

## Cloud 'Recovery' or Just The Same Old Thing?

*This post is part of our ReadWriteCloud channel, which is dedicated to covering virtualization and cloud computing. The channel is sponsored by [Intel and VMware](#). As you're planning your Cloud Architecture, check out this helpful resource from our sponsors: [IBM Credits Virtualization With Helping Client Contain Server Sprawl](#).*

Cloud computing means many things, but almost all definitions include some key value propositions: scalable on-demand resources, a metered pay-per-use model, access over the Internet, and infrastructure management and optimization that is better than most data centers.

At a more conceptual level, cloud computing abstracts away all the undifferentiated IT tasks. Most businesses don't add any value to their customers or create any competitive advantage for themselves when they buy, build, configure, and manage servers and storage. This is doubly true for disaster recovery equipment and data centers.

Conversely, poor performance in these tasks can cost value and competitive advantage. There is no benefit in doing these tasks well, but there is cost to doing them badly. This is like the opposite of a financial call option - lots of downside risk, but no upside.

For companies planning their first disaster recovery data center, with the associated selection, build, and maintenance tasks for servers, storage, and networking, cloud computing seems like an obvious fit. They can trade the capital expense that buys them no new value, for a no-commitment operating expense that probably buys better operating practices than they could achieve themselves.

Solutions are beginning to grow up around this idea of cloud recovery. The name is a little optimistic because most offerings today are traditional backup solutions, with little or no ability to actually recover in the cloud. Although a lot of vendors in the backup industry are making cloud announcements, they are mostly just letting users store backups in the cloud. In order to really deserve the cloud recovery title a solution should have the following features.

1. **The ability to recover workloads in the cloud:** The cloud can offer more than just a place to dump your backup files. It can provide the computing systems to run your recovered systems, and after a production system fails, the ability to quickly

restart a complete replacement with data, applications, and complete configuration in the cloud.

2. **Effectively unlimited scalability with little or no up-front provisioning:** A few vendors can offer rapid, off-site recovery, but they don't really qualify for the cloud title unless they provide lots of stand-by capacity with no up-front reservations or configuration. While this seems like a lot to ask, this is the promise of cloud computing.
3. **Pay-per-use billing model:** A defining characteristic of cloud computing is that we only pay for the things we use. Use a little this week and pay a little; if we use a lot next week then we pay more, but only for that specific week.
4. **Infrastructure that is more secure and more reliable than the one you would build yourself:** When we decide to outsource any part of our operation, we worry about the security and reliability of our vendor. The best cloud providers have not only large scale equipment, but also large scale expertise. This means that they can be much better at security and reliability than any of their customers, and their data center is better than one we might build for ourselves.
5. **Complete protection and automated recovery:** Non-expert users should be able to recover everything they need by default - the current crop of solutions is shockingly bad at this. This is the cherry on top because it makes everything so easy. No one wants to go through a "run book" full of recovery procedures and bring in experts for each system to assist with getting recovered systems back online. Depending on the type of disaster, experts may be scarce, and the run book is probably out of date. Why not make the run book part of the automated system? Instead, simply push a button that says "recover now", wait for the files to copy, and then log in to the perfectly configured system, running right in the cloud.

In summary, the trouble with traditional backup solutions is that they are really focused on dumping the data onto tape or disk (and now onto cloud storage), and maybe restoring it back onto the original hardware. If you must recover to different hardware or a different virtualization platform, they don't generally do much to help with the inevitable incompatibilities. And for equipment failures, you really have to have some hardware standing by.

So while cloud computing means many things, it is fair to say that any cloud recovery customer who doesn't get all five of these features will be disappointed.