Clouds from Both Sides

Steven M. Bellovin
Columbia University

"Cloud computing" is the buzzword *du jour*. Everyone is either doing it or wants to; some technophiles even liken it to the Industrial Revolution. (Admittedly, those of us of a certain technical age have to squint to differentiate it from the time-sharing service bureaus of the 1960s.) But is the cloud "secure"?

The question as I've just phrased it is unanswerable because neither "secure" nor (especially) "cloud" have rigorous definitions. We also have to ask the first question in any security dialog: "What are you trying to protect against whom?" And one more question that's asked all too infrequently: "Secure compared to which alternatives?" This last question is often the most interesting.

Intuitively, the cloud can provide computing cycles (for example, Amazon's EC2) and/or remote storage; the latter can be just for the owner, or it can permit sharing. For "security," we can use the usual trio of confidentiality, integrity, and availability. Our questions, then, change: For remote storage and computing, does the cloud provide more or less confidentiality, integrity, and availability across a wide spectrum of attackers?

Availability is probably the easiest to answer. I assert that despite occasional well-publicized failures, a professionally run cloud service is more available than a typical in-house solution. The cloud service can run more redundant resources to resolve outages, whether malicious or accidental in origin. A good cloud service will use RAID disks and back them up. Put it this way: How recent are your backups? When did you last test your ability to recover from a disk crash? Do you have more servers than Amazon? Do you have more bandwidth than Google? Yes, a failure at a large provider will affect more users; conversely, we hear about such failures more than we hear about the routine (and frequent) outages at typical corporations. The difficult issue is whether your enterprise can function if all of it is cloud-resident. Diversity is always a good thing, but should you seek it in your own environment or via different cloud providers?

Integrity and confidentiality are somewhat harder to assess. Most (though, of course, not all) penetrations result from exploitation of holes for which patches are already available. Is your own in-house staff conscientious about installing all available fixes? Are your systems properly configured, especially for sharing data? Would a service provider do better? These questions aren't easy to answer. If the reason for a delay in patching is lack of resources, the cloud provider is likely to be better. On the other hand, many enterprises delay until they can assess the compatibility of their own, in-house applications with the new system—and cloud providers have many applications to worry about. Sharing resources with outsiders is almost certainly better done via the cloud because the cloud provider's access control mechanisms are tuned for that sort of scenario, and they've dealt with the underlying platforms' complexity.

In-house computing probably has the edge when considering possible attackers. Apart from a provider's own employees turning to the dark side, you run the risk of being collateral damage when some other customer is targeted. There are also legal issues to consider: under US law, at least, you arguably have less protection against "subpoena attacks" when your data isn't stored in-house.

I don't claim that the answer to cloud computing is simple. But I do assert that running your own systems isn't inherently better, even from a security perspective. You need to do a detailed assessment for your own particular situation.

Steven M. Bellovin *is a professor of computer science at Columbia University. Contact him via www.cs.columbia.edu/~smb.*