



FEATURE



## How to Choose a Datacenter

[By Scott Whitney, IT Manager, Journyx]

In 2005, notebook computers accounted for 50.1% of all computer sales. In 2006, shelf space for notebooks increased 44%, while desktop shelf space (and sales) went down by 23%. What does this have to do with a datacenter? Everything.

At Journyx, where I manage IT, we presently have about 25 employees. Of those, 11 have laptops issued to them as their primary machines. One employee works remotely in another state. Therefore, half of our employees need constant remote access to our business. Well, it's possible they don't *need* it, but they sure do whine about it an awful lot when they don't have it. So for me, in my little fiefdom known as "IT," that pretty much amounts to the same thing.

As with most companies, we store the bulk of our data internally on our network here at the corporate headquarters, but we also store a fair bit of it at our datacenter. We have software-as-a-service (SaaS) applications which we host for our customers, as well as for ourselves. We have our website, of course, which must be up and running 24/7 or my CEO calls me up in a panic. We have an FTP server for support, as well as one for the public, etc. You get the picture. We've got resources that are needed by our remote employees as well as our customers. In essence, we need a reliable, 24/7, redundant, and fast way for our people and the world to access our data. If this sounds familiar to you, you might be in the same boat that we were in. We needed a datacenter.

I'm oversimplifying our needs a bit, since we are a hosted service provider for literally

hundreds of organizations around the world. You see, with the software that Journyx creates, you can either host it locally on one of your own servers, or you can ask us to do it for you, taking away that overhead. Since we host our customers' data in addition to our own (in different time zones around the world), I was in the joyful, enviable position of evaluating datacenters (again). It was either that or get a root canal, and that was the excuse I used last time, so I decided to man up and take on the challenge.

I say "again" because my previous datacenter experience was a true fiasco. You see, this company — we'll call them "Evil" — had bought up my existing provider, and, in an effort to either cause the 100 or so customers significant pain for no reason whatsoever or to cut costs without evaluating the actual opportunity cost of the move, they decided to close the facility in which we were housed and move us across town to their "better" datacenter. Well, Evil and Evil's Minions had no idea how to run a datacenter. Without going much into their inexperience, let's just say that we knew we needed to move when at 5:30 p.m. on a Friday, one of the Minions *shut down all physical and logical access into and out of the datacenter because several of the collocated customers had a virus*. We were unable to get back up and running until Monday morning. This was

one indication that perhaps there were better choices available to us out there in the world.

Vowing to myself, in my best Roger Daltrey voice, that I wouldn't get fooled again, I put on my Due Diligence Hat (my boss makes me wear it from time to time to avoid situations like the above) and sat down to determine how to choose a datacenter.

Following are the major points which you absolutely cannot ignore if you hope to be successful. I wish I had this article when I was going about my business. Here, I hope to provide, in no particular order, a definitive list of investigation points that should lead you to the best collocation provider for your needs in your area.

### 1) Halt! Who goes there?

With the Sarbanes-Oxley Act of 2002, a lot of attention became focused on fraud and fraud prevention. Part of this particular Enron-created hell is the wonderful and invigorating SAS70 audit which, in the simplest terms, is a proctological exam wherein the external auditors and your internal management pokes and prods and searches around until they can pull sufficient controls out to ensure that customer data is kept relatively safe.

As I mentioned above, we host our own application for our use plus that of paying



FEATURE

customers. It collects time, expense, and travel data for users, and that data gets billed to projects, among other things. For many of our customers, it would be a catastrophe if any of that information was readily available to their competitors. While logical security is, of course, my purview, physical security at a datacenter can play a huge role in satisfying SAS70 requirements, as well as letting you sleep at night. Some things that you might consider for security in your quest for the perfect datacenter are:

- How many cameras does the datacenter have and where are they placed? How is the data recorded and how long is it kept?
- Is everyone who goes into and out of the datacenter required to sign in and sign out?
- Are there two or more specific stop-points on the way into the datacenter?
- Is the datacenter staffed 24 hours a day? Is it staffed with security personnel, and if not, what are the procedures for the onsite staff to deal with security threats?
- Who has access to the logs and videos and what is the procedure to get them?
- Is the datacenter insured against loss due to theft or vandalism, or must you carry your own insurance?

**2) By the power of Greyskull!**

Power is one of the main reasons to go with a datacenter instead of the “host it myself” plan. It is obviously crucial to the proper operations and uptime of systems,

and there’s an awful lot to know about it in order to keep systems up and running 24/7. The network operations center (NOC) manager should be able to speak fluently about power requirements, loads, circuits to your equipment, unimpressive power supply (UPS), power distribution units (PDUs), and so forth.

**Basic power information that you should know about your datacenter**

- **Open transfer versus closed transfer** The first thing to consider is the ATS, or automatic transfer switch. I’m going to assume that any datacenter worth the name has them. The ATS senses when there is an interruption to the city-provided power, and it switches to generator power. When the original power is restored, it switches back to the city grid. Open transfer is a mechanical transfer which provides a (very) small amount of downtime and a (relatively) large amount of electromotive force. Closed transfer, on the other hand, provides no downtime. To be quite honest, which one they have almost certainly won’t affect you, but you should ask the NOC manager anyway just to see if he knows. If not, this would be your first indication that they don’t truly speak “Powerese.” The real questions here are whether it’s a well-known, name-brand ATS, and what does the NOC manager know about it. Can he operate it manually? How often is it tested? And so on.
- **Power exclusive to your rack** When you’re hosting your application in a datacenter, the most common thing to do is to get your own little secure area, known as a “rack” or

a “cabinet.” If it’s less, it can be a half-rack or a third. In any case, it’s important to know how much power comes into your rack. Our cabinets have 20amp circuits by default, and we have more than one full rack. This means I have 48U (units) of space in each rack. Which, in theory, means I can place about 40 single-unit machines in there with a little space leftover, right? Wrong. Each one takes about 1.2amps to run (well, the ones we use, anyway), and that would put us somewhere around 48amps on a 20amp circuit. I figured this little gem out when I plugged one too many machines in and took down our entire cabinet. My boss was not very pleased about this. Afterwards, our datacenter (CoreNAP) quite easily added in a 2<sup>nd</sup> 20amp circuit to our cabinet, thus doubling the amount of equipment I could put in there.

Brian Achten, NOC manager for CoreNAP ([www.corenap.com](http://www.corenap.com)), says specifically:

“NEC (National Electric Code) says you should not run any circuit over 80%. Your 20A 120V circuits should never go above about 16 amps for straight power. In an A&B power situation, all circuits should be run at 40% to allow for one of the two circuits to fail without affecting the device (assuming each of your devices has two power supplies). You should make sure your datacenter provider over-sizes the conductors on all runs. This is not only more efficient, but helps to eliminate the neutral overloading.”

Yeah, sounds like gobbledygook to me, too, but these are the sorts



FEATURE

of questions a good NOC manager should be able to answer off the top of his head.

- **Multiple grids with multiple entry points** Everyone sits on a “city grid.” You’ve seen this in bad action movies. Someone screams, “Shut the whole grid down!” and the city goes dark. Have you ever noticed that they don’t do this during the day? Regardless, having a datacenter that sits on more than one city grid with more than one entry point from each grid is nice, if you can find it.
- **UPS/PDU/monitoring** Between the city power and the generators sit the UPSes, or the unimpressive power supplies. These are basically large batteries that power your systems in the event of “brown-outs” or short interruptions or fluctuations in power. Actually, the UPSes power your systems all the time because, among other things, they contain line conditioners to keep the power steady instead of spiking. A PDU is a power distribution unit, and it does precisely what it says it does. These are arguably the most important parts of a datacenter, since absolutely no power can corrupt data absolutely, especially if your stuff was up and running and writing to disk when the power outage occurred. So what do you need to know about this very involved, complex area of power distribution?

1. What kind of UPSes does the datacenter use? Are they internally redundant?
2. Does the datacenter have a

proactive management plan for these?

3. Are they on contract with the UPS/PDU manufacturers/ service people?
4. What’s the worst power story they have, and how did they handle it?
5. What’s their process for post-mortems on failed devices and customer notification?
6. How often do they test this critical equipment?

All of the above questions also apply to the diesel (or other backup) generator(s) at the datacenter.

**3) An inconvenient truth**

As Al Gore would tell us, climate monitoring and climate change is crucial to the ongoing survival of our equipment. Or species. Or salamanders. Or something. Over the past four years, I’ve replaced the motherboard in one of my home computers six times. It just gets too hot where it sits. This is bad. It also happens to be a personal machine with non-critical data in it that I haven’t actually turned on in about six months now. Of course, that’s mostly because it blew the board again, but the point is that there’s nothing critical on that machine. I cannot afford to be lackadaisical, however, with my customers’ data and business-critical machines, so I have to run my rack somewhere near optimal temperature. Our cabinets at the datacenter each contain roughly 10 machines, and they generate more than a little bit of heat. Not only that, but we are not the only customer in that room. Therefore, there must be an HVAC

system capable of cooling the entire area it controls with all machines running and generating heat. How much is that? Well, that’s a good question to ask the NOC manager along with some others.

- What kind of HVAC system do they have? How many?
- What’s the total tonnage of air that pumps into the room?
- Do they have a hot-aisle/cold-aisle layout?
- How do they monitor them?

The most important question concerns the hot-aisle/cold-aisle layout. Since machines are placed in a rack with front-to-front and back-to-back, the back end of the computers spits out all the heat. That would be a “hot-aisle” where cold air is not pumped. Above would be A/C return air vents which suck out the hot air. The corresponding cold air would preferably be pumped up through the floor in front of your systems so the fans suck in the cold air.

**4) Burn, baby, burn!**

Fire suppression is a pretty important point for any IT organization. If your machines go up in flames, bad things happen to the data that was on them. Correspondingly, while I’ve never tried this, I suspect that putting a running system under a sprinkler head would lead to similar bad things. Many cities, however, have ordinances that require water-based fire suppression. Here in sunny Austin, Texas, for example, you can have a Halon suppression system, but you also *must* have a working water-based system. What this means is that while you can have a “dry-pipe” solution (a fire suppression system which has no water in the pipes



FEATURE

until the sprinklers are activated for that zone), it doesn't really matter, because if the rack next to mine goes up in smoke, my equipment is still going to see a strong chance of showers extending right up until the point that everything shorts itself out. I strongly recommend that you consider all the options in this area. To me, it's far more important to have a good fire *monitoring* system than it is to have an automatic fire *suppression* system that destroys my equipment. Ask the NOC manager about what happens if the guy's equipment two doors down smokes out, or if a monitor on a mobile cart bursts into flames 10 feet away from your rack.

**5) Hey, can I call you sometime?**

This one is the absolute basic due diligence check that every responsible person should perform. Ask the datacenter for references, and take the time to talk to their existing customers, rather than just reading the testimonials on the websites. Call the customers and ask them where they were prior to this datacenter, why they left, and why they like where they are now, if they do. If they do not, listen carefully as to why. Just because a datacenter has all the right technology does not mean that they do business in a fashion that you'd like to join.

**6) Congratulations! You're hired! And you. And you. And you.**

Whenever a company has high employee turnover, there's something going on, and with a datacenter, that can be an enormous red flag. Is management hiring the wrong people? If so, how can you have confidence that the people working for them now are the right people? If people are quitting left and right, why is that? Hiring for expansion is one thing, but employee turnover is another.

**7) It really is size that counts, baby!**

Ok, size and redundancy. And redundancy. Sometimes, a bigger pipe is just what you need. If the water main going into your house were the size of the pipes going to your sprinkler system, it would take you quite some time to take a shower or do a load of laundry. Similarly, if the datacenter has a single 100Mbps link coming in from the outside world, and they have 100 customers each displaying equal and similar usage statistics over a 24-hour period, the math says that their datacenter is not capable of getting you much above the 1Mbps mark, which is one-third the speed of your average T1. One of the main reasons for using a datacenter instead of hosting it yourself has to do with the significant bandwidth advantages provided. However, this also comes with a cost, and this particular part of the equation gets pretty dicey. How can you determine what is best for your needs? I'm going to assume that you have some familiarity with the amount of bandwidth you'll need, so I'll only concentrate on the datacenter aspects of the bandwidth issues.

- What does the datacenter have coming in to them?
- What are the plans they can offer you for your equipment?
- Do they bill on a 95% rule?
- Can they ensure that you'll never go over the bandwidth plan you've chosen?

**8) Sign here, here, and here, and initial here.**

Most datacenters will provide you with an SLA (service level agreement) which will detail out their requirements for your

uptime. Read this carefully. There are only two things I recommend in this area:

- Does the SLA have an "out" clause for non-performance, if you're in a contract situation?
- Does the SLA provide for reimbursement or reduction of charges in the case of non-performance?

**9) Facility**

Though this one is a no-brainer, it is also an often overlooked piece of the puzzle. Years ago, when comparing datacenters, one of the short-list competitors had their datacenter on the 6<sup>th</sup> story of a building. The datacenter overlooked beautiful Austin to the west through large floor-to-ceiling plate glass windows. The view was stunning. I stood there, admiring it, knowing that my machines would be happy having windows where they could look out over the city — right up until about 2 p.m. when the sun came blaring in, turning the whole area into a large oven. Another company which made the short list had their data center in an actual bank vault below the level of the street. While I'm not sure that is necessary, I do think there are a few things you should look for in a facility.

- Moisture detection. How? Where? How is this monitored?
- What's the roof like? Is it a dual-roof with moisture detection?
- Raised flooring?
- Overhead and easily-accessible cabling?
- Make sure it's not located next door to a gun shop, mental facility, or jail.



FEATURE

- Parking. This might surprise you, up until you have to carry a few 100-pound servers six blocks because you didn't bring a dolly.

With all of the above, I was able to pare our choices in Austin down to two facilities that really provided the proper amount of information for me and gave me confidence that they really knew what they were doing. This last point helped me choose between them easily.

**10) Take me to your leader, or I'll atomize your face!**

As many datacenters are privately-held businesses, and often "small" businesses, if the worst happens and I have to move up

the chain, where does it end? In the case of the two providers that made the short list, I spoke with the owners of each business. In one case, the owner sat down with me and chatted. I asked the questions above to assess his technical knowledge, not really listening to the answers, but being blown away by his actual know-how. In the other case, the owner was late to the meeting by 20 minutes, knew nothing technical, and acted as if he were very put-upon by the fact that I was in his building. Guess which one I chose. It's very important to know escalation procedures in the case of emergency, and a company where the owner is willing to take time out of his day to clarify this and is materially and emotionally invested in his business gets my vote each time.

Hopefully, armed with the above knowledge, choosing your datacenter will be a piece of cake. Happy hunting!

If you've got questions or comments, feel free to email me either at [swhitney@journyx.com](mailto:swhitney@journyx.com) or [justscott@gmail.com](mailto:justscott@gmail.com) (the latter having a slightly less aggressive spam filter).

EmploymentCrossing is the largest collection of active jobs in the world.

We continuously monitor the hiring needs of more than 250,000 employers, including virtually every corporation and organization in the United States. We do not charge employers to post their jobs and we aggressively contact and investigate thousands of employers each day to learn of new positions. No one works harder than EmploymentCrossing.

Let EmploymentCrossing go to work for you.