

## Practical virtualization with ESX Server

# VIRTUAL BENEFITS



Improved resource balancing, central administration, and service consolidation – you can only win with virtualization. Linux Magazine visits a communal data center for a practical look at virtualization in the workplace. **BY CHARLY KÜHNAST AND JÜRGEN BACKES**

**W**e call them “electric heaters”: servers that are indispensable for daily operations but never show a system load of more than 10 percent. They could be DHCP, DNS, print, file and fax servers, and the occasional, venerable application server that only three people need – but you don’t dare switch them off. Add the inevitable test systems that nobody dares touch. The hardware is typically outdated, but why would you want to buy something newer if the old machines are not working at anything like full capacity?

At the same time, these machines give you that sinking feeling: one thing is for sure, old hardware is bound to crash some time. Just imagine the chaos if the DHCP server goes down on a Monday

morning; this is when the load of just a few percent would really start to tell.

This is what convinced us to look into virtualization not so long ago. Both of us, my colleague Jürgen Backes and I, Charly Kühnast, work as administrators at a communal data center that provides services to district and municipal offices, as well as official institutions supporting about 11,000 desktops.

When we reached our decision, there weren’t as many competitive virtualization technologies as there are today, and some of them did not give us the kind of flexibility we needed. On the other hand, we had gained some positive first-hand experience with VMware; a fact that drew our attention to the other members of the VMware family. It quickly became

clear that the ESX Server would not give us the performance we needed, and this pushed our decision in favor of the VMware ESX Server.

## Layered Model

The ESX Server’s functional design is pretty much the same as that of the well-known VMware Workstation, although the server also provides the operating system. In both cases, a virtualization layer separates the host system from the guest systems. The guests are presented with an isolated and idealized hardware emulation. Each guest system thinks it has a CPU of its own, or even two if the host system has the virtualSMP extension, which is sold as an add-on product. What really happens is that all of these logical CPUs are mapped to the physical CPUs in the virtualization layer. Main memory is allocated by VMware and is either shared by multiple guest systems or even swapped out. Of course, the guest systems remain blissfully unaware of this, as the whole process is completely transparent to them.

The guest systems also see hard disks and network adapters as their own physical resources, with each virtual network card having its own MAC address. The VMware Tools, which you may be familiar with from VMware Workstation, provide the required drivers for the graphics card, NICs, and so on.

## Hardware Updates Required

To run an ESX Server, you need a system with at least 2 CPUs – or a CPU with two cores at the very least. You can never have enough memory or hard disk space. Two NICs are also a requirement, as one is reserved for ESX Server management, whereas the guest systems share the other card.

Older documentation or forums tell you that an old 10Mbps card is perfectly okay as the primary network card, as it is only required for server management. But this is a bottleneck you will probably want to avoid, as you need this interface to back up images of guest systems; so



**Figure 1: Central monitoring of the resource consumption for multi-parallel virtual instances.**

go for the best hardware your budget will give you.

For our first experiments, we used 2 IBM x Series (x445) machines with eight 2.8GHz Xeon CPUs, and 24GBytes RAM apiece, with a redundant SAN connector via two fibre channel cards. The cards were detected when we installed the ESX Server, and they were automatically configured to fail over.

No sooner had we gotten the system running than our colleagues, bosses, and customers all acquired a taste for it, and this meant that the job quickly became too much for the hardware. The new hardware comprised three HP Proliant DL580 G2 machines, with four 2.0 GHz Xeon CPUs, and 16GBytes RAM each. Again, these systems had two QLOGIC SAN adapters. At the next stage, we added three HP Proliant DL 740 G1 machines with eight 2.70GHz CPUs, 32GBytes RAM, and two SAN adapters. Our most recent addition was a HP Proli-

ant DL 585 G1 with four 2.2 GHz Opteron DualCore CPUs.

This turned out not to be a good idea, as it proved difficult to migrate the guest systems between the various CPU architectures no matter what kind of tricks we tried. On the upside, the Opteron showed surprisingly good performance.

Besides the SAN, all the machines have two mirrored local disks for the virtual systems to store the operating systems, any ISO images we need, and other installation components. The reason for this was both cost (SAN space is expensive), and the fact that the 1.5.x version that was available at the time did not support booting from SAN. This is not an issue with the current version.

In the meantime, our server farm has grown to include 9 production servers (with a total of 60 CPUs) that host over 200 virtualized guest systems. The guest systems are Windows XP, Windows 2000 Server, Windows 2003 Server, Suse SLES 8, Suse SLES 9, Suse Linux 9.x, Fedora Core 2, and Debian. The last two are not officially supported but work fine.

### Savings All Round

The savings achieved by using virtualization solutions result from improved hardware utilization (especially of CPU

and RAM) and from port cost savings (LAN, SAN connections), which are very often overlooked. In addition to this, virtualization gives us the ability to react quickly to growing demand for more systems (assuming scalable hardware). In a production environment, it makes a big difference to be able to have an urgently needed system up and running in a couple of hours thanks to virtualization, rather than waiting for a delivery from your hardware dealer. To speed up this process, you can either clone systems – this is a good idea for Linux guest systems – or use Windows unattended installation to set up a Windows system with just a few clicks. This saves a lot of time and money in all departments.

And let’s not forget VMware’s ability to create backup images of guest systems automatically, either offline or online, without the need to invest in an expensive third-party backup software. After virtualizing a server, you remove the need for future hardware migrations for all time.

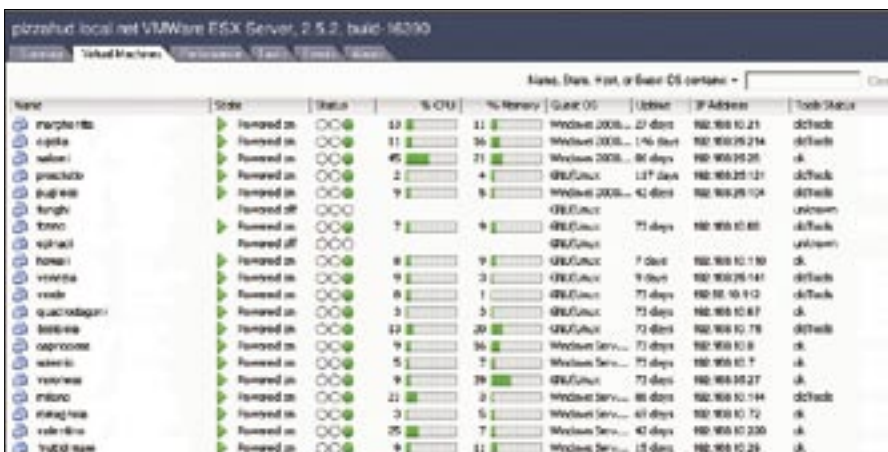
Even if you do need to replace the host hardware some time in the future, this change has no effect on the guest systems, thanks to the virtualization layer. Instead, you just move the guest systems onto the new hardware. And if you use VirtualCenter and VMotion for this process, you can even make the move on the fly!

### Virtual Center

VirtualCenter is an administrative interface for VMware GSX and ESX that gives you an tree view of all your server and guest systems, along with status and resource utilization data. A simple scheduler supports automating various tasks, including cloning guest systems.

VirtualCenter thus provides rudimentary system management functionality. Larger installations will find it hard to do without VirtualCenter, which is a precondition for using Vmotion to migrate guest systems to new host hardware on the fly.

ESX Server supports nothing like the range of virtual operating systems supported by GSX Server or VMware Workstation, but if you only need to virtualize Windows servers, this will not be a disadvantage. Linux derivatives are restricted to commercial distributions by RedHat and Suse. ■



**Figure 2: An overview of all the virtual Pizza boxes on this host. It is easy to identify failures and runaway systems.**