UNIVERSITEIT VAN AMSTERDAM

Xen hypervisor security in VM isolation

Author:
Yanick de Jong
E-mail:
yanick.dejong@os3.nl

Supervisors:
Fred Mobach
Mendel Mobach

Version:
Final
Date:
30 January, 2009

Abstract

Nowadays cost reduction, efficiency, flexibility, manageability and green IT are items that would be important for organizations. Virtualization is a good solution here for. Examples of virtualization tools are Xen, VMware, KVM, Virtualbox, and so many more. With virtualization servers would use the CPU power, memory, disks, etc... more efficient, so reducing the amount of servers that are running This is help for cost reduction, but also for a more green nature.

This research will primary look if it safe to merge some Xen host, who are in different network segments, with different security levels. If it is safe then research delivers also a contribution to further cost reduction, efficiency, etc...

Contents

1 Introduction

Nowadays cost reduction, efficiency, flexibility, manageability and green IT are items that would be important for organizations. Virtualization is a good solution here for. Examples of virtualization applications are Xen, VMware, KVM, Virtualbox, and so many more. With virtualization servers would use the CPU power, memory, disks, etc... more efficient, so reducing the amount of servers that are running This is help for cost reduction, but also for a more green nature.
This research will contribute to further cost reduction, efficiency, etc... to look if it safe to merge some Xen host, who are in different network segments, with different security levels. The report will gave an advice, if it safe to merge a virtual machines to one Xen host.

1.1 Research Focus

This research give an advice, if it is safe when merging virtual machines (VM's) on one Xen host that has connectivity to more network segments with different security levels. This research will focuses on the question, when you using Xen for virtualization, all other virtualization applications become out of consideration. Each chapter will treat a few of the questions that are relevant to answer the research question. So the eventual conclusion if it isn't safe or safe under some conditions is based on all these chapters.

Research Question: *What are the risks involved with merging Xen servers in different segments of the network and put all virtual machines together on one machine?*

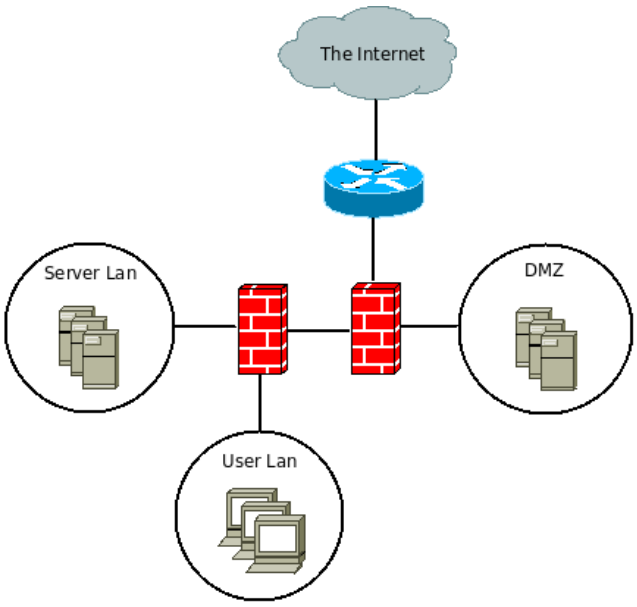

Figure 1: Current and safe network overview

In figure 1 you see a network design that is make following the standards. This design works fine, also if you use virtualization, but at some consideration you may want the merge all virtual machines to one Xen host, that would then serve virtual DMZ machines and virtual LAN server machines. There a many reasons why you should do this, for example cost reduction, power saving, less manageability, more flexibility, etc..., but each advantage has also his disadvantage.

Figure 2 presents what the new network design would be if you make use of virtualization, and maybe you must do it and maybe not, but that can you a the conclusion.

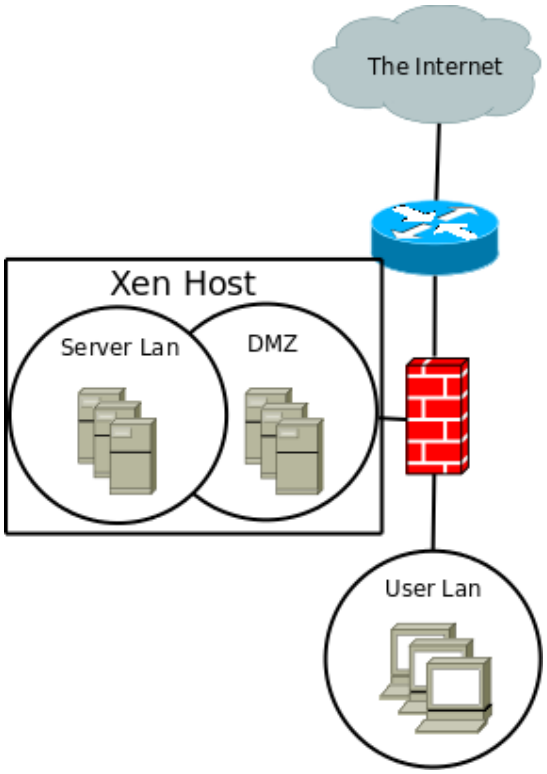

Figure 2: New and maybe safe network overview

During the research all things are tested with OpenSuse 11.1 and Xen version 3.3. The virtual machine are also running OpenSuse 11.1.

1.2 Structure of the report

In this introduction you could read why this research would be use full in today's context of virtualization. In chapter 2 you could some information for this research. The other chapters *Network & System* 3 *Disk Alloaction* 4 *Memory* 5 *Bridging* 6 *DMA* 7 gives an answer to the subquestions and the result of the experiments. Finally all the information will came to a final conclusion 8 , that gives an answer on the research question.

2 Virtualization

Virtualization is simply the logical separation of the virtual request for some service from the physical resources that actually provide that service. Virtualization provides the ability to run applications, operating systems or system services in a logically distinct system environment that is independent of a specific piece of hardware. Obviously all of these have to be running on a certain computer system at any given time, but virtualization provides a level of logical abstraction that liberates applications, system services, and even the operating system that supports them from being tied to a specific piece of hardware. The virtualization focuses on logical operating environments rather than physical ones makes applications, services, and instances of an operating system portable across different physical machines. [4]

There are many kinds of virtualization, but most people think about server or machine virtualization, when they hear the word *virtualization*. Other kinds of virtualization are application, desktop, network, storage and system-level virtualization. The research would only look at machine virtualization.
The two kinds of server or machine virtualization are:

Full Virtualization Full virtualization emulates the entire hardware platform of a guest computer. That can be effective for running an otherwise incompatible operating system, like Windows on a Linux server. The advantage you have the same hardware as the Xen host, but the disadvantages it works slower. [1]

Paravirtualization Para-virtualization uses a customized kernel, compatible with the host's kernel and "hyper visor", that speaks compatibly and much more directly to your host's hardware. It's much lighter weight, allows memory to be re-allocated among guest domains so a server can run far more guest domains,, and provides a really noticeable to any guest operation that has to talk to the disk. But it requires a compatible kernel on the guest OS, compatible with the Xen version of the host OS. The hardware will be virtualized by standard, so there is implemented a standard video card and cabinetwork and so on. This has the advantage that paravirtualization works fast. [1]

2.1 Advantages

Xen has advantages, and they would be enumerate below:

- Efficient use of the capacity of the hardware. Processor uses increases of 20 to 53 percent.
- Consequences of the above item would be advantage in case of reduction of (new) hardware, management costs, power and cooling supplies
- Hardware independent
- Virtual machines (VM's) with all kind of operating systems and application are transportable and could be moved to another virtualization support server.
- Consequences of the above item would be advantage in handling flexible with VM's so you could recover easily crashed VM's and minimized the downtime to almost zero. Virtualization simplifies system and application testing, system installation and deployment, support of legacy systems and application, system-level development

  [2] [3] [4]

2.2 Disadvantages

Xen has also disadvantages and they would be enumerate below:

- Increase the administrative complexity and debugging time.
- Single point of failure problems
- Server sharing and performances issues
- Increase in network complexity and debugging time.
- Network and IT engineers haven't or too little knowledge about virtualization. So the engineers need education. [2] [4]

3 Network & System

A system is nothing without a network today and vice versa. The networks will connectivity to the system, but greatest also risks for the system, because it could be reached by people on the network. For that we segment network and define for each segment network security levels, so system and network engineer could protect there systems. System and network engineer do all kinds of security, to hold there system and network secure and manageable, so no other people could take over control.

Security of a infrastructure became more and more important, because there is interesting information inside a company for outsiders to make money with. So the risks through the years of computing became larger, because more people get and use a computer for all kind of aims. Organizations don't want any risks, because it costs there profit or there would be a chance that some other organization earns moment with a product that they have made. So organizations want no risks, but that is impossible so you want to minimize the risks. A risk depends on the impact, how great the chance would be that the risks is happening and the damage. So you could make for each risks an analysis and take some precaution measures depends on the analysis.

Is it safe to place *DMZ servers* and *LAN servers* in one network segment?

We want to merge all virtual servers from the network segment called the Demilitarized zone (DMZ) with the Local Area Network (LAN) server segments. In case of a non virtualized network it would be sin, and all system and network engineer won't agree to implement it this way. This would be in fact the first argument why you also don't want to do it with virtualized systems.

Second fact if you merge all the virtual machines to one network segment you lose you security levels and make you network more weak. You could solve it a bit by making good firewall rules based in IP addresses of virtual machines, who has access for where and with which protocols. It won't be the ideal network.

There is also a system and not only the network, but also to the systems. The third fact, you create all single point off failure even if you have redundancy. It would know far more interesting to hack the Xen host and get control over all virtual machines. In the past you must hack one host to control it and

now you can get all lot of running virtual machines to control, what has far more impact. And the network and system hanging together because you would satisfied with less security segments it may also easier to hack the Xen host.

In practice a network engineer would also make physical distinction between his internal server LAN segment and his DMZ that is open for the dangerous world on Internet. So do you want to merge your virtual machine from the DMZ and the LAN server segment together in one network segment. No you won't. [6][7][8][5]

What happens if some one take control of the operating system where Xen is running on?

Virtualization is based on a few layers see figure 3. In practice you couldn't take control over the hardware, because it stands safely in a data center. The data center is only accessible for authenticated persons and the hardware is lock up in cabinets.

On the hardware runs a host operating system. None of the operating systems, which one also would be used is 100 percent secure. There is always a bug or something so a hacker could crack the host operating system. If you could crack the host operating system, then you have full control of the DOM0 and all DOM U (virtual machines).

So what is the risks that an operating system would be hack? This risks would be right for a virtualized system as for a non virtualized system, but the impact and the damage would be different. If a virtualized system is hacked then all virtual machine could be shut downed and your hole service is gone and that costs a lot of money. In fact of a non virtualized system, it would maybe one system, but the other service will still functioning.

With merging all virtualized machine on one Xen host in one dangerous network segment you would dig your one sepulcher and no judicious system or network engineer would do that. [9]

What are the risks if you can take over a virtual machine?

Machine is a machine if it is virtualized or not, because it has still a function to accomplish. If some one could take over the virtual machine and damage if functioning then a service wouldn't be available, what is still the same to a

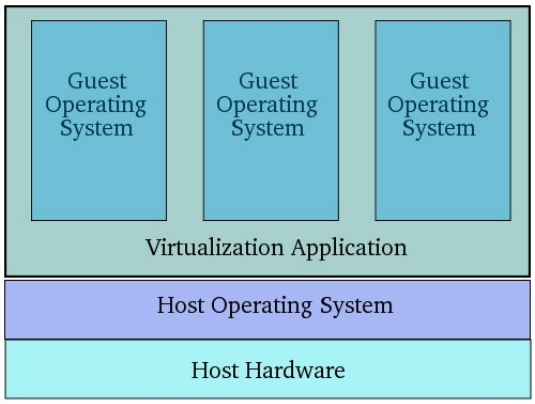

Figure 3: Overview of Virtualization

non virtualized system. The risks are the same, but the impact and the damage are different. You lost maybe confidentiality, availability and integrity. Depending on that facts you should decide the impact and the damage.

It is still not proved that you come out of virtual machine, so that is safe. In case of virtualization you could kill the virtual machine a build a new one, based on backup of the image you had made. So it is not fine thought that some one takes over a virtual machine, but is less worse then when its a normal machine and you do a new installation, having downtime, etc...

3.1 Conclusion

If secure is the most important point, then you merge place DMZ servers and LAN servers in the same network segment, because it has consequences for other hardware and software that running on the network. So you can conclude that virtualization is good product, but with security points where you must be aware of as system and network engineer.

4 Disk allocation

A VM needed disk space to install the operating systems and write and read files just like a non virtualized system. When you create a new virtual machine there would be automatically made a new virtual disk by hand of what you have configured in the configuration of the VM. On a non virtualized system would it possible to write some blocks of data somewhere, even it you haven't any rights to write it.

Question: *Is it possible to read and write outside of the allocated virtual machine disk allocation?*

4.1 Experiments & Results

The write outside of the allocated disk space of a virtual machine (VM), we first need to know which blocks are allocated to the VM. Making a dump of the virtual disk of the VM gives the information about which blocks would be allocated or know. Below you could see the begin of dump and the end of dump, what were in the middle is left out.

Begin of the allocated blocks

```
linux-virthost10:~ # cat /dev/xvda2 | hexdump -C | head -n 10
00000000  eb 48 90 00 00 00 00 00  00 00 00 00 00 00 00 00  |.H..............|
00000010  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |................|
*
00000030  00 00 00 00 00 00 00 00  00 00 00 00 00 00 03 02  |................|
00000040  ff 01 00 80 fc 89 7a 00  00 08 fa 90 90 f6 c2 80  |......z.........|
00000050  75 02 b2 80 ea 59 7c 00  00 31 c0 8e d8 8e d0 bc  |u....Y|..1......|
00000060  00 20 fb a0 40 7c 3c ff  74 02 88 c2 52 be 81 7d  |. ..@|<.t...R..}|
00000070  e8 36 01 f6 c2 80 74 56  b4 41 bb aa 55 cd 13 5a  |.6....tV.A..U..Z|
00000080  52 72 4b 81 fb 55 aa 75  45 a0 41 7c 84 c0 78 3e  |RrK..U.uE.A|..x>|
00000090  75 05 83 e1 01 74 37 66  8b 4c 10 be 05 7c c6 44  |u....t7f.L...|.D|
```

End of the allocated blocks

```
linux-virthost10:~ # cat /dev/xvda2 | hexdump -C | tail -n 10
140248f90  b2 97 4c c1 9c 0d 18 00  04 bc 61 c1 f4 06 14 00  |..L.......a.....|
140248fa0  72 5a 6b c1 08 07 14 00  38 8b 82 c1 1c 07 10 00  |rZk.....8.......|
140248fb0  2e f7 91 c1 2c 07 14 00  32 26 aa c1 40 07 10 00  |....,...2&..@...|
140248fc0  f4 68 bc c1 50 07 1c 00  b8 3e dd c1 6c 07 18 00  |.h..P....>..l...|
140248fd0  bc 04 14 c2 84 07 14 00  24 dd 31 c2 98 07 14 00  |........$.1.....|
140248fe0  04 4f 40 c2 ac 07 1c 00  be 47 56 c2 c8 07 18 00  |.O@......GV.....|
140248ff0  fe 8f 82 c2 78 0c 1c 00  3e 15 a0 c2 e0 07 18 00  |....x...>.......|
140249000  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |................|
*
```

Here below would be first written to an allocated block of the VM and that is legal and works fine.

Writing to a allocated block

```
linux-virthost10:~ # dd if=file count=140249000
yanickyanickyanickyanickyanickyanickyanickyanickyanickyanickyanickyanickyanickyanickyanickyanick
...
yanickyanickyanickyanickyanickyanickyanickyanickyanick
2+1 records in
2+1 records out
1495 bytes (1.5 kB) copied, 0.000533016 s, 2.8 MB/s
```

Now write to unallocated piece of disk space and look what happen. The tool gives an error and that error came from the fact that the piece of disk space where is written to isn't allocated in the VM and not know.

Writing to a non allocated block

```
linux-virthost10:~ # dd if=file count=140249a00
dd: invalid number ‘140249a00’
```

4.2 Conclusion

Writing outside of the allocated disk space seem to like impossible in my case, so that would be safe. In theoretical it must also be impossible, because which location isn't allocated, isn't know by the kernel of the VM and kernel couldn't allow the action to write to another block. So is disk full, also disk full in a VM, you can't write any data any more.

5 Memory

Xen allocates a piece of the physical memory for its own use. DOM0 the first domain started by the Xen hypervisor on boot manages the memory. DOM0 has special privileges, like starting, stopping, pausing domains and is also able to access hardware directly. In the top 64MB on 32-bit systems, the top 168MB on Physical Address Extension (PAE) systems and on 64 bit systems a portion in the middle of the address space will reserved for every virtual address space. The other unreserved memory is available for the domains or virtual machines (VM). Each VM could have a maximum physical memory allocation and if the guest OS runs a 'balloon driver', the memory can dynamically adjust to limit. [10] [12]
Xen makes a distinction between machine memory and pseudo-physical memory.

Machine memory Machine memory is the physical memory that is installed in the machine and that would be used by host operating system and the DOM0 that serves the virtualization.

Pseudo-physical memory Pseudo-physical is ordered per-domain abstraction. It gives the operating system its memory allocation that consist of a contiguous range of physical page frames starting at physical frame 0. This allocated memory is only know by the guest operating system and could be sparsely allocated and in any order on the host operating system.

The machine memory table which records the mapping from the machine page frame to the pseudo-physical ones. Each domain is supplied with a physical-to-machine table which performs the inverse mapping. The machine-to-physical table has size proportional to the amount of RAM installed in the machine, while each physical-to-machine table has size proportional to the memory allocation of the given domain. [11]
So every domain or VM has his own pieces of memory that is mapped to the physical memory. The memory allocated for a domain won't consist of a contiguous range of physical page frames starting at physical frame 0, but it would be located every where address space is free. This means that it is possible to read out the memory from the DOM0, and you can see everything in memory what also is used by a domain.
But is it also possible to read out the memory from inside of the VM. Theoretical it isn't possible, because you break out of the virtual machine and in many researches it is proved that it almost impossible. Maybe there is a memory issue so the following questions are part of the research.

- Is it possible to read the memory that is not allocated to the virtual machine, such as memory of the host in which the virtual machine is running?
- Is it possible to read the memory of a virtual machine that has been stopped, from another virtual machine?
- Is it possible to read the memory of a virtual machine that has been stopped, from another virtual machine that has been started later and allocated the same memory?

In the next paragraph the experiments and the result will be shown.

5.1 Experiments & Results

Experimenting if the memory safe for reading and writing outside of a machine is the issue of this experiment. For this experiment there is one remark, because otherwise the experiment could also fail and you don't no why. Some Linux operating systems like Ubuntu, Red hat, Cent Os, etc... will make use of Security Enhanced Linux or SE Linux. What does SE Linux? SE Linux So first write something to the memory that a domain or virtual machine has allocated from the DOM0. Below you could see that some string is written to the memory.

Writing to memory

```
linux-virthost10:~ # echo yanickyanickyanickyanick > /dev/mem
```

There is written something in to the memory of the virtual machine and now find it back in the memory. Below you could see how to get the information out of the memory, make it a little readable en search for some string in all that information.

Try to find *yanick* in memory

```
linux-virthost10:~ # cat /dev/mem | hexdump -C | grep yanick
```

Unfortunately nothing was found, so maybe you can't read the memory at all. So below the memory will be read without any conditions. The result is that you only get back some zeros and a star the star means that the rest would still zero. So in domain or virtual machine (VM) you can't read any memory. So it might be impossible, but there is nothing to find about this in manuals, on the internet or somewhere. So far I could conclude that the answers on my questions would be no, it is not possible to read the memory when your in a virtual machine.

Try to read memory in VM

```
linux-virthost10:~ # cat /dev/mem | hexdump -C
00000000  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |................|
*
```

In domain or VM you can't read the memory, but there is another way even if it may not issue because a user of VM doesn't have any rights on the DOM0. On DOM0 you could manages the domains or VM's. With the *xm save* command you save the state of the domain to a file, see below. In the file will also put all memory information, that is in the memory at the moment of making a save.

Searching for ...

```
yjong:~ # xm save -c opensuse11 /home/checkpointfile/checkpointfileopensuse11-2
```

So let's discover if we can find in the file of what was written to memory. So take the file and look if there is something readable and yes there is the string that is written to memory. Below you could see that there make a the results to a file, because it is big.

Searching for ...

```
yjong:/home/checkpointfile # cat checkpointfileopensuse11-2 | hexdump -C |
grep yanick > dump
```

Below the result of the dump, the first and the last ten lines show that the string is in there.

Begin of dump out of VM checkpointfile

```
yjong:/home/checkpointfile # cat dump | head -n 10
048a03b0  61 6e 69 63 6b 79 61 6e  69 63 6b 79 61 6e 69 63  |anickyanickyanic|
048a03c0  6b 79 61 6e 69 63 6b 79  61 6e 69 63 6b 79 61 6e  |kyanickyanickyan|
048a03d0  69 63 6b 79 61 6e 69 63  6b 79 61 6e 69 63 6b 79  |ickyanickyanicky|
048a03e0  61 6e 69 63 6b 79 61 6e  69 63 6b 79 61 6e 69 63  |anickyanickyanic|
048a03f0  6b 79 61 6e 69 63 6b 79  61 6e 69 63 6b 79 61 6e  |kyanickyanickyan|
048a0400  69 63 6b 79 61 6e 69 63  6b 79 61 6e 69 63 6b 79  |ickyanickyanicky|
048a0410  61 6e 69 63 6b 79 61 6e  69 63 6b 79 61 6e 69 63  |anickyanickyanic|
048a0420  6b 79 61 6e 69 63 6b 79  61 6e 69 63 6b 79 61 6e  |kyanickyanickyan|
048a0430  69 63 6b 79 61 6e 69 63  6b 79 61 6e 69 63 6b 79  |ickyanickyanicky|
048a0440  61 6e 69 63 6b 79 61 6e  69 63 6b 79 61 6e 69 63  |anickyanickyanic|
```

End of dump out of VM checkpointfile

```
yjong:/home/checkpointfile # cat dump | tail -n 10
1faec970  79 61 6e 69 63 6b 79 61  6e 69 63 6b 79 61 6e 69  |yanickyanickyani|
1faec980  63 6b 79 61 6e 69 63 6b  79 61 6e 69 63 6b 79 61  |ckyanickyanickya|
1faec990  6e 69 63 6b 79 61 6e 69  63 6b 79 61 6e 69 63 6b  |nickyanickyanick|
1faec9a0  79 61 6e 69 63 6b 79 61  6e 69 63 6b 79 79 1b 5b  |yanickyanickyy.[|
1faec9b0  31 35 3b 31 48 61 6e 69  63 6b 79 61 6e 69 63 6b  |15;1Hanickyanick|
1faec9c0  79 61 6e 69 63 6b 79 61  6e 69 63 6b 79 61 6e 69  |yanickyanickyani|
1faec9d0  63 6b 79 61 6e 69 63 6b  79 61 6e 69 63 6b 79 61  |ckyanickyanickya|
1faec9e0  6e 69 63 6b 79 61 6e 69  63 6b 79 61 6e 69 63 6b  |nickyanickyanick|
1faec9f0  79 61 6e 69 63 6b 79 61  6e 69 63 6b 79 61 6e 69  |yanickyanickyani|
1faeca10  79 61 6e 69 63 6b 20 1b  5b 6d 1b 5b 33 33 6d 3e  |yanick .[m.[33m>|
```

5.2 Conclusion

So making up the conclusion it isn't possible to read the memory from inside the virtual machine (VM). So the memory the memory in Xen would no issue or nevertheless. I also post some questions about memory and Xen on the mailing list and the community of `www.xen.org` and no answers on my questions. The questions where:

- What happens with the memory when a virtual machine would be shutdown?
- Is the memory empty when a virtual machine get the memory allocated?

The strange thing is, if I ask more common stuff I will get a normally answer. Why? Is it because not many people no about memory in combination with Xen or is it because there is some bug or leak in Xen.

6 Bridging

Bridging is a technique that is used to forward packets in packet-switched computer networks. Bridging makes no assumptions about where something is locate in a network, it depends on broadcasting. If the device is located by broadcasting, the location is recorded in a Address Resolution Protocol (ARP) table where the Media Access Control (MAC) address is stored in combination with the IP address. Bridging serves only connections on a local area network.

In Xen bridging won't do anything else then what is describe here above. So bridging makes it possible for virtual machines (VM's) to communicate with each other of with a devices on the outside op the Xen host.

- Which bridging components are able to qualify for security research?

  Bridging would only be used in Xen for network, and maybe it is unnecessary if networks get smart intelligence inside virtual machines (VM's). [13]

- Is it possible to sniff or eavesdrop when bridging between DOM0 and DOM U, or when a virtual network interface would be bridged to a real network interface?

  No it isn't possible to sniff or eavesdrop, on one conditions and that would be that DOM0 also makes a part of bridges. If DOM0 makes part of a bridge, that also is used by virtual machine (VM) then you could read in the DOM) all the traffic of the virtual machine. The condition is that the network must build on a clean way and that every virtual machine and DOM0 gets his own virtual interfaces that makes part of a bridge, see figure 4.

- Can your send network packets from the virtual machine and let it seem that the Xen host has sent the packets?

  In basics you always see that the packets came from the Xen host and not from a virtual machine that is running on the Xen host. If you look in figure 5 there is ping form a virtual machine to an ip address on the outside of the Xen host. The result that in every packet the Media Access Control (MAC) address of the network card of the Xen host would be shown. In this case is the MAC address *00:15:C5:E1:43:41*. If you sniff at the DOM0, then you will see the real interface and that could be for example an xenbridge0.

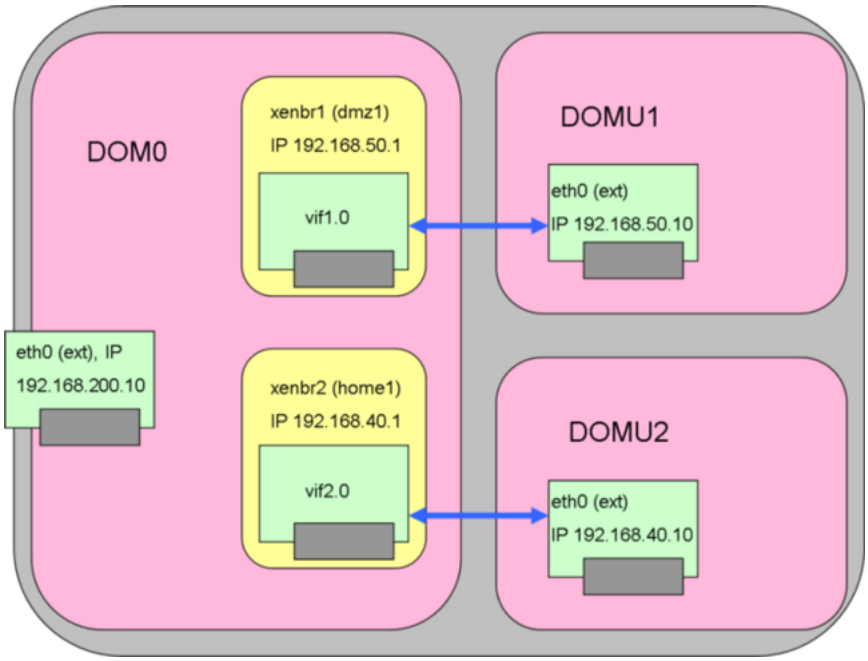

Figure 4: Bridging Network with Virtual Interfaces

6.1 Experiments & Results

By building and testing the networks there is a chance to discover that something is wrong or not safe with the networking in Xen. First there make a network where all VM's are on the same bridge and has connectivity with the Internet. The result is that all VM's, DOM0 can ping each other an also known for the outside world. It function as a normal network, when you using no virtualization.

So far so good, but take into account with that every one that manages DOM0 could read the traffic, because all traffic go through DOM 0. To prove it, in figure 6 you see traffic from the VM to bridge and you see also traffic from some the outside of the Xen host. For example 145.100.96.70 would be a web server.

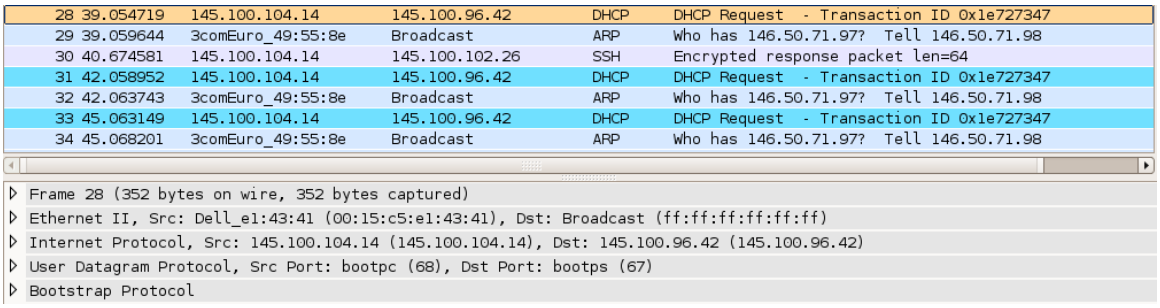

Figure 5: Packets came form Xen host

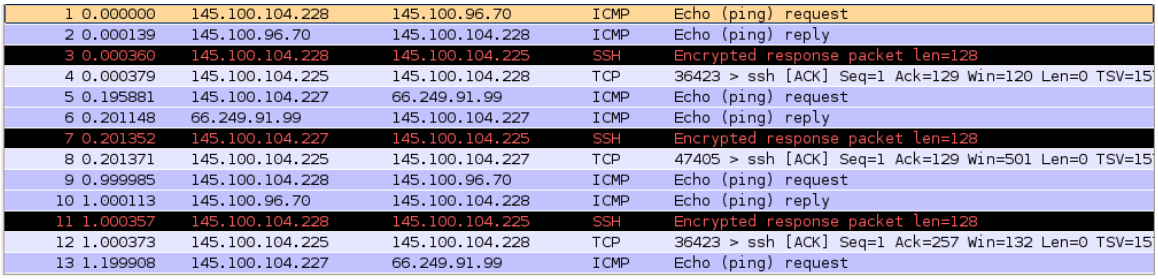

Figure 6: Ping from VM to bridge

In figure 7 you see some traffic of visiting a website that support Secure Socket Layer (SSL) support, and also this traffic you could see on the bridge.

Now people would think that the network would be an issue in Xen, but

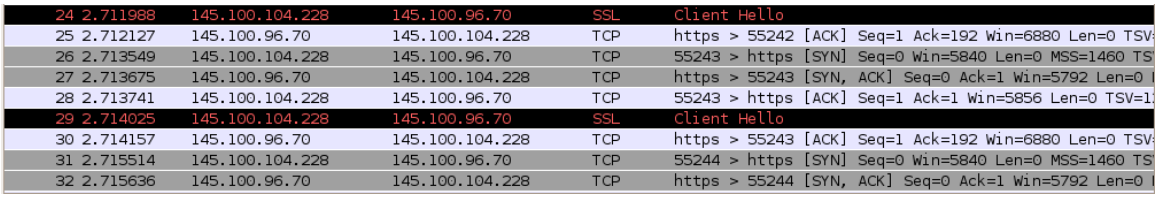

Figure 7: Visit website from inside VM

it doesn't. Xen engineers must configure the network on a different way. So they must use for every virtual machine a virtual interface, also for DOM0 to make a simple but same network. Other options are to use vlan's. Most companies using vlanning and can then also configure the Xen networking part so that a virtual machine has network on only one VM, for example the DMZ vlan.

6.2 Conclusion

If you configure the networking on the Xen host in the right way is it safe and you can minimize the risk to low. It would then be easier to take control op the host operating system, because you can manage the DOM0 and all the DOM U or virtual machines (VM) that are on the system. The fact that the Xen host for the virtual DMZ machine and for the virtual LAN server machine has open Internet connectivity would be worser.

7 DMA

Is it possible to rewrite to memory location if Direct Memory Access (DMA) is used?

It should be possible to rewrite the memory location if DMA is used and the consequences could be that you memory would be read through a Firewire port. There is no experiment done, because the little amount of time and the complexity of the problem. The fact that you can read your memory through a Firewire port is an issue on its own, what would described below But there are still problems with DMA in some cases in combination with Xen. [14]

7.1 Reading memory through Firewire

Reading and writing the memory of a computer by a Firewire port is a fact. It won't easy the secure the machine here for, because Firewire uses his own processor. The main processor and the operating system doesn't know anything about the the Firewire processor and what he is doing. This is all possible, because the Firewire could directly communicate with the Direct Memory Access (DMA) controller without any mediation of the operating system.

The Open Host Controller Interface (OHCI) is the controller of Firewire. The OHCI controller manages all the traffic that go through the Firewire cable and also communicates with the DMA controller. So the problem is that there is no detection in the operating system or some where in the hardware, that will gave an alarm if you reading the memory through the Firewire port. In figure 8 how the OHCI and DMA controller communicates with each other.

There is one condition why this principle works and that is, because it is possible that a laptop with Linux could behave it self as an iPod. Most operating systems trust an iPod, you are in and could every thing you like. [15][16][17][18]

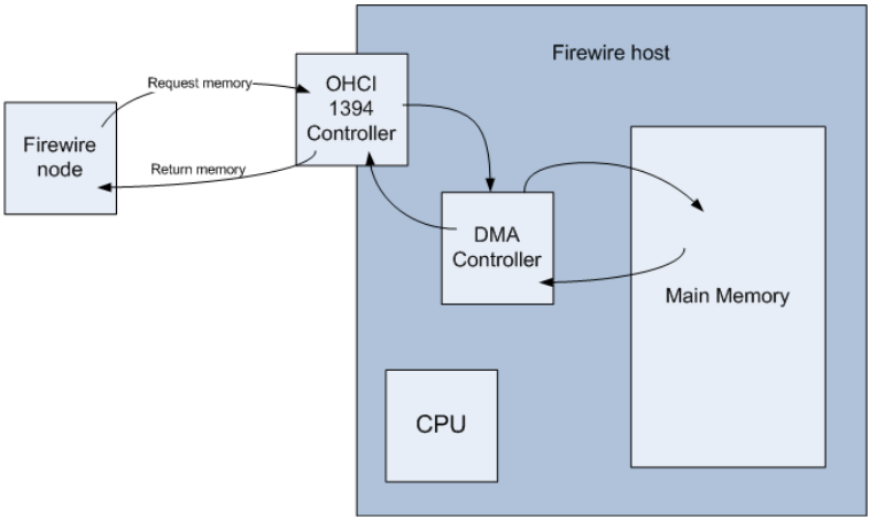

Figure 8: Firewire overview

7.2 Conclusion

It could be done by virtualized machines, but also by non-virtualized machines. So the risk is from the same scope, maybe something bigger, because you Xen host could be approach from the Internet. In the case if you want to merge the virtual machines to one Xen host it would not so important fact, because it now also there on a non virtualized system.

8 Final Conclusion

Finally conclusion is that it won't be safe for organizations to merge the virtual machines of the DMZ network segment and the LAN server segment into one segment on Xen hosts, because the following:

- You merge systems out of different networks segments with different security levels to one segment, and that has also consequences for the systems who are in that network segment.
- Virtualized systems would be a single point of failure, the risk will be the same as a normal system, but the impact raised enormous, because there are more virtual machine running on a host.

9 Acknowledgement

While conducting this research project, we received the generous help from the following people, Fred Mobach (Systemhouse Mobach)and Mendel Mobach (Systemhouse Mobach) to express my gratitude for their information, guidance and explanation. I would like to thank the education System and Network Engineer for the facility's, like the experimental machine and a workplace.

References

[1] Paravirtualization Full virtualization *Xensource Mailinglist* " http://lists.xensource.com/archives/html/xen-users/2007-07/msg00500.html "

[2] Datacenters steeds beter benut door virtualisatie *State of the Data Center Onderzoek door Symantec - Webwereld.nl* " http://new.webwereld.nl/article/view/id/54335 "

[3] Onverwachte voordelen van Server Virtualisatie Herkent u dit? *Presentation of Netwell consultancy* " http://www.netwellconsultancy.nl/presentaties/Onverwachtevoordelen.pdf "

[4] Professional Xen Virtualization *Book from William von Hagen*
ISBN: 978-0-470-13881-3

[5] Netwerkbeveiliging voor professionals *Book form Eric Maiwald, Chief Technology Officier, Fortrex Technologies, Inc.*
ISBN: 90-395-1865-3

[6] Network Infrastructure Security " http://system.vccs.edu/ITS/models/NetworkInfrastureSecurityModel.htm "

[7] Internet security and critical infrastructures " http://www.eurescom.de/message/messagesep2004/Internet_security_and_critical_infrastructure.asp "

[8] Designing Network Infrastructure Security " http://www.tech-faq.com/designing-network-infrastructure-security.shtml "

[9] Common Virtualization Information and Diagrams *Virtuatopia website* " http://www.virtuatopia.com/ "

[10] Memory Allocation *Linuxtopia - Xen 3.0 Virtualization Interface Guide* " http://www.linuxtopia.org/online_books/linux_virtualization/xen_3.0_interface_guide/linux_virualization_xen_interface_9.html "

[11] Pseudo-Physical Memory *Linuxtopia - Xen 3.0 Virtualization Interface Guide* " http://www.linuxtopia.org/online_books/linux_virtualization/xen_3.0_interface_guide/linux_virualization_xen_interface_10.html "

[12] DOM0 *DOM Wiki Xensource.com* " http://wiki.xensource.com/xenwiki/Dom0 "

[13] Development Roadmap for new version *New feautures and support for Xen* " http://www.xen.org/download/roadmap.html "

[14] [XEN-IOMMU] Proposal of DMA protection/isolation support *There are some issue with DMA, that would be could by a buggy driver* " https://lists.linux-foundation.org/pipermail/iommu/2008-January/000088.html "

[15] Paper - Reading RAM through firewire *Paper from Arthur Schutijser en Lili Poupa* Paper is op verzoek verkrijgbaar

[16] Information about Reading RAM through firewire *Website about Reading RAM through firewire* " http://www.storm.net.nz/projects/16 "

[17] Information about Reading RAM through firewire *Presentation about Reading RAM through firewire* " http://www.storm.net.nz/static/files/ab_firewire_rux2k6-final.pdf "

[18] Information about Reading RAM through firewire *Website about Reading RAM through firewire* " http://computer.forensikblog.de/en/2008/02/acquisition_5_firewire.html "