

Contents

- [1 4th Dec 2006 Xen Report](#)
- [2 08th Jan 2007 Xen Report](#)
- [3 16th Jan 2007 Xen Report](#)
 - ◆ [3.1 Conclusion to eth0
Crash/Hang at 4gb TX](#)
- [4 24th April 2007 Xen Report](#)
- [5 30th May 2007 - Xen Update](#)
- [6 03rd October 2009](#)

4th Dec 2006 Xen Report

Today after trying to kill a vnc session

```
vncserver -kill :1
```

The domU hung at 99.9% cpu (which was seen using `xm top` on dom0. After waiting a nice few minutes to see if it'd sort itself out - it didnt :-/, nor did it let me login via `xm console`. I then proceeded to:

```
xm shutdown domU_name
```

and gave it the usual few minutes. I then do `xm list` - it now reports:

Zombie-domU_name and listed the stats.

`xm destroy domU_name` returned - invalid integer. This was probably because the dying domain was named to **Zombie-domU**. Investigations on google led me to:

<http://lists.xensource.com/archives/html/xen-users/2005-12/msg00433.html>

I didnt really want to reboot dom0. I then read that even though `xm list` and `xm top` reported the **Zombie domU** using memory etc - it wasn't. I said feic it - i'll start up `domU_name` again and see.

```
xm create domU_name
```

it booted up fine. `xm list` still shows:

```
Zombie-domU_name          3          360          1 ---s-d 32864.8
```

`xm top` shows:

```
ds----- 32864 0.0          8 0.0 368640 35.2 1 1 4194300 66341
```

I checked for errors in `/var/log` and I got the following in `/var/log/xend.log`

```
[2006-12-04 17:57:55 xend.XendDomainInfo] DEBUG (XendDomainInfo:877) XendDomainInfo.handleShutdownWa
[2006-12-04 17:58:25 xend.XendDomainInfo] INFO (XendDomainInfo:867) Domain shutdown timeout expired:
[2006-12-04 17:58:25 xend.XendDomainInfo] DEBUG (XendDomainInfo:1327) XendDomainInfo.destroy: domid=
```

Ongoing_Experiences_with_Xen

```
[2006-12-04 17:58:25 xend.XendDomainInfo] DEBUG (XendDomainInfo:1335) XendDomainInfo.destroyDomain(6
[2006-12-04 18:07:31 xend.XendDomainInfo] DEBUG (XendDomainInfo:178) XendDomainInfo.create(['vm', ['
```

Anyways - thats how that ended up. I'll schedule a reboot for dom0 later on.

08th Jan 2007 Xen Report

Seems today that the same above box encountered another problem. This time - eth0 crashed. After transferring a substantial amount of network data, the server dropped off the network. Going in via "xm console vm" I noticed that there was no connectivity. There was still a lot of netstats -n but they were all TIME_WAIT. I decided to restart eth0:

```
ifdown -a
ifup -a
```

I got an error after trying to bring up eth0 saying:

```
eth0: full queue wasn't stopped!
```

eth0 wouldn't come back. I gave it a few minutes - same error. I decided to do a reboot. The virtual machine hung on reboot using 99% of cpu as identified via "xm top". After another few minutes - I decided to "xm shutdown domain" - to no avail. I done a "xm destroy domain" which put it as a Zombie - freeing up memory. I managed to get the machine back up working.

Other posts and info I found on it:

<http://lists.xensource.com/archives/html/xen-devel/2004-09/msg00113.html>

<http://blog.gmane.org/gmane.comp.emulators.xen.user/day=20061208>

Will see how it goes. Might try and take a look into it again when I have time.

16th Jan 2007 Xen Report

The same problem as occurred on 08th Jan, 2007 happended again, so I decided to investigate, and replicate on a different domU. Here is some background info firstly:

```
debian-host:~# uptime
12:48:19 up 99 days, 7:34, 1 user, load average: 0.09, 0.04, 0.07

xentop - 12:39:16 Xen 3.0.2-3
7 domains: 1 running, 3 blocked, 0 paused, 0 crashed, 1 dying, 2 shutdown
Mem: 1047100k total, 845988k used, 201112k free CPUs: 1 @ 2199MHz
  NAME STATE CPU(sec) CPU(%) MEM(k) MEM(%) MAXMEM(k) MAXMEM(%) VCPUS NETS NETTX(k) NETRX(k)
  ds---- 32864 0.0 8 0.0 368640 35.2 1 1 4194300 663
  d----- 18091 0.0 8 0.0 368640 35.2 1 1 4194302 3409
  ds---- 3208 0.0 8 0.0 368640 35.2 1 1 4194303 964
vm01 --b--- 134757 0.1 196404 18.8 196608 18.8 1 1 2094927 3453
Domain-0 -----r 34787 0.2 128008 12.2 no limit n/a 1 8 881476 420
vm03 --b--- 32980 0.2 130920 12.5 131072 12.5 1 1 4194303 2185
vm02 --b--- 178 0.0 368492 35.2 368640 35.2 1 1 4194303 91
```

Ongoing_Experiences_with_Xen

The dying or zombie domU's are from previous domU crashes as described further above.

vm02 and vm03 have no network connectivity. eth0 is crashed/hung etc. Notice also (in bold) that the NETTX(k) is **4194303**, and is common. Once the NETTX(k) reaches this figure of 4194303, eth0 on the domU dies/hangs.

```
xentop - 12:44:01 Xen 3.0.2-3
7 domains: 1 running, 3 blocked, 0 paused, 0 crashed, 1 dying, 2 shutdown
Mem: 1047100k total, 846180k used, 200920k free CPUs: 1 @ 2199MHz
NAME STATE CPU(sec) CPU(%) MEM(k) MEM(%) MAXMEM(k) MAXMEM(%) VCPUS NETS NETTX(k) NETRX(k)
ds----- 32864 0.0 8 0.0 368640 35.2 1 1 4194300 663
Net0 RX: 679333888bytes 21266499pkts 0err 1237drop TX: 4294963680bytes 19848474pkts
d----- 18091 0.0 8 0.0 368640 35.2 1 1 4194302 3409
Net0 RX: 3491791700bytes 9120614pkts 0err 1028drop TX: 4294966008bytes 7097985pkts
ds----- 3208 0.0 8 0.0 368640 35.2 1 1 4194303 964
Net0 RX: 987874486bytes 2811437pkts 0err 64drop TX: 4294966748bytes 3853122pkts
vm01 --b--- 134761 0.1 196476 18.8 196608 18.8 1 1 2095216 3454209
Net0 RX: 3537110202bytes 35480426pkts 0err 20drop TX: 2145501552bytes 42044159pkts
Domain-0 -----r 34788 0.2 128040 12.2 no limit n/a 1 8 881695 421
Net0 RX: 431354279bytes 5498978pkts 0err 0drop TX: 902856133bytes 3528712pkts
Net1 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
Net2 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
Net3 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
Net4 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
Net5 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
Net6 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
Net7 RX: 0bytes 0pkts 0err 0drop TX: 0bytes 0pkts
vm03 --b--- 32981 0.2 130860 12.5 131072 12.5 1 1 4194303 2186204
Net0 RX: 2238673582bytes 7363366pkts 0err 324drop TX: 4294966953bytes 6500733pkts
vm02 --b--- 179 0.0 368640 35.2 368640 35.2 1 1 4194303 91568
Net0 RX: 93766088bytes 1224399pkts 0err 672drop TX: 4294966578bytes 2890778pkts
```

The above shows the Network information within xm top. Definately there is a problem when the NETTX(k) reaches 4194303kb or 4294966953bytes. Lets look at some other information:

```
debian-host:~# ifconfig -s
Iface MTU Met RX-OK RX-ERR RX-DRP RX-OVR TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0 1500 0 5503676 0 0 0 3530737 0 0 0 BMRU
lo 16436 0 56 0 0 0 56 0 0 0 LRU
peth0 1500 0 71877410 0 0 0 078077723 0 0 0 BMORU
vif0. 1500 0 3530737 0 0 0 5503687 0 0 0 BMRU
vif4. 1500 0 42047643 0 0 0 035486757 0 28 0 BMRU
vif5. 1500 0 6500733 0 0 0 7366281 0 372 0 BMRU
vif9. 1500 0 2890778 0 0 0 1224399 0 3635 0 BMRU
xenbr 1500 0 944793 0 0 0 6 0 0 0 BMRU
```

There doesnt appear to be too much of benefit in the above. Lets move on and look at the ifconfig -s of the domU's where eth0 crashed:

```
vm02:~# uptime
12:41:26 up 17:58, 2 users, load average: 0.00, 0.00, 0.00
```

```
vm02:~# ifconfig -s
Iface MTU Met RX-OK RX-ERR RX-DRP RX-OVR TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0 1500 0 1222745 0 0 0 2891034 0 0 0 BMRU
lo 16436 0 2241 0 0 0 2241 0 0 0 LRU
```

```
vm03:~# uptime
12:47:20 up 87 days, 21:45, 2 users, load average: 0.00, 0.09, 0.27
```

```
vm03:~# ifconfig -s
```

Ongoing_Experiences_with_Xen

Iface	MTU	Met	RX-OK	RX-ERR	RX-DRP	RX-OVR	TX-OK	TX-ERR	TX-DRP	TX-OVR	Flg
eth0	1500	0	7365681	0	0	0	6500989	0	0	0	BMRU
lo	16436	0	6873	0	0	0	6873	0	0	0	LRU

Nothing appears out of the ordinary.

What I cant figure out is:

Is TX-OK (as in ifconfig -s on domU) meant to match with NETTX(k) (as in xm top for that domU)??

The other vm01 runs fine and has the following information:

```
vm01:~# uptime
12:49:46 up 91 days, 22:07, 5 users, load average: 0.00, 0.03, 0.04
```

```
vm01:~# ifconfig -s
Iface  MTU Met  RX-OK RX-ERR RX-DRP RX-OVR  TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0   1500 0    35488893 0 0 0 042048148 0 0 0 BMRU
lo     16436 0    14006676 0 0 0 014006676 0 0 0 LRU
```

Conclusion to eth0 Crash/Hang at 4gb TX

If the counter NETTX(k) as given in xm top, reaches 4194303, eth0 on that vm dies. This NETTX(k) does not seem to correlate with TX-OK as given in ifconfig -s on domU.

Does NETTX(k) as given in xm top, flush itself every so often? It does appear to be an incremental counter, however may reset or flush itself over a given period of time. If however a total of 4gb are downloaded from the domU in a short space of time (1-2hours), this NETTX(k) does not reset, reaches 4194303, and eth0 dies. I do realise that I am using Xen 3.0.2-3, which is not the current one. Xen 3.0.3 did not want to work with my SATA controller as far as I can remember. Currently the 3.0.2-3 I am using is from Debian Backports. When Debian Etch is finally released, I may convert over all.

Oh - also - the Exact same problem is described here:

<http://lists.xensource.com/archives/html/xen-users/2006-12/msg00252.html>

Email be at: sburke [at] burkesys.com if you have any heads-up. Thanks. -steve

24th April 2007 Xen Report

I installed Xen from Debian Etch. The 4gb problem is now fixed - Excellent. A huge thanks to the folks at Xen and Debian.

The how-to install Xen with Debian Etch is here: [Debian Etch Xen Install](#)

I tested the 4gb problem out. It was interesting watching the NETTX via "xm top". It hit 4GB, but instead of hanging etc. it reset itself back to 0 and started counting up again. There was a slight delay between NETTX hitting 4GB and the counter resetting to 0, but it worked seamlessly.

Note the Xen Kernel I am using with Debian Etch is: Xen 3.0.3-1 via linux-image-2.6-xen-686 kernel.

30th May 2007 - Xen Update

I recieved an email today from someone who experienced the same "eth0 Crash/Hang at 4gb TX" problem". They reported that this problem occurred for **both** Xen-3.0.2/Sarge and Xen-3.0.3/Etch!! Although I have tested Etch and Xen 3.0.3 using the methods previously described, thankfully, I have not experienced the same problem.

Patrik did however outline a temporary solution which worked. The details of which are:

```
ethtool -K eth0 tx off
#run the above on the domU
```

The original link/thread can be found here: <http://eagain.net/blog/2006/05/21/xen-tcp-hangs.html>

Hopefully, someone will find the above solution useful.

03rd October 2009

Xen is performing very well. Don't have much time to experiment these days. Anyways, the network lock-up happend on my own VM's. They were up and running fine for around 220 days, and then all of a sudden dropped off the net. Nagios informed me quickly. I had a VM running for a friend who had services running, and no init scripts (which I mentioned to him on several occasions to write). Anyways, instead of rebooting his VM, the following worked sweet as:

```
xm save vm-name /tmp/name-of-vm
xm list
#vm was taken down and saved. The size of the checkfile saved was the ram of the vm.
xm restore /tmp/name-of-vm
```

Happy Days. I must see if this will work across reboots of dom0.