

ESX Hosting Primer for SMB

By Gil Vidals / CEO VM Racks

Read this document for an overview of ESX server virtualization. Even if you are not leaning toward a particular virtualization technology, this information is practical in educating you on how to select a virtual server for your business or organization.

On the average, one can expect to **save approximately 25%** on a virtual server compared to a fully dedicated server. The cost savings is a result of the primary benefit of virtualization, which is resource sharing. The customer benefits by cost savings, and the web hosting company benefits by reducing power consumption and other overhead.



A virtual server is for all practical purposes the same as a physical server. As you can see in the funnel diagram, a physical server is "cut up" into several virtual servers. Each virtual server has access to the physical server's resources including the CPU, memory (RAM), and hard drives. The advantage for the hosting company is significant. In order to understand the positive impact virtualization technology provides, it is

important to note that studies have shown that the average server is idle about 80% of the time.

If a typical server is idling most of the time, then the web hosting company is paying for electricity and rack space for a machine that is unproductive. To drive the point home, suppose you purchased a vacation home in the country side for your family. You use the home only four weeks out of the year, but you still have to pay the mortgage each month. In order to make the vacation rental financially palatable, you decide to rent your vacation home out to friends, family and other guests. The idea is to rent the vacation home during the long stretches of time when it is sitting empty or "idle".



Virtualization uses the same concept as a vacation rental. Since the average business server, known as the *host*, is idle most of the time, why not allow several clients to take advantage of the idle time. And in fact, this is exactly what a web hosting company does. Invite guests and let them use the idle resources of the server. Each client is given their own guest machine, known as the *Guest OS* or virtual machine *VM*, in virtualization parlance. The *host* has several *VMs* piggy backing on it. As a vacation rental, the idea is to utilize the idle time of the host, but being careful, not to overtax it by bringing on too many guests. You wouldn't want to show up at the vacation rental to find your mother-in-law still there, would you? The web hosting company also has to be careful in balancing the number of *VMs* with the available "free" time of the host. Adding too many will cause the *VMs* to fight for limited resources and degrade their performance, while adding too few

means the host is underutilized.

If you followed the analogy of the vacation rental, you should now have an intuitive understanding of what the virtualization hype is about. It's important to point out that just like the guest who pays a smaller fee compared to the owner of the vacation home, the renter of the Guest OS will pay less than having purchased or leased a fully dedicated server. The cost savings vary, but is **approximately 25%**.

Don't be fooled when you see an ad for a virtual server that says it only costs \$29.99. This would be like you renting a vacation home, but the landlord restricts you to the bathroom during your stay. For small fee, perhaps \$29.99, you are going to be restricted in utilizing minimal resources of the host. Read the fine print, you may only get 5 GB of disk space and 100MB of RAM for example. These resources are just too restrictive to be effective.



At this point you understand that your virtual machine will be fully functional just like a physical server, but you pay less. Continuing the analogy of a vacation rental, you know that vacation homes are different depending on where they are located and the amenities they include. Some vacation rentals are on the slopes of Vale and others on the beaches of the Bahamas. Some vacation rentals are lavish and extremely expensive and some are simple and cheaper. Virtualization technology also varies. Some virtualization technology is considered enterprise class while others are for the smaller business man or woman. How does one decide which virtualization technology is right for their particular situation?

Currently there are dozens of virtualization solutions. Of these, three stand out as the leaders thus far in this relatively new industry:

- VMware ESX (market leader)
- Microsoft Hyper-V
- Citrix XEN

Because virtualization is complex it is best to stay with a market leader that is tried and true. Select a smaller player whose virtualization is immature and unproven and you could find yourself in a pickle with a server that isn't functioning properly or worse yet is down frequently.


SECURITY

A virtualized server like any other computer must be secured or protected from outsiders. Typically when thinking of security, *firewalls* come to mind. A properly configured firewall is the first level defense. A security expert would chime in and say that the first level of defense is actually to "harden" the server. *Server hardening* refers to securing ports on the server – that is to close up unnecessary access. Most servers need access to the internet, but what about bluetooth or printing? Since bluetooth and printing are usually turned on by default, the systems administrator (*sysadmin*) must configure the settings so that these utilities are not turned on when the server is rebooted, thus eliminating an opening that a hacker could use to gain unauthorized access to the server.

However, even server hardening and firewalls are not enough. Firewalls are like a sieve. The holes in the sieve allow access to the server.

Typical *pinholes*, as they are called in security lingo,

provide access to the server through ports for



web (port 80), email (port 25) and FTP (ports 20 & 21). Hackers are aware of these pinholes and have found ways to sneak in. Another level of security is needed to examine the traffic that is flowing through the pinholes. Since most attacks happen over the web, a *Web Application Firewall (WAF)* is a primary line of defense, especially for servers whose core function is serving web pages.

A Windows server will need anti-virus software installed, but typically Linux servers are relatively impervious to viruses, so no anti-virus software is typically installed.

MINIMAL CONFIGURATION

Briefly stated, a Linux server should start off with at least 512 MB of RAM, 1 CPU, and 25 GB of hard drive space with 500 kbps of bandwidth (0.5 Mbps). Bandwidth can be tricky to calculate because it can be measured in kilobits per second or megabits per second or even gigabytes per month. Use BWCalc located at <http://www.valkaryn.net/bwcalc/bwcalc.cgi> to help you compare different web hosting companies. If one company sells bandwidth in megabits per second and the other sells it in gigabytes per month, this tool will allow you to compare the two side by side. It even provides a way to include the cost per unit quoted by the web hosting company, so you can easily compare across disparate bandwidth values on a dollar basis.

A windows server will typically require more RAM and CPU, so a starting point would be 1 GB of RAM and 2 CPUs.

These settings are simply a minimum starting point. A more custom tailored virtual machine requires that you list the applications the server will be running.

MIGRATION

Once you select a web hosting company as a home for your virtual machines, then the next

step is to have the hosting company create a server from scratch or to migrate your existing server(s). Naturally setting up a new server is much easier than migrating a production server.

VMware provides a “physical to virtual” tool (P2V) that converts a physical server running Microsoft Windows into a virtual server. The virtual server is comprised of several files.

BACKUPS

Hard drives will fail within approximately 1-3 years. When the hard drive fails, how long can you afford to be down? If the answer is less than a day, then you need to understand the differences between the various backup methods and select wisely.

Backs ups provide a means to recover a file that was accidentally overwritten or even an entire directory that was wiped out unintentionally. However what happens when the hard drive fails completely? Is it possible to efficiently recover the entire system?

The answer depends on what time of back up system is used. A typical back up system operates at the file level and is good for restoring files, but is not well-suited for recovering entire hard disks. For recovering failed drives a *bare metal recovery (BMR)* system is recommended. BMR operates at a level beneath the files, capturing disk partitions instead of files. Essentially this means a snapshot image of the hard drive is taken and all that is needed to restore the drive is to copy the saved image to a new drive.

MISSION CRITICAL

If you are running mission critical applications, then your server must be available round the clock. High Availability (HA) and Distributed Resource Scheduling (DRS) are sophisticated methods of monitoring a server, and if it becomes unavailable, another server is started in

its place. Normally having a standby server is an expensive proposition; however, with virtualization, HA and DRS can be provided at a lower cost.

SUMMARY

A virtual server is the same as a dedicated server and requires all the same precautions and maintenance than a physical server. A cost savings of about 25% is realized with a virtual server since many of the same resources are

shared, including power and disk space. A virtual server will only consume about 3% of its resources in overhead, so the performance hit is minimal.

High availability and Distributed Resource scheduling are two enterprise level services that virtualization provides at a substantial cost savings, but is only important for mission critical services.