

## Choosing tools for effective virtualization

# GOING VIRTUAL

Good tools are half the battle – even if you are just managing virtual machines. This month we take a practical look at virtualization, and we show you a new threat to watch for in the virtual future.

BY JENS-CHRISTOPH BRENDEL, HENNING SPRANG, AND JÜRGEN QUADE

**S**ervers are not human. They don't live and breathe. They just consume power and take up space. Do we really need so many? The virtualization revolution is about saving money, time, and floor space. Today's virtualization tools provide an efficient environment for testing, running, and managing applications – with lower electric bills and fewer hardware headaches. But is virtualization all good, or does it also open the door to new kinds of threats?

In this month's cover story, we examine the practical side of virtualization. We start with an introduction to some of the virtualization tools available for Linux. Then we take a closer look at a pair of popular open source virtualization alternatives: Xen and VirtualBox. Finally, we settle in for a look at the dark side of virtualization: the mysterious world of virtualizing rootkits.

## Practical Virtualization

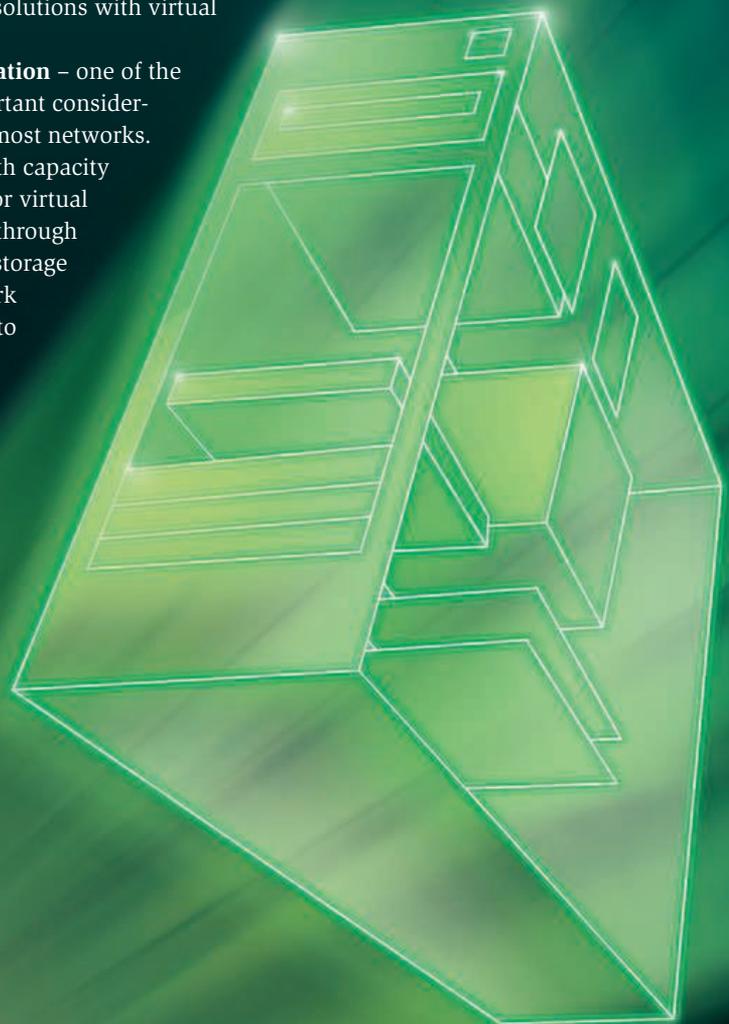
The virtualization paradigm has come down to Earth, leaving the lofty heights of Mount Olympus for real-world concerns like stability, performance, and ease of management. A virtualization system that wants to fulfill all of these requirements must be ready for:

- **Provisioning and deployment** – creating the required number of virtual instances as quickly as possible and with minimal need for manual configuration.
- **Migration** – converting physical machines to virtual machines – and vice versa – if necessary. The ideal is on-line conversion across the network, requiring as little personal intervention as possible. Equally important is ensuring freedom of movement for virtual machines during operations. Live migration is the basis of load balancing solutions with virtual machines.
- **Administration** – one of the most important considerations for most networks. It starts with capacity planning for virtual instances, through staging of storage and network resources, to countless settings

for virtual hardware, physical hardware, users, and privileges.

- **Workload management** – handling multiple virtual machines or physical hosts and replacing instances in case of failure.
- **Monitoring** – providing detailed, real-time information on the status and resource usage of individual virtual machines.

The best solution for your network depends on your needs and your budget. Table 1 introduces some of the more



## COVER STORY

Virtualizing with Xen ..... 30

Networking with VirtualBox ..... 34

Virtualizing Rootkits ..... 39

popular virtualization options for Linux environments.

## Management Tools

Many distributors have gone to great lengths to facilitate virtualization for their customers. Both the Novell and Red Hat enterprise distributions integrate a tool known as Virtual Machine Manager (or Virt-Manager). On SUSE, the tool is integrated with YaST. Virt-Manager gives users the ability to set up a Xen instance in a couple of simple steps. Besides Linux (Novell, Red Hat), Solaris 8-10, and Netware 6, the tool also works with various Windows versions – provided the CPU has the required virtualization support.

Red Hat Enterprise Linux 5.1 also uses Virt-Manager for managing virtual machines. Red Hat also offers a tool on top of Enterprise Server, called Advanced Platform, which builds clusters of virtual machines and is capable of migrating guests across the borders of a physical host.

Fedora 8 gives a clue to where things are headed for Red Hat. The latest Fedora includes the new Xen 3.1, and the Virt-Manager version supports both Xen and QEMU. Additionally, Red Hat is looking to improve the security of the

### Virtualization in Action

Amazon's S3 storage system lets customers store volatile data. Although this offering is still officially in the "Unlimited Beta" phase, some companies are already considering it for production use.

In November 2007, IBM announced that it would offer a similar service, Blue Cloud, based on Xen and IBM's own Power VM. Google also uses Xen: In a move that is atypical for the corporation, which tends not to reveal details of its IT systems, Google introduced the Ganeti management tool, which they have developed specially for this purpose.

Ganeti was released under the GPL in August 2007. The search giant uses the technology for its internal systems, but not for its search engine. Ganeti is best suited to systems with low resource requirements, said Google's Guido Trotter at the LISA 07 conference.

Lufthansa also uses Xen for test systems. Recently, both Oracle and Sun surprised the market with the Oracle VM and Sun xVM Xen variants.

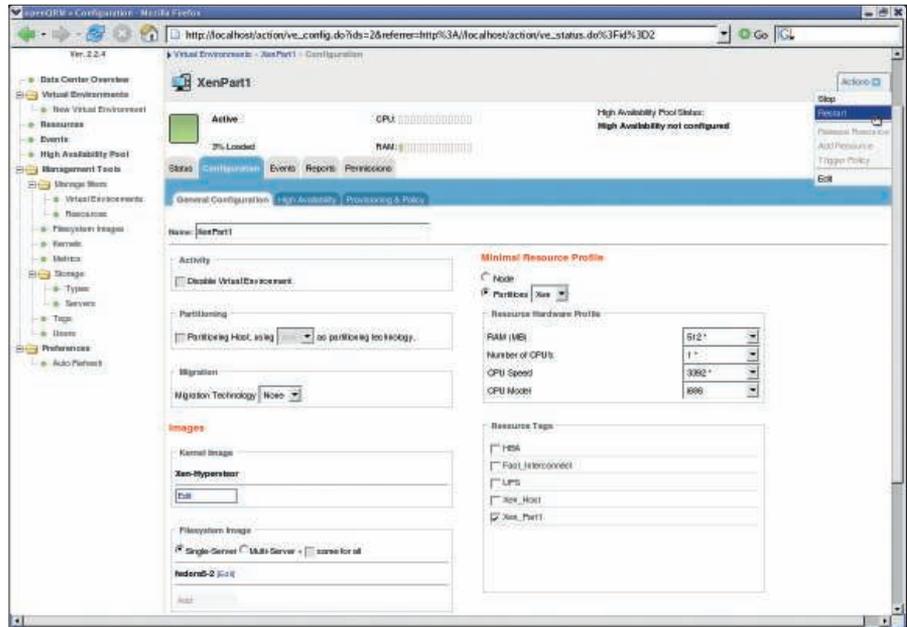


Figure 1: OpenQRM lets you manage physical and virtual computers with a single tool.

administration tool, something that has been a mere sideshow in the past.

Google also has a tool for virtual cluster management, Ganeti, which is released under the GPL. Installing Ganeti is complicated because the tool requires half a dozen Python modules, which Novell, for example, does not provide in a single package.

Ganeti does not offer the convenience of a GUI. This said, the text-based Ganeti commands lend themselves to scripting solutions, which provides a means for integrating the tool with other open source utilities.

Another tool for managing virtual machines is openQRM, a powerful utility that manages images for virtual and

physical machines on the same interface. Another contender is Qumranet's Solid ICE, which focuses entirely on desktop virtualization using the KVM kernel hypervisor.

At the other end of the scale are various small command-line tools, such as xen-tools, a collection of scripts used on Debian for creating and configuring virtual instances.

## Hosting Farms

In many cases, users do not need to simulate a full-fledged computer with a custom kernel for every one of the virtual guests. Professional hosters, who are simply concerned with keeping their customers' web offerings apart, are quite

### What Runs Where?

A rule of thumb dictates that flexibility requires either more performance or special hardware. In hardware virtualization, the physical CPU handles most of the guest's instructions: The hypervisor only steps in to avoid conflict. If the hardware is unable to detect conflict directly, para-virtualization takes the role of manager. This technology is fairly advanced with respect to the main processor, whereas virtualization of I/O components is still at an early stage of development. Emulators simulate almost any scenario and architecture, but their performance is comparatively slow because everything is handled by the software.

All of the virtualization solutions we discuss in this issue support Linux as a guest system, although some require changes to the guest kernel. These changes can cause a problem with support for some applications, especially if the service provider requires a special kernel version. VMware and VirtualBox will basically run on any recent, unpatched kernel, as will the emulators.

Guest support for different versions may be limited. The current Windows versions, XP and Vista, are supported by all server virtualizations; operating system virtualization tools, such as OpenVZ or VServer, work on a different principle and cannot offer this feature.

# Customised IT solutions for individual needs are our strength!



**With over 500 installed clusters, we offer everything – except inflexibility. We cover the entire portfolio from cost-effective entry clusters to enterprise solutions.**

transtec Entry HPC Clusters are suitable for small work groups and fulfill the need for superior computing performance. Our mid-range clusters ensure that there are sufficient resources available for multiple power users or mid-sized work groups. The Enterprise Clusters offer the computing performance of multiple TFLOPs and a large number of compute nodes. You will always have enough power for future requirements and to work efficiently.

**>> [www.transtec.co.uk/go/cluster](http://www.transtec.co.uk/go/cluster)**

**transtec**

## **Any questions? We are happy to help:**

- transtec Computers Ltd. • Suite A, Castle Link • 39 North Bar • Banbury • Oxfordshire OX160TH
- Tel: +44 (0) 1295/756 100 • [www.transtec.co.uk](http://www.transtec.co.uk)

**Microsoft**  
CERTIFIED  
Partner

**Novell**



**AMD**  
Smarter Choice

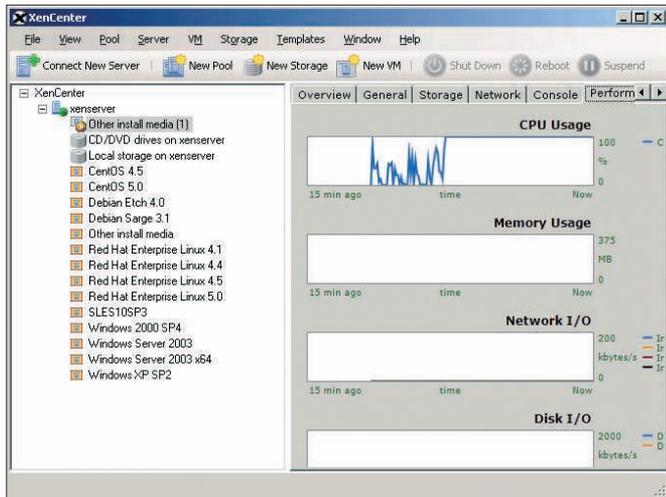


Figure 2: XenCenter, the Citrix Xen Server command center, only runs on Windows.

happy with a single kernel that the guests can share.

Virtualization solutions such as OpenVZ, Virtuozzo, and VServer use

OpenVZ and Virtuozzo, first acquired a majority shareholding in Parallels, the desktop virtualization company, early in 2007 and is now using the Parallels

this approach to the virtualization problem, giving users an amazing application density on normal hardware, with low overhead.

Many professional hosters use the commercial Virtuozzo product to give customers who want to manage their own systems access to virtual consoles. SWsoft, the company behind

brand name. The OpenVZ virtualization tools are available with any major distribution.

### Shopping for Server Virtualization?

The big players in the server virtualization market each offer advantages, although a few of the tools are only for Windows.

#### VMware

Market leader VMware offers an excellent line of virtualization products. However, VMware's VirtualCenter management tool, which plays the role of managing the guests on the ESX and GSX servers, only runs on Windows at present – and it doesn't look like a Linux port will be on the roadmap any time in the near future.

If you prefer a Linux-only solution, you have to make do with the free VM-

Table 1: Virtualization Solutions

Product	Current Version	Vendor	Support	Technology	Typical Deployment Scenarios	Host OS	Guest OS
XenServer	4	XenSource/Citrix	XenSource/Citrix	Hypervisor, para-virtualization, and full virtualization with hardware support	Virtualization for Linux and Windows servers	Modified Red Hat included	Linux, Windows
Xen GPL	3.1.2	Linux distributors	Distributions, independent vendors	Hypervisor, para-virtualization, and full virtualization with hardware support	Virtualization for Linux and Windows servers	Linux, Solaris, FreeBSD, NetBSD	Linux, Windows, FreeBSD, and NetBSD in part para-virtual, Solaris (not all supported on all host systems)
VMware VI	3	VMware	VMware, partner	Hypervisor with full virtualization	Server virtualization, virtual desktops	Linux	Linux, Windows, BSD, ...
VMware Workstation	6	VMware	VMware, partner	Emulation	Desktop virtualization	Windows, Linux	Linux, Windows, BSD, ...
Virtual Iron	4	Virtual Iron	Virtual Iron	Xen hypervisor	Server virtualization	Linux	Linux and Windows
OpenVZ	2.6.22-ovz005	SWsoft and OpenVZ Community	Community	Virtualization at operating system level	Server virtualization on Linux	Linux	Linux in host kernel version
Virtuozzo	4	SWsoft	SWsoft	Virtualization at operating system level	Server virtualization on Linux/Windows	Windows, Linux	Linux in host kernel version, Windows on Windows hosts
VServer	2.2.0.5	Linux VServer project	Distributions, independent vendors	Virtualization at operating-system level	Server virtualization on Linux	Linux	Linux in host kernel version
KVM	55	Qumranet	Community and Qumranet for commercial product	Hardware virtualization	Server and desktop virtualization	Linux	Linux, Windows, BSD, ...
QEMU	0.9.0	Fabrice Bellard	Community	Emulation, optionally accelerated via kernel module for direct access to resources	Software testing, desktop virtualization	Linux, Windows	Linux, Windows, BSD; other processor architectures possible (not with KQEMU)
VirtualBox	1.5.2	Innotek	Innotek	Emulation, hardware virtualization	Optional desktop virtualization	Linux, Windows, OS X	Linux, Windows, BSD in part

## DESTINATION : LINUX

www.linux-magazine.com



News

Article Archive

Online Exclusive Content

All of Linux.

Video Streaming Archive

Linux Magazine Online helps you stay connected to the Linux universe.

Our new website offers exclusive articles on challenging technical topics and provides instant access to the latest Linux news. You can browse for answers by subject or search our extensive article archive.

If you're hunting for the perfect tool, or if you just need some help with tackling a network problem, tune in to Linux Magazine Online!

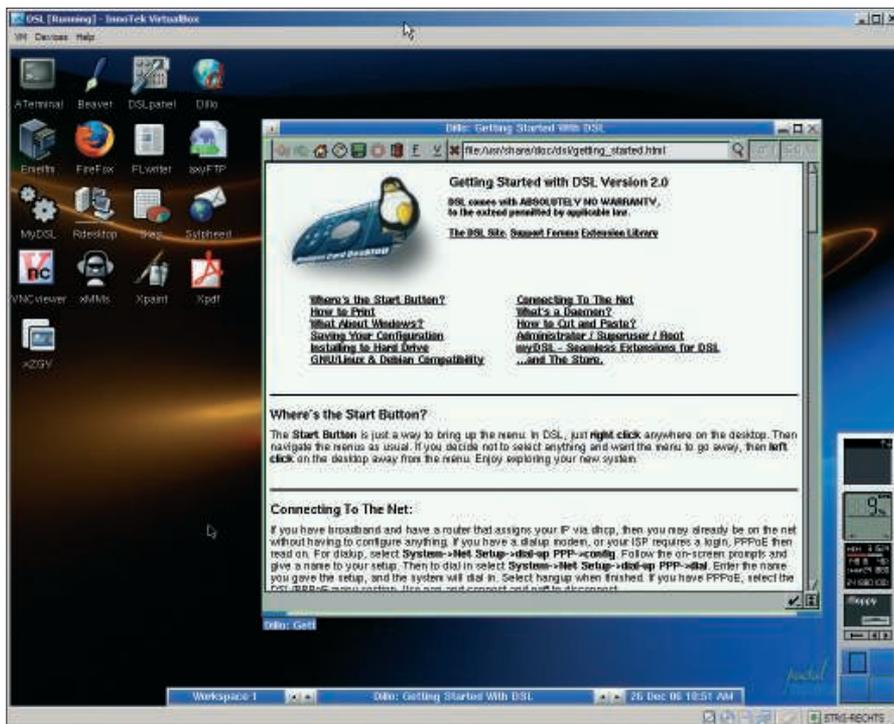


Figure 3: VirtualBox supports either Windows or Linux guests.

ware server. This entry-level solution still has much to offer. The biggest visual change compared with the previous 1.x version on VMware server is the new web interface, Virtual Infrastructure Web Access. (The former standalone server console is no longer usable.) This web interface requires a separate browser plugin (Figure 1).

Provisioning and deployment of VMs is supported by a large number of templates – for 17 Windows versions and 17 Linux distributions (including 64-bit variants) – as well as NetWare and So-

laris 10. Although VMware Server does not let you migrate VMs between physical machines (in contrast to its big brother ESX Server), and definitely not at run time, at least virtual machines can reside in a shared memory area exported by NFS. Load balancing is also restricted to the commercial versions; however, the free version has excellent centralized server administration and monitoring.

### Virtual Iron

Virtual Iron is hot on the heels of VMware. Its Virtual Iron 4 product, or at

least the most expensive Extended Enterprise Edition, not only supports cloning of virtual instances and live migration, but also dynamic capacity management of virtual machines at run time.

The Virtual Iron virtualization environment also includes high-availability features, such as failover for virtual machines on a reserve host ( $N + 1$  clusters). Virtual Iron comes with sophisticated monitoring and logging for virtual instances and policy-based capacity management, which gives administrators the ability to swap out virtual machines automatically if a host exceeds a specific load threshold.

On the storage side, Virtual Iron supports iSCSI, SAN (Fibre Channel), and NAS. VMware's lead on its competitors is no longer very substantial; however, the price differences are still significant. Virtual Iron starts at US\$ 500 per socket, whereas the price for VMware's ESX Server starts at approximately three times that.

### Xen Source

Xen Source Server used to be a Linux virtualization solution that was very

## INFO

- [1] VMware: <http://www.vmware.com>
- [2] Virtual Iron: <http://www.virtualiron.com>
- [3] Novell: <http://www.novell.com>
- [4] Red Hat: <http://www.redhat.com/>
- [5] Xen: <http://www.xen.org/>
- [6] Citrix: <http://www.citrix.com/>
- [7] VServer: <http://linux-vserver.org/>
- [8] Qumranet: <http://www.qumranet.com/>
- [9] KVM: <http://kvm.qumranet.com/>
- [10] Solid ICE: <http://web1.qumranet.com/>
- [11] QEMU: <http://fabrice.bellard.free.fr/qemu/>
- [12] VirtualBox: <http://www.virtualbox.org>
- [13] S3: <http://www.amazon.com/gp/browse.html?node=16427261>
- [14] Parallels: <http://www.parallels.com/>
- [15] Ganeti: <http://code.google.com/p/ganeti/>
- [16] SWsoft: <http://www.swsoft.com>
- [17] OpenVZ: <http://openvz.org/>
- [18] openQRM: <http://www.openqrm.com>
- [19] xen-tools: <http://xen-tools.org>



Figure 4: KVM is popular with both users and kernel developers.

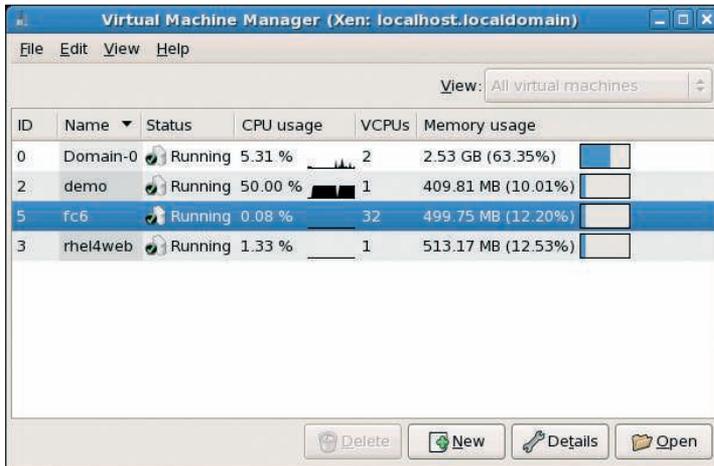


Figure 5: Tracking virtual systems with Virtual Machine Manager.

close to the VMware concept. Citrix recently acquired Xen Source. The server product, which is now known as Citrix Xen Server, is still a Linux system at its core. Unfortunately, the Citrix Xen Server command bridge, which is known as Xen Center, has been transformed into a Windows-only application (Figure 2). One point worthy of notice is that Xen now has a live migration feature, Xen-

Motion, which resembles VMware's VMotion.

### Conclusions

Linux virtualization has entered the real world, and the choice of tools comes down to real-world questions of stability, efficiency, and management. The commercial vendors, with their well-integrated, graphical controls, are currently

the major contenders. On the other hand, a single host is easily managed through less sophisticated alternatives.

Before choosing a virtualization system, take a careful look at your needs and build a comprehensive solution that is easy to manage, monitor, and deploy. Also, leave yourself room to grow, because you're no longer limited by the size of your server room.

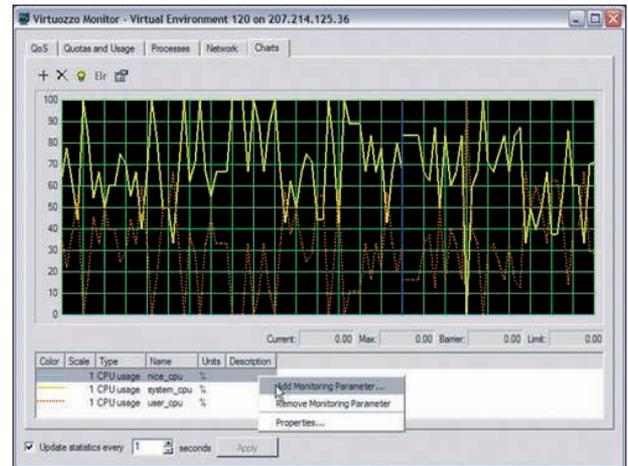
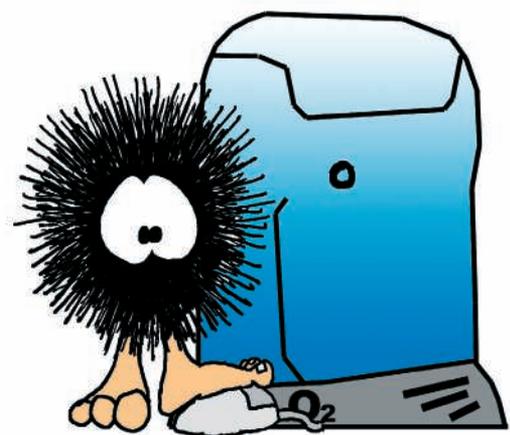


Figure 6: Monitor your virtual environment with Virtuozzo.

# THE MATHEMATICS OF HUMOUR

**TWELVE** Quirky Humans,  
**TWO** Lovecraftian Horrors,  
**ONE** Acerbic A.I.,  
**ONE** Fluffy Ball of Innocence and  
**TEN** Years of Archives  
**EQUALS**  
**ONE Daily Cartoon** that Covers the  
 Geek Gestalt from zero to infinity!

Over Two Million Geeks around the world can't be wrong!  
**COME JOIN THE INSANITY!**



UserFriendly.org