

Know Your Enemy: Building Virtual Honeynets

Honeynet Project 2002-08-20

Know Your Enemy: Building Virtual Honeynets

by [The Honeynet Project](#)

last updated August 20, 2002

This article has been contributed to SecurityFocus by the Honeynet Project. For more information on honeypots and honeynets, please visit the Honeynet Project at <http://www.honeynet.org>.

Over the past several years, honeynets have demonstrated their value as a security mechanism, primarily to learn about the tools, tactics, and motives of the blackhat community. This information is critical for organizations to better understand and protect against the threats they face. Among the problems with honeynets is that they are resource intensive, difficult to build, and complex to maintain. Honeynets require a variety of both physical systems and security mechanisms to be effectively deployed. However, the [Honeynet Project](#) has been researching a new possibility, virtual honeynets. These systems share many of the values of traditional honeynets, but have the advantages of running all the systems on a single system. This makes virtual honeynets cheaper to build, easier to deploy, and simpler to maintain.

What is a Honeynet?

Honeynets are one type of [honeypot](#). A honeypot is a resource who's value is in being probed, attacked or compromised. A honeynet is a high-interaction honeypot, meaning it provides real operating systems for attackers to interact with. This high interaction can teach us a great deal about intruders, everything from how they break into systems to how they communicate and why they attack systems. Honeynets accomplish this by building a network of systems. This network is highly contained, where all inbound and outbound traffic is both controlled and captured. Each system within the network is really a honeypot, a system designed to be attacked. However, these honeypots are fully functional systems, the same found in most organizations today. When these systems are attacked, honeynets capture all of the attacker's activity. This information then teaches a great deal about the threats we face today.

For the technical details on honeynets, you are encouraged to review [Know Your Enemy: Honeynets](#). This paper describes different ways of building virtual honeynets. This is not meant to be a how-to on building virtual honeynets. Detailed how-tos will follow. From this point on, it

is assumed you have a understanding of honeynet technologies and their requirements, specifically data control and data capture.

Virtual honeynets take the concept of Honeynet technologies, and implement them into a single system. This article will describe several different ways of building virtual honeynets. This is not meant to be a how-to on building virtual honeynets, such articles will be published at a later date. From this point on, it is assumed that the reader has a understanding of honeynet technologies and their requirements, specifically data control and data capture.

Virtual Honeynets

Virtual Honeynets take the concept of honeynet technologies, and implement them into a single system. Virtual honeynets are not a new concept, instead they take the existing concept of Honeynets and implement them in a different fasion. This implementation has its unique advantages and disadvantages over traditional honeynets. The advantages are reduced cost and easier management, as everything is combined on a single system. However, this simplicity comes at a cost. First, you are limited to what types of operating system you can deploy by the hardware and virtualization software. Second, virtual honeynets come with a risk, specifically that an attacker can break out of the virtualization software and take over the Honeynet system, bypassing data control and data capture mechanisms.

There are two types of virtual honeynet: self-contained and hybrid. We will first define these two different types, and then cover the different ways that virtual honeynets can be deployed.

Self-Contained Virtual Honeynet

A self-contained virtual honeynet is an entire honeynet network condensed onto a single computer. The entire network is virtually contained on a single, physical, system. A honeynet network typically consists of a firewall gateway for data control and data capture, and the honeypots within the honeynet. You can see a [diagram of such a deployment](#) here. Some advantages of this type of virtual honeynet(s) are:

- **Portable** - Virtual Honeynets can be placed on a laptop and taken anywhere. The Honeynet Project demonstrated this funtionality at the [Blackhat Briefings in August, 2002](#).
- **Plug and catch** - You can take the one box and just plug it in to any network and be

ready to catch those blackhats. This makes deployment much easier, as you are physically deploying and connecting only one system.

- **Cheap in terms of time and money** - You only need one computer, so it cuts down on your hardware expenses. It also has a small footprint and only takes one outlet and one port! For those of us with very limited space and power, this is a life saver.

Virtual honeynets also have several disadvantages, such as:

- **They create a single point of failure** - If something goes wrong with the hardware, the entire honeynet could be out of commission.
- **They require a high-quality computer** - Even though a self-contained honeynets only requires one computer, it will have to be a good one. Depending on your set-up, users might need a lot of memory and processing power;
- **They may weaken security** - Since everything might be sharing the same hardware, there is a danger of an attacker getting at other parts of the system. Much of this depends on the virtualization software, which will be discussed later; and,
- **Limited software** - since everything has to run on one box, you are limited to the software you can use. For instance, its difficult to run Cisco IOS on an Intel chip.

Hybrid Virtual Honeynets

A hybrid virtual honeynet is a combination of the classic honeynet and virtualization software. Data capture, such as firewalls, and data control, such as IDS sensors and logging, are on a seperate, isolated system. This isolation reduces the risk of compromise. However, all the honeypots are virtually run on a single box. You can see [a diagram of such a deployment here](#). Some advantages to this setup are:

- **Secure** - As we saw with the Self-Contained Virtual Honeynets, there is a danger of an attacker getting to the other parts of the honeynet (like the firewall). With Hybrid Virtual Honeynets, the only danger would be that the attacker accessing to the other honeypots.
- **Flexible** - Users are able to use a myriad of software and hardware for the classic parts of the hybrid network. An example would be that you can use the OpenSnort sensor or a Checkpoint appliance on the network. You can also run any kind of honeypot you want because users can just drop another computer on the network (in addition to the virtual honeypot's box).

Among the disadvantages of self-contained virtual honeypots is that they are:

- **Not portable** - Since the honeynet network will consist of more than one box, it makes it more difficult to move.
- **Expensive in terms of time and space** - You will have to spend more in terms of power, space and possibly money since there is more than one computer in the network.

Virtualization Software

Hybrid virtual honeynets can allow you to leverage the flexibility of classic honeynets and let you increase the amount of honeypots by using virtualization software. Now that we have defined the two general categories of virtual honeynets, let's highlight some of the possible ways to implement a virtual honeynet. Here we outline three different technologies that allow you to deploy your own. Undoubtedly there are other options, such as [Bochs](#); however, the Honeynet Project has used and tested all three methods. No one solution is better than the other. Instead, they each have their own unique advantages and disadvantages, it is up to you to decide which solution works best. The three options we will now cover are VMware Workstation, VMware GSX Server, and User Mode Linux.

VMware Workstation [VMware Workstation](#) is made for the more casual user. It is a long-used and well-established virtualization option that is available for Linux and Windows platforms. Some advantages to using VMware Workstation as a virtual honeynet include:

- **Wide range of operating system support** - You are able to run a variety of operating systems within the virtual environment (called GuestOSs), including Linux, Solaris, Windows and FreeBSD honeypots.
- **Networking options** - Workstation provides two ways to handle networking. The first is bridged, which is useful for hybrid virtual honeynet networks because it lets a honeypot use the computer's card and appear to be any other host on the honeynet network. The second option is host-only networking, this is good for self-contained virtual honeynets because you are able to better control traffic with a firewall.
- **VMware Workstation creates an image of each Guest Operating System** - These images are simply a file, making them highly portable. This means you can transfer them to other computers. To restore a honeypot to its original condition, you can just copy a backup into its place.
- **Ability to mount VMware virtual disk images** - You are able to mount a VMware

image just like you would a drive using `vmware-mount.pl`.

- **Easy to use** - VMware Workstation comes with a graphical interface (both Windows and Linux) that makes installing, configuring, and running the operating systems very simple.
- **Support** - As a commercial product, VMware Workstation comes with support, upgrades, and patches.

Some disadvantages of VMware Workstation include:

- **Cost** - VMware workstation costs around \$300 per license. This might be a bit much for the hobbyist or the unemployed student.
- **Resource requirements** - VMware Workstation must run under an X (or GUI) environment and each virtual machine will need its own window. So on top of the memory that is allocated for the GuestOSs (the OS that is installed on a virtual machine), there is the overhead of the GUI system.
- **Limited number of GuestOSs** - With VMware you can only run a small amount of virtual machines (approximately 1-4). This might make for a small honeynet.
- **Closed Source** - since VMware is closed source, users can't really make any custom adjustments.
- **Fingerprinting** - It may be possible to fingerprint the VMware software on a honeypot, especially if the "VMware tools" are installed on the systems. This could give the honeypots away to the blackhat. However, VMware Workstation does have options that can make fingerprinting more difficult, such as the ability to set the MAC address for virtual interfaces.

You can learn more about deploying a self-contained virtual honeynet using VMware Workstation at <http://online.securityfocus.com/infocus/1506>.

VMware products also have some nice features, like the ability to suspend a virtual machine. You are able to pause the VM, and when you take it out of suspension, all the processes go on like nothing happened. In our research, a system was once compromised and the intruder started an ICMP fragment attack. The intruder was also logged into IRC servers. We did not want to cut the connection because we would lose valuable information. So we suspended the VM, adjusted the firewall to block the attack, then brought the VM back up.

An interesting use of VMware, and other virtualization software too, is the ease and speed of bringing up VMs. Once a honeynet is compromised, and we have learned as much as we can

from it, we want to start over. With a virtual honeynet, all we have to do is copy files or use the undoable disk or nonpersistent disk feature in VMware Workstation to discard any changes made. Another feature of VMware Workstation is the ability to run several networks behind the HostOS. So if you only have one box, you can have your honeynet and personal computers all on the same box without worrying about data pollution on either side.

If you would like to learn more about VMware and its capabilities for honeypot technology, check out Kurt Seifried's excellent paper [Honeypotting with VMware - The Basics](#).

VMware GSX Server

The [VMware GSX Server](#) is a heavy-duty version of VMware Workstation. It is meant for running many higher-end servers. As we will see, this is perfect for use as a honeynet. GSX runs on Linux and Windows, advantages include:

- **A wide range of operating system support** - GSX Server support Windows (including 95, 98, NT, 2000, XP and .NET server), various Linux distributions, and potentially BSD and Solaris (Not officially supported).
- **Networking** - GSX Server includes all of the options that Workstation has.
- **No X (GUI) means more GuestOSs** - GSX Server does not need X running in order to have VMware running. This allows you to run many more GuestOS's at the same time. However, it does require that some of the X libraries be installed if the Host is running Linux.
- **Web Interface** - GSX can be managed through a Web page interface. GuestOSs can be started, paused, stopped and created via the Web page.
- **Remote terminal** - This is one of the best features of GSX Server. Through the Web page and with some VMware software, you can remotely access the GuestOSs as if you were sitting at the console. You are able to do things like remote installs and checking out the system without generating traffic on the honeynet.
- Ability to mount VMware images, just like in Workstation.
- **Capacity** - VMware GSX Server supports more host memory (up to 8GB), more CPU's (up to 8), and more memory per virtual machine (2GB) than VMware Workstation.
- Includes a Perl API to manage GuestOSs.
- **Support** - Similar to Workstation, GSX Server is a supported product, including patches and upgrades.

Some disadvantages of VMware GSX Server include:

- **Cost** - a GSX license will run around \$3,500;
- **Limited types of GuestOSs** - Operating systems like Solaris X86 and FreeBSD are not officially supported (however, you may be able to install them). This can limit the diversity of your Honeynet.
- **Memory hog** - GSX recommends at least 256 MB just to run the GSX software. GUI-based operating systems, such as Windows XP, require another 256 MB for each instance.
- **Closed source** - just like Workstation.
- **Fingerprinting** - It may be possible to fingerprint the VMware software on a honeypot, especially if the "VMware tools" are installed on the systems. This could give the honeypots away to the blackhat. However, as with Workstation, there are configuration options that can reduce that risk.

VMware also makes an [VMware ESX server](#). Instead of being just a software solution, ESX Server runs in hardware of the interface. ESX Server provides its own virtual machine OS monitor that takes over the host hardware. This allows more granular control of resources allocated to virtual machines, such as CPU shares, network bandwidth shares and disk bandwidth shares, and it allows those resources to be changed dynamically. This product is even higher end than GSX Server. Some of its features are: it can support multiple processors, more concurrent virtual machines (up to 64 VMs), more host memory (up to 64GB) and more memory per virtual machine (up to 3.6GB) than GSX Server.

User Mode Linux

[User Mode Linux](#) is a special kernel module that allows you to run many virtual versions of linux at the same time. Developed by Jeff Dike, UML gives you the ability to have multiple instances of Linux, running on the same system at the same time. It is a relatively new tool with great amounts of potential. Some advantages to using User Mode Linux are:

- **Free and open source** - you have access to the source code.
- **Small footprint and fewer resource requirements** - User Mode Linux does not need to use X. It can also run extensive amount of systems with little memory.
- **Ability to create several virtual networks** - ULM allows users to create several virtual networks and even create virtual routers all inside the original virtual network.

- Supports both **bridging** and **networking**, similar to VMware.
- **UML has the ability to log keystrokes through the GuestOS kernel.** The keystrokes are logged right on the HostOS, so there are no issues with how to get the keystrokes off the honeypot in a stealth way.
- **UML comes with pre-configured and downloadable file systems**, making it fast and easy to populate your honeynet with honeypots. Like VMware, these file system images are mountable.
- **You can access UML consoles in a wide variety of ways**, including through pseudoterminals, xterms, and portals on the host which you can telnet to. And there's always screen. Run UML inside screen, detach it, and you can log in to the host from anywhere and attach it back.

Some disadvantages are:

- **Currently only supports Linux virtual machines.** UML currently only supports Linux virtual machines; however, a port to Windows is under development.
- **Bugs, other issues** - As a new tool, there are some bugs, documentation, security issues.
- **There is no GUI interface** - currently all configuration and implementation is done at the command line. This entails a steeper learning curve.
- **No support** - As an OpenSource tool, there is no official or commercial support.
- **Fingerprinting** - Similar to VMware, it may be possible to fingerprint a UML Honeynet due to the virtualization software. However, the maintainer, Jeff Dike, is taking measures to reduce that risk, such as modifications to the `/proc` on the guest operating systems.

Conclusion

The purpose of this paper was to define what a virtual honeynet is, the different types, and options for deploying them. Virtual honeynets take the technology of a honeynet and combine them on a single system. This makes them cheaper to build, easier to deploy, and simpler to maintain. However, they also share common disadvantages, including a single point of failure and limitation with both the physical hardware and virtual software. It is up to you to decide which solution is best for your environment. In the future, we intend to develop documentation detailing how to deploy these technologies.

Relevant Links

[The HoneyNet Project](#)

[Privacy Statement](#)

Copyright 2006, SecurityFocus