Subject: io-waits / openvz / network

Posted by arghbis on Tue, 23 Dec 2008 09:38:08 GMT

View Forum Message <> Reply to Message

Hi,

I've observed a strange behaviour on my openvz enabled servers. Most of the servers have a very high io-wait value in the cpu-usage graphs. The servers that have low io-wait percentage have only one or two VE, that are not exchanging data between them.

It is strange to notice that the more a server contains VEs that are communicating, the more the io-wait percentage is high.

On the biggest server, i can see in /proc/interrupts that eth0 is responsible for 1.5 10^9 interrupts on each cpu (8 cpus). It sounds abnormal to me.

Does anyone has an opinion on this topic?

The problem is that i want to be sure that these io-waits do not hurt too much the performances.

thanks for any advice

Subject: Re: io-waits / openvz / network

Posted by piavlo on Tue, 23 Dec 2008 21:08:25 GMT

View Forum Message <> Reply to Message

Regarding the /proc/interrupts thing I've compared it with non openvz kernels and they also have about the same high amplitude high number interrupts as my openvz systems.

Also If yo use MSI then you i'll see that all cpus are involved

equaly in handling interrupts, meaning you have 2.6.24 ovz kernel, since with 2.6.18 ovz even with MSI

enabled in the kernel most net interrupts are handled by one cpu using APIC (probaly only more recent kernels wich have CONFIG_ARCH_SUPPORTS_MSI fully utilze the MSI)

Also try online monitoring of iowait with mpstat -P ALL 1 100

also have alook at iostat to make sure that iowait is not due to disk io. iostat 1 100