

Commentary on ESX 3.5 'VI3 in a Box' on VMware Workstation 6

1 Introduction

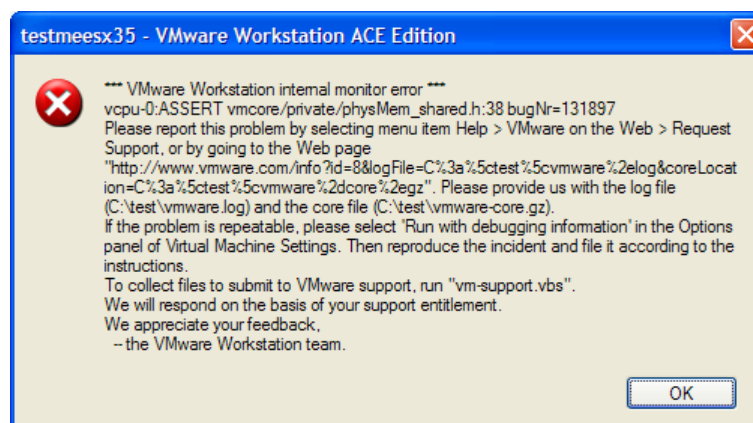
Note: *The views expressed in this article are those of Xtravirt Limited and are purely observational based upon internal research and publically available information.*

In 2007, we, like others in the VMware community, developed a common methodology to install and run VMware ESX 3.x on Workstation 6 together with a shared storage solution based upon an iSCSI software target

(http://www.xtravirt.com/index.php?option=com_remository&Itemid=75&func=startdown&id=11).

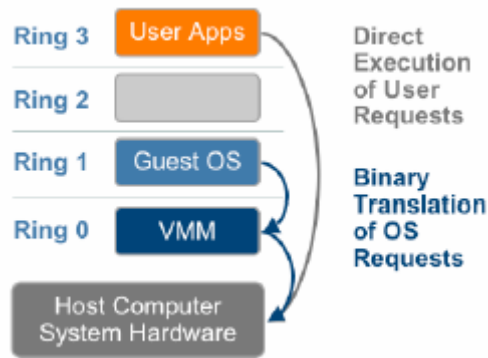
This allowed users and professionals in the IT community to create simple and cost effective VI3 infrastructures to test, learn about and demo these environments.

Upon the release of VMware ESX 3.5 we found that this methodology did not continue to hold true and although it is still possible to install ESX 3.5, when attempting to start a nested VirtualMachine, the ESX VM panics, and a stop message similar to the one below generated by Workstation appears:



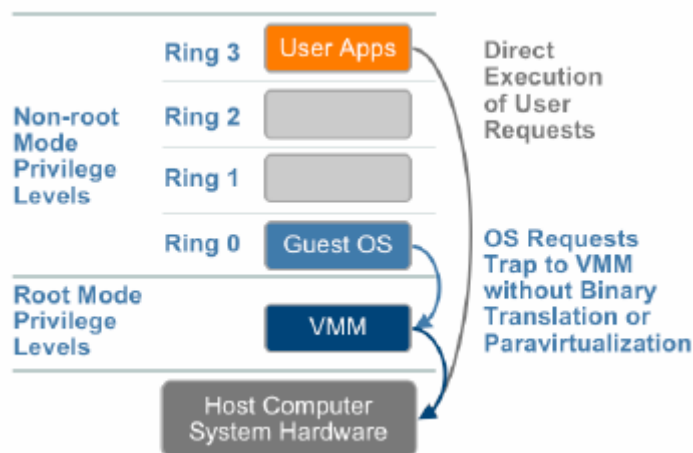
Without a public or officially definitive answer available at the date of writing, the following are a few observations we have made so far.

VMware ESX, x86 hardware virtualization consists of an architecture which uses levels of execution and isolation isolated as 'Rings'. The ESX kernel VMM and the physical hardware sit in the core ring (Ring 0) and VM Guest Operating Systems and applications sit outside of this in lower privilege rings. The VM's OS instructions are intercepted and translated by the VMM using binary to the host hardware. Application requests to hardware are direct which improves VM performance.



Source: VMware – *Understanding Full Virtualization, Paravirtualization, and Hardware Assist*

Currently, hardware manufacturers are embracing this technology and integrating new functionality into their CPU's, which is increasingly aiding virtualization. This is known in some areas as 'Hardware Assisted Virtualization'. It effectively removes some of the overhead previously used in the translation process.



Source: VMware – *Understanding Full Virtualization, Paravirtualization, and Hardware Assist*

Examples of Hardware Assisted Virtualisation are the newer range of CPU's by Intel and AMD with virtualization extensions. Intel-VT and AMD-V CPU technology allows one hardware platform to function as multiple virtual platforms. This technology was one of the pre-requisites for the original 'VI3 in a box' configuration.

Interestingly under ESX 3.5 the additional line in the ESX VM .vmx file "monitor_control.vt32 = TRUE" now works in a different context. This can be tested after the ESX VM has booted up by logging on and typing the following at the command line:

```
Esxcfg-info -w | grep vt
```

```
Esxcfg-info -w | grep hv
```

Under ESX 3.0.x the first command would have resulted in a line which ended with a numerical value to show whether ESX VT was enabled on the hardware. A value of '3' means that VT support is enabled.

Under ESX 3.5.x this technology was augmented by a new term 'HV' support (presumably 'Hardware Virtualization' – or similar) which is why the 2nd command produces a response under 3.5. ESX 3.5 enables support of new Hardware Assisted Virtualisation which may not exist in the majority of chips currently in circulation.

At the time of writing, a new class of CPU's are entering the marketplace providing a 2nd generation of VT extensions. Untested by us at this stage, these include Intel and AMD Quad Core CPU's.

So where is this all leading to?

The net result is that upon power on, the nested VM makes calls to the physical hardware via ESX 3.5 using the new advanced Hardware Assisted Virtualization CPU instruction sets whilst being translated by the existing Workstation 6 hosted architecture.

It appears that primarily the current version 1 VT enabled CPU's do not have the ability to understand and pass the new VT2 extended instruction sets thereby causing the application to crash.

Also VMware Workstation 6, (and all other existing hosted products we've tested to date), doesn't provide the logic to manipulate the new instructions being passed from ESX 3.5 and therefore there are also no available switches to handle compatibility for them.

Those in the VMware community who have been able to test one of the latest CPU's such as the AMD Quad Core 'Greyhound' have found that the CPU has been able to respond to the calls using its new virtualization extension functionality and the 'VI3 in a Box' solution works again.

There have been hints that future versions of these hosted virtualization architectures will be also updated to deal with this new technology, but at the moment it's a case of upgrading your hardware or watching this space.

Xtravirt