The Virtualization Illusion:
The Smoke and Mirrors of Cloud Computing

The one common concept that binds together all the various technologies, ideas, etc. behind cloud computing (and anything related to it) is the illusion that a single host computer can be split into many smaller and (as a group) faster computers. The maintaining and enhancing of this illusion with systems built on top of it is the heart and soul of cloud compting. This white paper examines all the major ingredients of this illusion and how they help maintain and/or enhance the illusion as well as the results of their failure to maintain their part of the illusion.

The main components of any cloud computing platform are (all optional components have been removed):

* Virtual computers
* Virtual networking
* Virtual storage
* The control and monitoring systems needed for a potentially large number of hosts to be deployed
* An API that allows the above to be accessed in a automated fashion

The hypervisor is the program that provides the actual virtual computers. In general, the other services are simply “plugged into” the hypervisor (see forth coming white papers for details on how this done). Once the proper services are connected the hypervisor then, with a predetermined number of CPU's and RAM, is started.

The disk (and other storage) needed for the instances is typically provided by subdividing a physical disk into different partitions (either maintained by the hardware or the host operating system).

The networking is provided by a combination of “on host” and “off host” cofigurations that guarantees (1) that all traffic to and from outside networks (such as the Internet) is sent to a instance and not the host and (2) that *basic* security measures have been taken to ensure that no really nasty surprises happen. At the simplest level this is nothing more then auto-assignment of addresses all the way up to the automated install, maintenance, and fault detection/recovery subsystem of a large data center.