
Subject: Syslog-ng on HN eats RAM and OOM Killer kills VE processess (CRITICAL)

Posted by [januszzz](#) on Thu, 30 Apr 2009 06:58:12 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

I have reproducible problem which takes my OpenVZ installation out of production.

I got HN with ~10 VEs. One of those VEs is syslog-ng which listens for remote connections and gathers all log messages from several dozens of machines, including some of those VEs.

This syslog-ng gathers information from HN too.

And here is the real problem:

- when I restart syslog-ng on the VE, the syslog-ng on HN silently starts to eat memory and after 6-12 hours or so eats every memory bit machine has (24GB plus 10GB swap).

Needless to say OOM killer enters then and kills all syslog-ng instances on every VE. Besides, it randomly kills some more programs, including critical MySQLs and even ssh.

Then machine gets back to life, but the services needs restarting.

I got two identical machines (thanks for vzmigrate tool, btw!) and on both machines all goes the same. The situation has happen 4th time, always when network admin gets his hands into his syslog-ng config file, and after when he restarts it (it looks like the trigger for eating ram is syslog restart in VE, but I cannot be certain).

I can clone machines and simulate this behaviour, but please let me know what should I catch so I can give some feedback.

Thanks, Janusz.

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processess (CRITICAL)

Posted by [kir](#) on Thu, 30 Apr 2009 08:04:11 GMT

[View Forum Message](#) <> [Reply to Message](#)

What is your UBC configuration?

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processess (CRITICAL)

Posted by [januszzz](#) on Thu, 30 Apr 2009 08:37:42 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi, sorry, here it goes:

```
OpenVZ2 conf # ls -l *conf | tail -n 12 | while read x; do vzcfgvalidate $x; done
```

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

Validation completed: success

OpenVZ2 conf #

uname -a

Linux OpenVZ2 2.6.18-028stab056-3 #4 SMP Thu Mar 12 16:18:18 CET 2009 x86_64 Intel(R)
Xeon(R) CPU E5430 @ 2.66GHz GenuineIntel GNU/Linux

Kernel Config included.

Janusz.

File Attachments

1) [config.gz](#), downloaded 266 times

2) [userbeancounters.txt](#), downloaded 336 times

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processes
(CRITICAL)

Posted by [kir](#) on Thu, 30 Apr 2009 08:40:57 GMT

[View Forum Message](#) <> [Reply to Message](#)

Thanks! And the output of vzmemcheck -v, please.

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processes
(CRITICAL)

Posted by [januszzz](#) on Thu, 30 Apr 2009 08:49:59 GMT

[View Forum Message](#) <> [Reply to Message](#)

Here it goes.

File Attachments

1) [vzmemcheck.txt](#), downloaded 334 times

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processess (CRITICAL)

Posted by [kir](#) on Thu, 30 Apr 2009 09:09:17 GMT

[View Forum Message](#) <> [Reply to Message](#)

januszzz wrote on Thu, 30 April 2009 12:49 Here it goes.

This pretty much answers it. Your mem+swap commitment is more than 100%, meaning you allow to allocate more RAM+swap than you actually have.

See http://wiki.openvz.org/UBC_systemwide_configuration which explains it in great details. In particular, it says "Commitment levels more than 1 means that the containers are guaranteed more memory than the system has. Such overcommitment is strongly not recommended, because in that case if all the memory is consumed, random applications, including the ones belonging to the host system, may be killed and the system may become inaccessible by ssh(1) and lose other important functionality". See, this precisely describes your situation.

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processess (CRITICAL)

Posted by [januszzz](#) on Thu, 30 Apr 2009 09:16:21 GMT

[View Forum Message](#) <> [Reply to Message](#)

Huh, I expect that is true.

When I had 16GB RAM I've used vzsplit -n8 and based on this I've created the configs.

Later, I've added more 8GB more and tuned the next configs accordingly, expecting that this won't fail.

EDIT: what is bothering me is that those machines should not use more than say, 8GB ram and the only process which is ballooning (it drinks 20-25 GB of RAM) is syslog-ng in HN. Why?

EDIT 2: As far as I understand this is more syslog-ng issue that it uses lots of ram. But while it goes in HN all of syslog-ng in each of the VE gets killed, am I right?

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processess (CRITICAL)

Posted by [kir](#) on Thu, 30 Apr 2009 09:31:31 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote: But while it goes in HN all of syslog-ng in each of the VE gets killed, am I right?

No.

What can happen though is if the init script of syslog-ng has invalid/bad logic (i.e when restarting it kills processes by name not by PID it obtained when the process was starting) then it can kill all

VE's syslog-ng processes. Solution to that problem is to fix initscript to save/use PIDs.

Subject: Re: Syslog-ng on HN eats RAM and OOM Killer kills VE processes (CRITICAL)

Posted by [januszzz](#) on Thu, 30 Apr 2009 09:36:55 GMT

[View Forum Message](#) < > [Reply to Message](#)

>> But while it goes in HN all of syslog-ng in each of the VE gets
>> killed, am I right?

> No.

Hmm, this is strange and exactly opposite to what I observe.

As I've already said, syslog-ng is ballooning in HN, not in a VE, and oom killer kills everything what is large, in all VEs. All mysql get killed, all sshd, all syslog-ng, all apaches, all Oracle listeners etc.

EDIT: COngrats to me, I've just edit the main post to clarify that this restart in VE of the syslog-ng triggers memory consuming in HN.

EDIT 2: So it looks like action taken inside VE silently touches a program in HN.
