

---

Subject: ioacct interpretation and unit

Posted by [Atosboy](#) on Tue, 19 May 2009 15:29:46 GMT

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

---

Hi,

We try to analyze the behavior of our system in OpenVZ environnement. In order to see it behavior, we periodically capture the content of /proc/bc/<veid>/ioacct file and try to interpret theses indicators.

We are from of several problems :

- + What represent each counter(read, write, vfs\_read, vfs\_write, ....)?
- + What is the unit for each counter(Kb, Block, ...)?

Thank for your help,

Pat

NB : For instance, the first snapshot is :

090515_14h48m06	read	1081344
090515_14h48m06	write	684281856
090515_14h48m06	dirty	684572672
090515_14h48m06	cancel	37818368
090515_14h48m06	missed	0
090515_14h48m06	syncs_total	0
090515_14h48m06	fsyncs_total	46
090515_14h48m06	fdatasyncs_total	1
090515_14h48m06	range_syncs_total	0
090515_14h48m06	syncs_active	0
090515_14h48m06	fsyncs_active	0
090515_14h48m06	fdatasyncs_active	0
090515_14h48m06	range_syncs_active	0
090515_14h48m06	vfs_reads	4737583
090515_14h48m06	vfs_read_chars	576083171
090515_14h48m06	vfs_writes	281962
090515_14h48m06	vfs_write_chars	829761534
090515_14h48m06	io_pbs	71

---

Subject: Re: ioacct interpretation and unit

Posted by [maratrus](#) on Wed, 20 May 2009 11:53:25 GMT

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

---

Hi,

to get full information you'd better look through kernel sources

I doubt that all of the output numbers help you very much but I'd like to explain a little bit "dirty"

and "write" numbers' meaning.

Writing to disk might be divided into two phases: writing to page and then writing to disk. "dirty" and "write" fields of /proc/bc/\$VEID/ioacct output correspond to the former and to the latter phases.

So, if a "dirty" field is continuously growing and "write" field is not (or much slower than "dirty") this means that i/o scheduler doesn't allow this particular VE to write to disk.

---

---

Subject: Re: ioacct interpretation and unit  
Posted by [Atosboy](#) on Tue, 26 May 2009 13:32:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Thank for your help,  
We only have analyzed read, write indicators which are documented by Linux Kernel documentation.

Pat

---