazine. James can be reached via email -- [james@o3magazine.com](mailto:james@o3magazine.com). If you are interested in writing an article, providing feedback or have any comments or suggestions regarding o3 or this article. Please feel free to contact James. If you liked this article or discovered an error with the content, please let James know directly rather than commenting on it privately in your blog.**