

Hi,

I've just written a small dstat plugin to monitor CPU usage per VE. You can find the plugin in the dstat subversion repository linked from:

<http://dag.wieers.com/home-made/dstat/>

Dstat is much like vmstat, but modular and versatile. So you can extend it with whatever counters you want to visualize, next to other counters. This helps to relate counters and find/troubleshoot bottlenecks.

This plugin can help to find the most consuming VE, or visualize how VE's are matching up to each other (and the system CPU usage).

I don't know whether the nice-value in /proc/vz/vestat is ever used (maybe I shouldn't add it in the list of counters).

I also don't know what other information in /proc/vz would be interesting to visualize (over a timeframe). I'm willing to add more plugins in dstat for openvz if I know how to make more sense of the current counters (and their usefulness). Feedback from developers welcome :)

Visually it looks like this (without colors though):

```
[root@itmphys01 dstat]# ./dstat -M cpu,vz -f 5
Module dstat_vz is still experimental.
-----cpu0-usage-----cpu1-usage----- --ve-301-usage----ve-302-usage-
usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
 0 0 99 0 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
25 15 50 10 0 0: 1 1 76 22 0 1| 26 15 60 0: 0 0 100 0
31 18 12 40 0 0: 2 2 70 25 0 1| 33 19 49 0: 0 0 100 0
16 14 35 35 0 0: 0 1 53 46 0 0| 16 14 70 0: 0 0 100 0
 5 5 59 30 0 0: 16 28 30 26 0 0| 20 32 47 0: 0 0 100 0
 0 0 100 0 0 0: 3 5 88 5 0 0| 2 5 93 0: 0 0 100 0
 0 0 98 1 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
 0 0 100 0 0 0: 0 0 98 1 0 0| 0 0 100 0: 0 0 100 0
 0 0 96 4 0 0: 0 3 97 0 0 0| 0 0 100 0: 0 0 100 0
```

```
[root@itmphys01 dstat]# ./dstat -M cpu,vz -C total -f 5
Module dstat_vz is still experimental.
----total-cpu-usage---- --ve-301-usage----ve-302-usage-
usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
 0 0 97 2 0 0| 0 0 100 0: 0 0 100 0
```

```
19 12 61 8 0 0| 37 23 40 0: 0 0 100 0
12 8 41 38 0 0| 25 15 61 0: 0 0 100 0
14 14 35 37 0 0| 28 26 46 0: 0 0 100 0
4 10 84 3 0 0| 7 20 73 0: 0 0 100 0
0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
0 0 100 0 0 0| 0 0 100 0: 0 0 100 0
```

All feedback welcomed.

Kind regards,

-- dag wieers, dag@wieers.com, <http://dag.wieers.com/> --
[all I want is a warm bed and a kind word and unlimited power]

Subject: Re: Dstat plugin for OpenVZ CPU statistics

Posted by [jbravo](#) on Sat, 14 Oct 2006 15:43:04 GMT

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

Hi Dag,

Good work, its great.

> Hi,

>

> I've just written a small dstat plugin to monitor CPU usage per VE. You
> can find the plugin in the dstat subversion repository linked from:

>

> <http://dag.wieers.com/home-made/dstat/>

>

> Dstat is much like vmstat, but modular and versatile. So you can extend it
> with whatever counters you want to visualize, next to other counters. This
> helps to relate counters and find/troubleshoot bottlenecks.

>

> This plugin can help to find the most consuming VE, or visualize how VE's
> are matching up to each other (and the system CPU usage).

>

> I don't know whether the nice-value in /proc/vz/vestat is ever used (maybe
> I shouldn't add it in the list of counters).

>

> I also don't know what other information in /proc/vz would be interesting
> to visualize (over a timeframe). I'm willing to add more plugins in dstat
> for openvz if I know how to make more sense of the current counters (and
> their usefulness). Feedback from developers welcome :)

```

>
> Visually it looks like this (without colors though):
>
> [root@itmphys01 dstat]# ./dstat -M cpu,vz -f 5
> Module dstat_vz is still experimental.
> -----cpu0-usage-----cpu1-usage----- --ve-301-usage----ve-302-usage-
> usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
> 0 0 99 0 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 25 15 50 10 0 0: 1 1 76 22 0 1| 26 15 60 0: 0 0 100 0
> 31 18 12 40 0 0: 2 2 70 25 0 1| 33 19 49 0: 0 0 100 0
> 16 14 35 35 0 0: 0 1 53 46 0 0| 16 14 70 0: 0 0 100 0
> 5 5 59 30 0 0: 16 28 30 26 0 0| 20 32 47 0: 0 0 100 0
> 0 0 100 0 0 0: 3 5 88 5 0 0| 2 5 93 0: 0 0 100 0
> 0 0 98 1 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 100 0 0 0: 0 0 98 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 96 4 0 0: 0 3 97 0 0 0| 0 0 100 0: 0 0 100 0
>
> [root@itmphys01 dstat]# ./dstat -M cpu,vz -C total -f 5
> Module dstat_vz is still experimental.
> ----total-cpu-usage---- --ve-301-usage----ve-302-usage-
> usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 97 2 0 0| 0 0 100 0: 0 0 100 0
> 19 12 61 8 0 0| 37 23 40 0: 0 0 100 0
> 12 8 41 38 0 0| 25 15 61 0: 0 0 100 0
> 14 14 35 37 0 0| 28 26 46 0: 0 0 100 0
> 4 10 84 3 0 0| 7 20 73 0: 0 0 100 0
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 100 0 0 0| 0 0 100 0: 0 0 100 0
>
> All feedback welcomed.
>
> Kind regards,
> -- dag wieers, dag@wieers.com, http://dag.wieers.com/ --
> [all I want is a warm bed and a kind word and unlimited power]

```

Subject: Re: Dstat plugin for OpenVZ CPU statistics
Posted by [kir](#) on Mon, 16 Oct 2006 10:57:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dag,

Thank you! I actually use dstat on my Gentoo note.

Speaking of counters, /proc/user_beancounters is of interest to any

OpenVZ user. Format is described at http://wiki.openvz.org/proc/user_beancounters. "held" and "failcnt" columns are most "dynamic". You can use the values either directly, or make some consolidated figures based of formulae in wiki.

Other thing that might be of interest is `/proc/fairsched{,2}`. These files are from OpenVZ Fair CPU scheduler. The only problem is looks like the format is not documented (well, not counting the source code).

Also, Kirill Korotaev will give us some suggestions...

Dag Wieers wrote:

```
> Hi,
>
> I've just written a small dstat plugin to monitor CPU usage per VE. You
> can find the plugin in the dstat subversion repository linked from:
>
> http://dag.wieers.com/home-made/dstat/
>
> Dstat is much like vmstat, but modular and versatile. So you can extend it
> with whatever counters you want to visualize, next to other counters. This
> helps to relate counters and find/troubleshoot bottlenecks.
>
> This plugin can help to find the most consuming VE, or visualize how VE's
> are matching up to each other (and the system CPU usage).
>
> I don't know whether the nice-value in /proc/vz/vestat is ever used (maybe
> I shouldn't add it in the list of counters).
>
> I also don't know what other information in /proc/vz would be interesting
> to vizualise (over a timeframe). I'm willing to add more plugins in dstat
> for openvz if I know how to make more sense of the current counters (and
> their usefulness). Feedback from developers welcome :)
>
> Visually it looks like this (without colors though):
>
> [root@itmphys01 dstat]# ./dstat -M cpu,vz -f 5
> Module dstat_vz is still experimental.
> -----cpu0-usage-----cpu1-usage----- --ve-301-usage----ve-302-usage-
> usr sys idl wai hiq siq:usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
> 0 0 99 0 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 25 15 50 10 0 0: 1 1 76 22 0 1| 26 15 60 0: 0 0 100 0
> 31 18 12 40 0 0: 2 2 70 25 0 1| 33 19 49 0: 0 0 100 0
> 16 14 35 35 0 0: 0 1 53 46 0 0| 16 14 70 0: 0 0 100 0
> 5 5 59 30 0 0: 16 28 30 26 0 0| 20 32 47 0: 0 0 100 0
> 0 0 100 0 0 0: 3 5 88 5 0 0| 2 5 93 0: 0 0 100 0
> 0 0 98 1 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 100 0 0 0: 0 0 98 1 0 0| 0 0 100 0: 0 0 100 0
```

```
> 0 0 96 4 0 0: 0 3 97 0 0 0| 0 0 100 0: 0 0 100 0
>
> [root@itmphys01 dstat]# ./dstat -M cpu,vz -C total -f 5
> Module dstat_vz is still experimental.
> ----total-cpu-usage---- --ve-301-usage----ve-302-usage-
> usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 97 2 0 0| 0 0 100 0: 0 0 100 0
> 19 12 61 8 0 0| 37 23 40 0: 0 0 100 0
> 12 8 41 38 0 0| 25 15 61 0: 0 0 100 0
> 14 14 35 37 0 0| 28 26 46 0: 0 0 100 0
> 4 10 84 3 0 0| 7 20 73 0: 0 0 100 0
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
> 0 0 100 0 0 0| 0 0 100 0: 0 0 100 0
>
> All feedback welcomed.
>
> Kind regards,
> -- dag wieers, dag@wieers.com, http://dag.wieers.com/ --
> [all I want is a warm bed and a kind word and unlimited power]
```

Subject: Re: Dstat plugin for OpenVZ CPU statistics
Posted by [Dag Wieers](#) on Mon, 16 Oct 2006 14:24:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 16 Oct 2006, Kir Kolyshkin wrote:

```
> Dag Wieers wrote:
>
>> I've just written a small dstat plugin to monitor CPU usage per VE. You can
>> find the plugin in the dstat subversion repository linked from:
>>
>> http://dag.wieers.com/home-made/dstat/
>>
>> Dstat is much like vmstat, but modular and versatile. So you can extend it
>> with whatever counters you want to visualize, next to other counters. This
>> helps to relate counters and find/troubleshoot bottlenecks.
>>
>> This plugin can help to find the most consuming VE, or visualize how VE's
>> are matching up to each other (and the system CPU usage).
>
> Thank you! I actually use dstat on my Gentoo note.
```

Great :)

> Speaking of counters, /proc/user_beancounters is of interest to any OpenVZ
> user. Format is described at http://wiki.openvz.org/proc/user_beancounters.
> "held" and "failcnt" columns are most "dynamic". You can use the values either
> directly, or make some consolidated figures based of formulae in wiki.

Well, I was wondering how the beancounters could be of any use. I guess it may be useful to see when (and maybe where) they happen, but how would it be displayed the most useful ? An aggregated number of all failcnt (either global, or per VPS) ? That way one can see if any failcnts happen, but they'll have to look into the beancounters themselves to know exactly which ones happened.

The held-changes are mostly useless to display in a dstat-fashion. Unless maybe if you are a developer ? There's so much information that needs an in-depth knowledge.

> Other thing that might be of interest is /proc/fairsched{,2}. These files are
> from OpenVZ Fair CPU scheduler. The only problem is looks like the format is
> not documented (well, not counting the source code).
>
> Also, Kirill Korotaev will give us some suggestions...

By any means, if you know what is interesting and how it could be made visible on a single line (aggregated or not) we can make different openvz plugins depending on the use-case.

Kind regards,

-- dag wieers, dag@wieers.com, <http://dag.wieers.com/> --
[all I want is a warm bed and a kind word and unlimited power]

Subject: Re: Dstat plugin for OpenVZ CPU statistics

Posted by [kir](#) on Mon, 16 Oct 2006 15:27:03 GMT

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

Dag Wieers wrote:

> On Mon, 16 Oct 2006, Kir Kolyshkin wrote:

>

>

>> Dag Wieers wrote:

>>

>>

>>> I've just written a small dstat plugin to monitor CPU usage per VE. You can
>>> find the plugin in the dstat subversion repository linked from:

>>>

>>> <http://dag.wieers.com/home-made/dstat/>

>>>
>>> Dstat is much like vmstat, but modular and versatile. So you can extend it
>>> with whatever counters you want to visualize, next to other counters. This
>>> helps to relate counters and find/troubleshoot bottlenecks.
>>>
>>> This plugin can help to find the most consuming VE, or visualize how VE's
>>> are matching up to each other (and the system CPU usage).
>>>
>> Thank you! I actually use dstat on my Gentoo note.
>>
>
> Great :)
>
>
>
>> Speaking of counters, /proc/user_beancounters is of interest to any OpenVZ
>> user. Format is described at http://wiki.openvz.org/proc/user_beancounters.
>> "held" and "failcnt" columns are most "dynamic". You can use the values either
>> directly, or make some consolidated figures based of formulae in wiki.
>>
>
> Well, I was wondering how the beancounters could be of any use. I guess it
> may be useful to see when (and maybe where) they happen, but how would it
> be displayed the most useful ? An aggregated number of all failcnt (either
> global, or per VPS) ?
Makes much sense, especially if using red color to denote that number is
increasing. Makes more sense if used per-VPS.
> That way one can see if any failcnts happen, but
> they'll have to look into the beancounters themselves to know exactly
> which ones happened.
>
> The held-changes are mostly useless to display in a dstat-fashion.
Well, showing kmemsize or privvmpages or physpages is making the same
sense as showing output of, say, "free" in a loop.
> Unless
> maybe if you are a developer ? There's so much information that needs an
> in-depth knowledge.
>
>
>
>> Other thing that might be of interest is /proc/fairsched{,2}. These files are
>> from OpenVZ Fair CPU scheduler. The only problem is looks like the format is
>> not documented (well, not counting the source code).
>>
>> Also, Kirill Korotaev will give us some suggestions...
>>
>
> By any means, if you know what is interesting and how it could be made

> visible on a single line (aggregated or not) we can make different openvz
> plugins depending on the use-case.
>
> Kind regards,
> -- dag wieers, dag@wieers.com, http://dag.wieers.com/ --
> [all I want is a warm bed and a kind word and unlimited power]

Subject: Re: Dstat plugin for OpenVZ CPU statistics
Posted by [dev](#) on Tue, 17 Oct 2006 16:04:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dag,

I think the following statistics are helpful for OpenVZ:

- per VE UBC fail counters (failcnt in /proc/user_beancounters).
I would draw a delta between 2 screen updates and if it is non-zero draw it in red.
- per VE memory limit (kmemsize + privvmpages) and current usage (kmemsize + oomguarpages). Same, if kmemsize or privvmpages failcnt is increasing, then it can be drawn in red.
- per VE CPU usage. This one can be calculated from /proc/fairsched2 here are some details on this:
<http://forum.openvz.org/index.php?t=tree&th=479&mid=2771&&rev=&reveal=>
- per VE and global CPU and memory latency.

- per VE disk I/O stats (stats not available yet).
- per VE network stats (there is no generic stats yet).

Thanks,
Kirill

> Dag,
>
> Thank you! I actually use dstat on my Gentoo note.
>
> Speaking of counters, /proc/user_beancounters is of interest to any
> OpenVZ user. Format is described at
> http://wiki.openvz.org/proc/user_beancounters. "held" and "failcnt"
> columns are most "dynamic". You can use the values either directly, or
> make some consolidated figures based of formulae in wiki.
>
> Other thing that might be of interest is /proc/fairsched{,2}. These
> files are from OpenVZ Fair CPU scheduler. The only problem is looks like
> the format is not documented (well, not counting the source code).
>
> Also, Kirill Korotaev will give us some suggestions...


```

>
> Dag Wieers wrote:
>
>>Hi,
>>
>>I've just written a small dstat plugin to monitor CPU usage per VE. You
>>can find the plugin in the dstat subversion repository linked from:
>>
>> http://dag.wieers.com/home-made/dstat/
>>
>>Dstat is much like vmstat, but modular and versatile. So you can extend it
>>with whatever counters you want to visualize, next to other counters. This
>>helps to relate counters and find/troubleshoot bottlenecks.
>>
>>This plugin can help to find the most consuming VE, or visualize how VE's
>>are matching up to each other (and the system CPU usage).
>>
>>I don't know whether the nice-value in /proc/vz/vestat is ever used (maybe
>>I shouldn't add it in the list of counters).
>>
>>I also don't know what other information in /proc/vz would be interesting
>>to visualize (over a timeframe). I'm willing to add more plugins in dstat
>>for openvz if I know how to make more sense of the current counters (and
>>their usefulness). Feedback from developers welcome :)
>>
>>Visually it looks like this (without colors though):
>>
>>[root@itmphys01 dstat]# ./dstat -M cpu,vz -f 5
>>Module dstat_vz is still experimental.
>>-----cpu0-usage-----cpu1-usage----- --ve-301-usage----ve-302-usage-
>>usr sys idl wai hiq siq|usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
>> 0 0 99 0 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
>> 25 15 50 10 0 0: 1 1 76 22 0 1| 26 15 60 0: 0 0 100 0
>> 31 18 12 40 0 0: 2 2 70 25 0 1| 33 19 49 0: 0 0 100 0
>> 16 14 35 35 0 0: 0 1 53 46 0 0| 16 14 70 0: 0 0 100 0
>> 5 5 59 30 0 0: 16 28 30 26 0 0| 20 32 47 0: 0 0 100 0
>> 0 0 100 0 0 0: 3 5 88 5 0 0| 2 5 93 0: 0 0 100 0
>> 0 0 98 1 0 0: 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
>> 0 0 100 0 0 0: 0 0 98 1 0 0| 0 0 100 0: 0 0 100 0
>> 0 0 96 4 0 0: 0 3 97 0 0 0| 0 0 100 0: 0 0 100 0
>>
>>[root@itmphys01 dstat]# ./dstat -M cpu,vz -C total -f 5
>>Module dstat_vz is still experimental.
>>----total-cpu-usage---- --ve-301-usage----ve-302-usage-
>>usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic
>> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
>> 0 0 97 2 0 0| 0 0 100 0: 0 0 100 0
>> 19 12 61 8 0 0| 37 23 40 0: 0 0 100 0

```

```
>> 12 8 41 38 0 0|25 15 61 0: 0 0 100 0
>> 14 14 35 37 0 0|28 26 46 0: 0 0 100 0
>> 4 10 84 3 0 0| 7 20 73 0: 0 0 100 0
>> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
>> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
>> 0 0 99 1 0 0| 0 0 100 0: 0 0 100 0
>> 0 0 100 0 0 0| 0 0 100 0: 0 0 100 0
>>
```

>>All feedback welcomed.

>>

>>Kind regards,

>>-- dag wieers, dag@wieers.com, http://dag.wieers.com/ --

>>[all I want is a warm bed and a kind word and unlimited power]

Subject: Re: Dstat plugin for OpenVZ CPU statistics
Posted by [Dag Wieers](#) on Mon, 23 Oct 2006 07:25:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 17 Oct 2006, Kirill Korotaev wrote:

> I think the following statistics are helpful for OpenVZ:

>

> - per VE UBC fail counters (failcnt in /proc/user_beancounters).

> I would draw a delta between 2 screen updates and if it is non-zero

> draw it in red.

I wrote a new vzubc plugin for this and renamed the vz plugin to vzcpu.

It only displays the per VE failcnt or the total failcnt.

Feedback welcome.

```
[root@itmphys01 dstat]# ./dstat -c -Mvzcpu,vzubc
```

Module dstat_vzcpu is still experimental.

Module dstat_vzubc is still experimental.

```
----total-cpu-usage---- -total-ve-usage total
```

```
usr sys idl wai hiq siq|usr sys idl nic|_fcnt
```

```
0 0 98 1 0 0| 0 0 100 0|8471
```

```
0 0 100 0 0 0| 0 0 100 0| 0
```

```
0 0 100 0 0 0| 0 0 100 0| 0
```

```
0 0 97 4 0 0| 0 0 100 0| 0
```

```
1 0 100 0 0 0| 0 0 100 0| 9
```

```
0 0 97 2 0 0| 0 0 100 0| 3
```

```
0 0 100 0 0 0| 0 0 100 0| 0
```

```
[root@itmphys01 dstat]# ./dstat -c -Mvzcpu,vzubc -f
```

Module dstat_vzcpu is still experimental.

Module dstat_vzubc is still experimental.

```
-----cpu0-usage-----cpu1-usage----- --ve-301-usage----ve-302-usage- -301---302-
```

```
usr sys idl wai hiq siq:usr sys idl wai hiq siq|usr sys idl nic:usr sys idl nic|_fcnt:_fcnt
0 0 99 1 0 0: 0 0 98 2 0 0| 0 0 100 0: 0 0 100 0|8470 0
0 0 100 0 0 0: 0 0 100 0 0 0| 0 0 100 0: 0 0 100 0| 0 0
0 0 100 0 0 0: 1 0 99 0 0 0| 0 0 99 0: 0 0 100 0| 4 0
0 0 100 0 0 0: 0 0 100 0 0 0| 0 0 100 0: 0 0 100 0| 5 0
0 0 100 0 0 0: 0 0 100 0 0 0| 0 0 100 0: 0 0 100 0| 3 0
0 0 96 4 0 0: 1 1 91 7 0 0| 0 0 100 0: 0 0 100 0| 2 0
```

Non-zero is displayed in color.

- > - per VE memory limit (kmemsize + privvmpages) and current usage
- > (kmemsize + oomguarpages). Same, if kmemsize or privvmpages
- > failcnt is increasing, then it can be drawn in red.

How do I have to look at this ?

```
used = kmemsize + oomguarpages
total = kmemsize + privvmpages
```

```
percentage used : used * 100 / total
```

- > - per VE CPU usage. This one can be calculated from /proc/fairsched2
- > here are some details on this:
- > <http://forum.openvz.org/index.php?t=tree&th=479&mid=2771&&rev=&reveal=>

Is my current implementation wrong ? It seems to provide correct information. Except maybe for the nice-value. I have never seen it different from zero.

I also don't understand what the idle/strv/uptime values are.

- > - per VE and global CPU and memory latency.

I have no clue to what CPU and memory latency means. What is the value displayed in ? How to visualize ?

- > - per VE disk I/O stats (stats not available yet).
- > - per VE network stats (there is no generic stats yet).

I'm obviously interested in both as well.

Are there any other counters that proof useful, like eg. number of interrupts, number of context-switches, loadavg etc. Are things like that available from openvz ?

dstat needs some infrastructure to pass options/arguments to plugins. At this point that is not possible, limiting the user to specify either the total, or the per VE for `_all_` VEs. This of course is pretty annoying if you have 50 VEs on a system :)

I'm considering a syntax like this:

```
dstat -M vzcpu:total;301;302 -M vzubc:total;301;302
```

Kind regards,

-- dag wieers, dag@wieers.com, <http://dag.wieers.com/> --
[all I want is a warm bed and a kind word and unlimited power]
