

Location, Location, Location: Where Should You Place Your Precious Equipment?

The Final Installment in a 4 Part Series on Data Centers

By [Bruce Whetstone](#), Executive IT Management Consultant with Taos' Office of the CIO Practice

Whether you plan to build your own data center, or source it to a co-lo, "ping-power-pipe" data center or managed services provider, one of the most important questions to answer is, "Where on the planet should you place the equipment?" There are a number of critical things to consider in selecting a location for your company's data center. There will inevitably be issues and trade-offs around the need for global applications performance vs. network latency. You should also consider the effect of natural disasters (seismic and weather-related events) and energy costs in your selection process.

Whether you are part of a very large multi-national company, or even a small company that currently houses all of your computer services under one roof, you need to ask yourself, "Are we in an earthquake zone? Is it volcanically active? Are we in a flood plain? Is the area prone to tropical storms or tornadoes? What about frequent power outages due to freezing rain or blizzards? Are we sitting in the flight path of a major airport?"

The last one is the easiest to answer. A lot of data centers are located near major airports, and not surprisingly. Warehouse-type facilities are generously located near airports, since the airport noise is not a major impact on this type of facility (which makes airports and warehouses good neighbors to each other).

The question always comes up before a company decides to invest millions in a data center, "What if a plane crashes into the data center?" The question is perfectly valid. And the answer is that you had better have a good disaster recovery plan in place. However, the number of plane crashes per year is extremely low compared to the threat of natural disasters, relative to impacts on data centers. And the number of planes that have ever crashed into a data center? Not aware of any, if you exclude the 9/11 attacks on the World Trade Center. (If you know of any occurrences, let me know; I haven't located any news items or data in the crash records directly related to a data center.) It is hardly worth the discussion, BUT, all things equal (including cost), when choosing between two facilities – one in the flight path, the other not – I'm sure you would pick the one that isn't.

It is tough to find a data center location that “has it all”:

- Cheap real estate (if you’re building your own)
- Cheap and reliable power
- Excellent network communications access facilities
- Well-trained pool of low-cost labor to support the facilities
- Proximity to a major transportation hub (to get people and equipment in and out)
- Little or no risk from natural disasters including earthquakes, volcanoes, and weather

If you do find a location that meets all of these criteria, chances are you will find the world’s largest corporations have beaten you to the spot. That’s why Google has a major facility on the Columbia River in Oregon, near the Bonneville hydroelectric facilities ([The Information Factories - WIRED 14.10](#)), and Microsoft, Intuit, Yahoo! and others have all built data centers in Central Washington state ([Data Center Operations Flock to Central Washington](#)). Even these facilities have compromises relative to the transportation facilities and the available labor pool.

If you are late to the party, what advice should you apply in searching for data center locations?

The first consideration is cost (of course). If you are in a low-cost location, park your data close to home. But if you are in a high-cost location, look for a location or facility that offers as many of the factors listed above, according to weighting criteria that you consider to be the most (to least) important (in rank order).

When it comes to U.S. power costs, a handy reference is the Official Energy Statistics from the U.S. Government (available at [Energy Information Administration](#) website).

When evaluating 3rd party data centers (assuming you have decided not to build), look for the data center uptime tier of services you require, but then throw in a standard list of weighted evaluation criteria:

- Cost
- Company viability
- Service and support
- Reliability statistics
- Evaluations from other customers
- Network performance criteria (network access facilities, latency)

If you are a small or medium-size company, with no customer-facing Internet applications, find a low-cost data center provider in a relatively safe location. (Reference: [Natural Disaster Threat Maps](#)). I tend to give preference to data centers in Plains states with good communication hubs (e.g., Dallas/Ft. Worth, Denver, Omaha).

However, if you have high-performance requirements for your applications over the network for either internal company or external customer-facing applications, you have to consider the ‘speed of light over glass’ latency from your workstation (or the customer’s location) to the data center. Even with the best possible routing and least number of relay points in the network path, sheer distance from the U.S. West Coast to Delhi, India, will result in around a 250 millisecond roundtrip delay. More than once I have heard executives that should know better state emphatically that network performance is all about carrier peering relationships and finding the best possible path through the Internet. While there certainly is a kernel of truth to ensuring you have a good network service provider, “exceeding the speed of light” has nothing to do with network peering. The Warp Drive has not been invented yet, and we do not have a Transporter Beam to send the data instantaneously to another country half-way around the globe.

Application performance over a wide-area network (WAN) is directly affected by network performance. Some applications are sensitive to latency only, some are sensitive to network packet loss, and many are sensitive to both occurring simultaneously. Latency is a function of round-trip distance and the delays contributed by ‘handling’ of data packets en route by data communications equipment (e.g., routers and firewalls). Packet loss is a function of data circuit quality – poor quality providers equal poor quality circuits. Good quality providers still have issues from time to time that result in poor quality transmission. That’s when network operations and monitoring systems need to alarm the degraded condition of the network link and flag it for maintenance or repair. This function can be provided by the carrier (managed service), staffed internally or outsourced to a 3rd party (as a managed service).

On another note, relative to applications performance over the WAN, many people confuse bandwidth (how much data can be transferred per second) with latency (how long does it take for one bit of data to get from the source to the destination), and erroneously believe that a “bigger pipe” will solve the latency over distance problem. It won’t. So, if you have latency sensitive applications, you may need to work with a data center provider that has data centers located strategically around the world, so that you and your customers can get reasonably good performance regardless of location. This implies that you may have to replicate web servers, application servers and/or database servers in multiple data centers. (Note: Implementing regional systems (e.g., distributed

databases) can seriously complicate data integration, replication and synchronization, and thereby result in increased support costs.)

The point is: Don't assume that "throwing bandwidth at the problem" of applications performance over the network will make the problem go away. I have very specific examples of large corporations blaming the global network for application performance problems, which required months to troubleshoot, only to discover that the application itself was not engineered to run properly over a wide-area network. The application was regenerating web-enabled drop-down lists for every user inquiry, every time the application was accessed, and as the application database grew, the processing time to create the drop-down list on the fly went up exponentially (to unacceptable levels). Not a network problem.

While it isn't possible to include an entire data center sourcing RFP and RFP evaluation methodology in a short article like this one, there are plenty of examples available on the Internet, and consultants are standing ready to help you through the process.

We certainly hope that you have found this series of articles interesting, informative and entertaining. If we have attained our goal, you should feel much more comfortable with your understanding of key issues affecting data centers, and will be much less likely to embarrass yourself discussing these topics with your peers.

[Bruce Whetstone](#) is an Executive IT Management Consultant with Taos' Office of the CIO Practice. Contact him at bwhetstone@taos.com .