A Winthrop & Weinstine blog dedicated to bridging the gap between legal & marketing types.

Cloud Computing Residue

October 28, 2011 by Brad Walz

One of the often touted benefits of cloud computing is the ability to scale. Scale up when demand increases and scale down when demand recedes. Scalability is about doing what you do in a bigger way. It is all about allowing more people to use your application. Generally, when people refer to scalability in the cloud context they are referring to vertical scalability. "Vertical Scalability" involves adding resources within the same logical unit to increase capacity. An example of this would be to add a server. What is happening when you scale up is that more servers are made available for your use and then taken away when you no longer need them. And while scalability is truly an advantage, there is little discussion about the residue that results from this benefit.

When you scale down, data is left on the server that is no longer in use. What happens to this data? I have not encountered a service level agreement from a cloud provider that addresses this issue. Arguably, the data belongs to the cloud provider since it is the owner of the server. However, this data should remain the property of the customer. As the customer, you should negotiate in the service level agreement that when you scale down the cloud provider will send you any data residing on the servers that are no longer in use and then reformat the hard drive and re-install the operating system to make sure that all of your data is wiped from the server