IBM Systems magazine

21st Century Time Shares

Should you find a home for your cloud with a managed service provider?

By Jim Young

Many companies are weighing whether they should take advantage of the benefits cloud computing can offer. But before you jump into it, you may want to review some considerations and best practices.

Some History

In the 1960s and '70s, what we now call cloud computing was called time sharing—though not to be confused with condos at the beach. Back then, computer hardware was so expensive that large scientific computers were rented out to multiple users on the basis of timed access because the average business couldn't afford to buy a high-performance computer for the projects that needed it. Sharing made sense for these specialized assets—kind of like renting a hardwood floor sander for a DIY project. You need to use one, but you surely don't want to own one.

At the birth of the time-sharing industry, one of the barriers to its applicability was the high cost and slow speed of telecommunications lines for computer access. The Internet and PC had not yet been invented. Access by Teletype or typewriter-like devices was slow, operating at a few characters per second, and the applications tended to be scientific and computationally

intensive compared to typical business apps that are more I/O intensive.

In today's telecom environment, high-speed, secure Internet access is ubiquitous and cheap. Nearly everyone has a PC with about the same horsepower as a '60s-era mainframe. And connection speeds are so high and unit costs of computing horsepower so low that virtually every kind of application is a candidate for the cloud. The technical and cost barriers to computer sharing are long gone.

Since the '60s, several developments have made time sharing easier, cheaper, and, in some cases, just plain better than dedicated, on-site computers and staff. For example, while compute and telecom costs have been reduced dramatically over the past four decades, staff costs have risen significantly. And while the typical server environment has undergone generations of sophistication in its operating systems and support software, these resources still don't manage themselves. Frankly, with the virtualization in use today taking advantage of cheap compute horsepower, the human resources needed to manage a server farm and attached storage are not only expensive, but, in some cases, becoming in short supply due to the graying of the talent pool for certain popular operating systems and programming languages. Similarly, the real estate needed to house, power and cool a data center has increased in cost even though the relative space required per unit of computing power has decreased. In some urban areas, data center growth is constrained by available electricity.

Cloud Is Hot

Cloud deployment makes sense for most organizations to at least consider, especially if constraints in data center space and human resources exist. But what other barriers—business, legal or otherwise—stand in the way of deployment? To learn more about the perspective of cloud computing providers, I spoke with several executives whose companies provide solutions across the spectrum of the submarket, spanning disaster recovery (DR), high availability (HA), managed hosting, infrastructure as a service, platform as a service and co-location.

My first impression is that due to the newness of the subject and the buzz surrounding cloud computing in general, this market is scalding hot right now. Hardware resellers who have made their livings off selling IBM Power Systems* server equipment and storage have watched their sales volumes diminish due to the incredible price and performance improvements in the Power* lineup over the past few years. For some, the cloud's allure is that it offers a way for resellers to

reinvent themselves into managed service providers (MSPs), thus turning sporadic transactional profits from selling servers into a predictable, long-term cash flow from regular monthly service fees. Meanwhile, the IBM i segment of the market is experiencing a spike in retirements of senior technical and IT management staff, with sparse bench strength to pick from as replacements. A competent MSP offers a way to retain applications that are doing a good job of running the business and help you worry less about staffing for the future. End-user staff can focus on helping improve the business, its competitiveness and agility, and let the MSP staff worry about operations, backup, DR, HA, storage management, power and cooling, server performance, network performance and security.

Considerations About MSPs

Partnering with an MSP may be a good way to go, but in addition to staffing, you'll want to be aware of several other considerations.

You may find comfort—justified or not—in a data center run by professionals who manage data centers for others for a living. The MSP execs I spoke with said that on the way to closing a deal, a data center visit was part of every sale. The companies relinquishing their on-site hardware want to kick the tires of the replacement, physically. Part of that evaluation process should include not only a site visit, but also verification of a provider's financials; proof of insurance, audits and certifications of data center security and processes; DR/HA provisions; license agreements; standby power arrangements; flood zone and elevation certificates; and proof of past DR/HA event outcomes.

Personally, I'd want to know about plan B—e.g., things don't work out as promised or planned and the services provider relationship must be unwound. What exactly will the vendor do to return your data and applications to you? As complex as a divorce is, this scenario is somewhat similar. Having a pre-nup makes a divorce simpler.

Another issue is dynamic LPAR provisioning, as supported by Power Systems servers running IBM i and AIX*. If your application workload spikes and utilization peaks higher than what's contractually provided and configured, what defines the fees for such a spike? Is it average utilization or average daily peak utilization? What is the average period, week, month or quarter?

Another important factor is the language of the service-level agreement (SLA) and its reporting formats, frequency, etc. SLA is the baseline for monitoring performance of the arrangement and, in most cases, will play a pivotal part in determining future fees and billing.

The economics may be your No. 1 motivator but the old axiom "Beware of bargains in brain surgery" can still ring true.

Cases of licensing issues with ISV suppliers of applications/tools and IBM as the provider of operating systems and licensed program products have been reported. It's important for prospective MSP customers to know their rights for each software product being used with respect to re-hosting the software on a CPU serial number other than the one originally specified in the license agreement. Some licenses are transferrable; some are not. Some licenses charge maintenance based on P groups (IBM-defined processor groups), others by cores, Commercial Processing Workload (CPW) rating of the machine serial number licensed, or even sometimes based on connected, concurrent users of the software. How those terms and conditions transfer and translate to a hosting environment will depend on scores of factors including who actually owns or controls the hardware and, in some cases, its physical location. A clear trend in the hosting business is the practice of re-hosting applications on smaller-rated, dedicated servers. The goal there is to reduce the software maintenance expense for application, OS and tools, since the maintenance fees for many of these products are priced based on either the P group or the CPU rating of the hardware hosting the software. This strategy has the additional benefit of reducing hardware maintenance over time, both by enjoying the warranty of new hardware and the reduction in hardware maintenance once the warranty period expires.

Of course, many of these tactics can be employed by end-users who want to reduce their operating expenses and still retain their own hardware. In many cases, decisions to go with an MSP are based more heavily on subjective factors than just a winning or losing economic proposal.

In the case of a graying staff, worries about the retirement of key employees and their lack of available replacements can be valid concerns. Just imagine—in the absence of the IT director, who is going to plan for the next generation of hardware or software? Who's going to apply PTFs? Who's going to make sure the computer room has enough power and cooling to handle new equipment and workloads? And who's going to keep current with trends in the industry to advise executives on future probable issues or concerns? In such an example, a relationship with a service provider would answer most, if not all, of these

concerns as well as provide at least a sense of backup or bench strength if needed. A services company can address these concerns with some convincing arguments. For instance, most providers have several people on staff capable of providing these functions on a time-shared basis. Nobody needs a full-time PTF or HA expert. But when it comes time to role-swap the host with the target, you surely want an expert—and your service provider should be able to convince you it has them in depth.

Four Best Practices

The best advice I received in discussing these issues with several MSPs boiled down to four basics:

1 Take time to evaluate all of the business and risk factors. Check references, look at financials and treat the task with all of the gravity a long-term contract for any other aspect of the business would be treated.

2 Put a toe in the water. Start with one task, like DR or HA, and try to negotiate a short-term, renewable contract. You don't have to negotiate all of the details in an all-encompassing contract for your first deal. Try for a contractual format so both firms can get to know each other with minimal business risk to either party. After all, for these deals to work out, both parties must enjoy benefits.

3 Remember that when you delegate a big responsibility, for the expectation of success to be valid, you must also delegate enough authority for the MSP to do its job. You're paying for expertise and for the MSP's capability to apply that expertise to solve your problems for you—ideally, without you ever being informed of a problem.

4 The economics may be your No. 1 motivator but the old axiom "Beware of bargains in brain surgery" can still ring true. Running your data center is still a critical element of enterprise success. If the deal seems too good to be true, it probably is. The happiest MSP clients enjoy a worthwhile long-term savings but go into the deal for more reasons than economics. The money, as the saying goes, should be the frosting, not the cake itself. P

Jim Young is the VP of Sales for Midrange Performance Group, providers of performance management and capacity planning software and services.