

Hi,

Is there a guide somewhere that discusses performance tuning for OpenVZ?

I am running CentOS 6.4 on a Dell PowerEdge server with 1.3TB disk space in a RAID 10 configuration with 32GB of memory.

I run ISPConfig 3 3.0.5.2 which has five containers:

mail.example.com - mail server
7GB RAM and 3GB VSwap

rock.example.com - web server 1
7GB RAM and 3GB VSwap

pb.example.com - web server 2
4GB RAM and 2GB VSwap

ns1.example.com - name server 1
1GB RAM and 512MB VSwap

ns2.example.com - name server 2
1GB RAM and 512MB VSwap

In some of the reading I did, I thought I wouldn't have to tweak beancounter parameters if I used RAM and VSwap settings. Did I misunderstand my reading?

Each container is running Debian Squeeze.

I have successfully installed the system, and have made some adjustments of kmemsize and privvmpages so all software could be installed. My mail server seems to take the most resources.

My goal is to do my best to have these parameters set well before the servers starts getting loaded, and then fine tuning parameters as needed.

Are there any guides to help me do this?

Here is the /proc/user_beancounters for the mail server.

Version: 2.5

uid	resource	held	maxheld	barrier	limit	failcnt
1152:	kmemsize	15543268	20049920	28745400		29580328
	0					

lockedpages	0	4	2048	2048	0
privvmpages	238896	268499	262144	278528	
shmpages	438	438	21504	21504	0
dummy	0	0	0	0	0
numproc	92	117	240	240	0
physpages	146312	208039	0	1835008	
vmguarpages	0	0	33792	9223372036854775807	
oomguarpages	103501	112197	26112	9223372036854775807	
numtcpsock	22	40	360	360	0
numflock	21	22	188	206	0
numpty	1	15	16	16	0
numsiginfo	1	27	256	256	0
tcpsndbuf	313920	441520	1720320	2703360	
tcprcvbuf	360448	434968	1720320	2703360	
othersockbuf	191896	247104	1126080	2097152	
dgramrcvbuf	0	21800	262144	262144	
numothersock	158	182	200	200	0
dcachesize	1754782	3624960	3409920	3624960	
numfile	1195	1410	9312	9312	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	48	48	128	128	0

If there is anyone who can help provide guidance, I would really appreciate it.

Thank you,
Darin