

---

Posted by [umask](#) on Wed, 02 Jul 2008 20:47:05 GMT

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

---

```
cat /etc/redhat-release
CentOS release 5.2 (Final)
```

```
Linux host01.test.ru 2.6.18-53.1.19.el5.028stab053.14 #1 SMP Thu May 8 20:43:27 MSD 2008
i686 i686 i386 GNU/Linux
```

3. VPS 1100

```
ONBOOT="yes"
```

```
# UBC parameters (in form of barrier:limit)
KMEMSIZE="2147483647:2147483647"
LOCKEDPAGES="2147483647:2147483647"
PRIVVMPAGES="2147483647:2147483647"
SHMPAGES="2147483647:2147483647"
NUMPROC="2147483647:2147483647"
PHYSPPAGES="2147483647:2147483647"
VMGUARPPAGES="2147483647:2147483647"
OOMGUARPPAGES="2147483647:2147483647"
NUMTCPSOCK="2147483647:2147483647"
NUMFLOCK="2147483647:2147483647"
NUMPTY="2147483647:2147483647"
NUMSIGINFO="2147483647:2147483647"
TCPSNDBUF="2147483647:2147483647"
TCPRCVBUF="2147483647:2147483647"
OTHERSOCKBUF="2147483647:2147483647"
DGRAMRCVBUF="2147483647:2147483647"
NUMOTHERSOCK="2147483647:2147483647"
DCACHESIZE="2147483647:2147483647"
NUMFILE="2147483647:2147483647"
AVNUMPROC="180:180"
NUMIPTENT="2147483647:2147483647"
```

```
# Disk quota parameters (in form of softlimit:hardlimit)
DISKSPACE="100000000:120000000"
DISKINODES="10000000:15000000"
QUOTATIME="0"
```

```
# CPU fair sheduler parameter
```

CPUUNITS="500000"

IP\_ADDRESS="192.168.1.10"  
HOSTNAME="a.host01.test.ru"  
VE\_ROOT="/vz/root/\$VEID"  
VE\_PRIVATE="/vz/private/\$VEID"  
OSTEMPLATE="centos-5-i386-minimal"  
ORIGIN\_SAMPLE="vps.basic"  
NAMESERVER="192.168.1.1"  
SEARCHDOMAIN="testru"  
CPULIMIT="1000"

# Kernel sysctl configuration file for Red Hat Linux

#

# For binary values, 0 is disabled, 1 is enabled. See sysctl( and  
# sysctl.conf(5) for more details.

# Controls IP packet forwarding  
net.ipv4.ip\_forward = 0

# Controls source route verification  
net.ipv4.conf.default.rp\_filter = 1

# Do not accept source routing  
net.ipv4.conf.default.accept\_source\_route = 0

# Controls the System Request debugging functionality of the kernel  
kernel.sysrq = 0

# Controls whether core dumps will append the PID to the core filename  
# Useful for debugging multi-threaded applications  
kernel.core\_uses\_pid = 1

# Controls the use of TCP syncookies  
net.ipv4.tcp\_syncookies = 0

# Controls the maximum size of a message, in bytes  
kernel.msgmnb = 65536

# Controls the default maximum size of a message queue  
kernel.msgmax = 65536

```
# Controls the maximum shared segment size, in bytes
kernel.shmmax = 4294967295

# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 268435456

# On Hardware Node we generally need
# packet forwarding enabled and proxy arp disabled
net.ipv4.ip_forward = 1
net.ipv4.conf.default.proxy_arp = 0
# Enables source route verification
net.ipv4.conf.all.rp_filter = 1
# Enables the magic-sysrq key
kernel.sysrq = 1
# TCP Explicit Congestion Notification
#net.ipv4.tcp_ecn = 0
# we do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0

# Local port range
net.ipv4.ip_local_port_range = 8192 65535

# Netfilter connection tracking table size
net.ipv4.ip_conntrack_max = 258068

# For servers that receive many connections at the same time,
# the TIME-WAIT sockets for new connections can be reused.
# This is useful in Web servers etc. See also net.ipv4.tcp_tw_recycle.
net.ipv4.tcp_tw_reuse = 1

# Enable fast recycling of TIME-WAIT sockets status
net.ipv4.tcp_tw_recycle = 1

# Tune VM subsystem to use swap only as last resort
vm.swappiness = 1

# Limit of socket listen() backlog, known in userspace as SOMAXCONN.
# Defaults to 128. See also tcp_max_syn_backlog for additional tuning
# for TCP sockets.
net.core.somaxconn = 2048

# The maximum number of queued connection requests which have still not
# received an acknowledgement from the connecting client. If this
# number is exceeded, the kernel will begin dropping requests.
# The default value of 256 is increased to 1024 when the memory present
# in the system is adequate or greater (>= 128Mb), and reduced to 128
```

```
# for those systems with very low memory (<= 32Mb). It is recommended
# that if this needs to be increased above 1024, TCP_SYNQ_HSIZE in
# include/net/tcp.h be modified to keep
# TCP_SYNQ_HSIZE*16<=tcp_max_syn_backlog, and the kernel be recompiled.
net.ipv4.tcp_max_syn_backlog = 1024
```

```
# Maximum number of packets in the global input queue.
# for 1 GBit links recommended value near 3000
net.core.netdev_max_backlog = 2500
```

```
# prevent time wait bucket table overflow
net.ipv4.tcp_max_tw_buckets_ub = 129034
net.ipv4.tcp_max_tw_kmem_fraction = 384
```

```
# This sets the max OS receive buffer size for all types of connections.
net.core.rmem_max = 16777216
```

```
# This sets the max OS send buffer size for all types of connections.
net.core.wmem_max = 16777216
```

```
# This sets the default OS receive buffer size for all types of connections.
net.core.rmem_default = 65535
```

```
# This sets the default OS send buffer size for all types of connections.
net.core.wmem_default = 65535
```

```
# TCP Autotuning setting. "The tcp_mem variable defines how the TCP stack
# should behave when it comes to memory usage. ... The first value specified
# in the tcp_mem variable tells the kernel the low threshold. Below this
# point, the TCP stack do not bother at all about putting any pressure on the
# memory usage by different TCP sockets. ... The second value tells the
# kernel at which point to start pressuring memory usage down. ... The final
# value tells the kernel how many memory pages it may use maximally.
# If this value is reached, TCP streams and packets start getting dropped
# until we reach a lower memory usage again. This value includes all
# TCP sockets currently in use."
net.ipv4.tcp_mem = 16777216 16777216 16777216
```

```
# TCP Autotuning setting. "The first value tells the kernel the minimum
# receive buffer for each TCP connection, and this buffer is always allocated
# to a TCP socket, even under high pressure on the system. ... The second
# value specified tells the kernel the default receive buffer allocated for
# each TCP socket. This value overrides the /proc/sys/net/core/rmem_default
# value used by other protocols. ... The third and last value specified in
# this variable specifies the maximum receive buffer that can be allocated
# for a TCP socket."
net.ipv4.tcp_rmem = 4096 131072 16777216
```

```
# TCP Autotuning setting. "This variable takes 3 different values which holds
# information on how much TCP sendbuffer memory space each TCP socket has to
# use. Every TCP socket has this much buffer space to use before the buffer
# is filled up. Each of the three values are used under different conditions.
# ... The first value in this variable tells the minimum TCP send buffer
# space available for a single TCP socket. ... The second value in the variable
# tells us the default buffer space allowed for a single TCP socket to use.
# ... The third value tells the kernel the maximum TCP send buffer space."
net.ipv4.tcp_wmem = 4096 131072 16777216
```

```
# This will ensure that immediately subsequent connections use these values.
net.ipv4.route.flush=1
```

```
# RFC 2018 TCP Selective Acknowledgements
net.ipv4.tcp_sack = 0
```

```
# RFC 1323 TCP timestamps
net.ipv4.tcp_timestamps = 0
```

```
net.ipv4.tcp_sack = 1
net.ipv4.tcp_fack = 1
```

```
# Enable TCP behaviour conformant with RFC 1337. When disabled,
# if a RST is received in TIME_WAIT state, we close the socket
# immediately without waiting for the end of the TIME_WAIT period.
net.ipv4.tcp_rfc1337 = 1
```

iptables off).

```
cat /etc/nginx/nginx.conf
```

```
user  nginx nginx;
worker_processes 4;
```

```
worker_rlimit_nofile 16384;
```

```
error_log /var/log/nginx/error.log debug;
```

```
events {
    worker_connections 16384;
    use epoll;
}
```

```

http {
    include      /etc/nginx/mime.types;
    default_type text/plain;

    log_format main '$remote_addr - $remote_user [$time_local] "$request" '
        '"$status" $body_bytes_sent "$http_referer" '
        '"$http_user_agent" "$http_x_forwarded_for" '
        '$request_time "$upstream_addr" [$upstream_response_time]';

    log_format compat '$remote_addr - $remote_user [$time_local] "$request" '
        '"$status" $body_bytes_sent "$http_referer" '
        '"$http_user_agent" "$http_x_forwarded_for"';

    sendfile      on;
    tcp_nopush    on;
    tcp_nodelay    on;

    client_header_timeout 60;
    client_body_timeout  60;
    send_timeout        30;
    keepalive_timeout 0;

    reset_timedout_connection on;

    server {
        listen 80 default backlog=16384 rcvbuf=4096 sndbuf=4096 deferred;
        server_name .test.ru;

        error_log /var/log/nginx/test.ru_error.log debug;
        access_log /var/log/nginx/test.ru_access_main.log main;

        location /nginx_status {
            stub_status on;
            access_log off;
        }

        fastcgi_intercept_errors on;
        proxy_intercept_errors on;

        error_page 500 502 503 504 /50x_empty.html;
        error_page 400 401 402 403 404 405 /50x_empty.html;
        location = /50x_empty.html {
            root /home/nginx/htdocs;
        }

        location ~ \.(wml|php)$ {
            proxy_read_timeout 3;

```

```

        proxy_connect_timeout 3;
        proxy_pass http://127.0.0.1:8080;
        proxy_set_header X-Real-IP $remote_addr;
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_set_header Host $host;
    }
}
}

```

httpd-2.2.8/2.2.9.

```

<?
$max = 0;
for($i=0;$i<1000;$i++){
    $t = microtime(true);
    file_get_contents("http://192.168.1.10/nginx_status");
    $t = microtime(true)-$t;
    if ($t>$max) $max = $t;
}
echo $max;
?>

```

```
while : ; do php dummy.php ; done | grep -e "[1-9]\.[0-9]"
```

```

3.0010089874268
3.0013828277588
3.001168012619
3.0015661716461
3.0009059906006
3.0006580352783
3.0018539428711
3.0014488697052
3.0009009838104

```

cat /proc/user\_beancounters

Version: 2.5

uid	resource	held	maxheld	barrier	limit	failcnt	
1100:	kmemsize	4334631	9382841	2147483647	2147483647		0
	lockedpages	0	0	2147483647	2147483647	0	
	privvmpages	245588	246632	2147483647	2147483647		0
	shmpages	1	1	2147483647	2147483647	0	
	dummy	0	0	0	0		
	numproc	118	122	2147483647	2147483647	0	
	physpages	34130	34505	2147483647	2147483647		0
	vmguarpages	0	0	2147483647	2147483647	0	
	oomguarpages	34130	34505	2147483647	2147483647		0
	numtcpsock	14	16	2147483647	2147483647	0	
	numflock	1	2	2147483647	2147483647	0	
	numpty	0	1	2147483647	2147483647	0	
	numsiginfo	0	2	2147483647	2147483647	0	
	tcpsndbuf	125216	125216	2147483647	2147483647		0
	tcprcvbuf	229376	229376	2147483647	2147483647		0
	othersockbuf	11180	13416	2147483647	2147483647		0
	dgramrcvbuf	0	0	2147483647	2147483647	0	
	numothersock	11	13	2147483647	2147483647		0
	dcachesize	0	0	2147483647	2147483647	0	
	numfile	2082	2158	2147483647	2147483647	0	
	dummy	0	0	0	0		
	dummy	0	0	0	0		
	dummy	0	0	0	0		
	numiptent	10	10	2147483647	2147483647	0	
0:	kmemsize	3298695	16916897	2147483647	2147483647		0
	lockedpages	1083	1083	2147483647	2147483647		0
	privvmpages	12459	14060	2147483647	2147483647		0
	shmpages	656	672	2147483647	2147483647		0
	dummy	0	0	2147483647	2147483647	0	
	numproc	73	85	2147483647	2147483647	0	
	physpages	4553	4994	2147483647	2147483647		0
	vmguarpages	0	0	2147483647	2147483647	0	
	oomguarpages	4553	4994	2147483647	2147483647		0
	numtcpsock	3	3	2147483647	2147483647	0	
	numflock	4	5	2147483647	2147483647	0	
	numpty	1	1	2147483647	2147483647	0	



numsiginfo	0	2 2147483647 2147483647	0
tcpsndbuf	35724	35724 2147483647 2147483647	0
tcprcvbuf	49152	32768 2147483647 2147483647	0
othersockbuf	154284	161420 2147483647 2147483647	0
dgramrcvbuf	0	8380 2147483647 2147483647	0
numothersock	122	126 2147483647 2147483647	0
dcachesize	0	0 2147483647 2147483647	0
numfile	1531	1771 2147483647 2147483647	0
dummy	0	0 2147483647 2147483647	0
dummy	0	0 2147483647 2147483647	0
dummy	0	0 2147483647 2147483647	0
numiptent	10	10 2147483647 2147483647	0

---

Posted by [umask](#) on Thu, 03 Jul 2008 06:52:39 GMT

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

---

eth0    Link encap:Ethernet HWaddr 00:15:17:1E:DF:6D  
 inet addr:192.168.1.1 Bcast:192.168.1.255 Mask:255.255.255.0  
 inet6 addr: fe80::215:17ff:fe1e:df6d/64 Scope:Link  
 UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
 RX packets:316696 errors:0 dropped:0 overruns:0 frame:0  
 TX packets:256175 errors:0 dropped:0 overruns:0 carrier:0  
 collisions:0 txqueuelen:1000  
 RX bytes:23930105 (22.8 MiB) TX bytes:157922369 (150.6 MiB)  
 Base address:0x2000 Memory:48180000-481a0000

lo     Link encap:Local Loopback  
inet addr:127.0.0.1 Mask:255.0.0.0  
inet6 addr: ::1/128 Scope:Host  
UP LOOPBACK RUNNING MTU:16436 Metric:1  
RX packets:0 errors:0 dropped:0 overruns:0 frame:0  
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

venet0   Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00  
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1  
RX packets:254038 errors:0 dropped:0 overruns:0 frame:0  
TX packets:254075 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:152950314 (145.8 MiB) TX bytes:15203174 (14.4 MiB)

eth0     Link encap:Ethernet HWaddr 00:04:23:B3:01:F6  
inet addr:192.168.1.20 Bcast:192.168.1.255 Mask:255.255.255.0  
inet6 addr: fe80::204:23ff:feb3:1f6/64 Scope:Link  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:195583406 errors:0 dropped:0 overruns:0 frame:0  
TX packets:195452910 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:1000  
RX bytes:1762512072 (1.6 GiB) TX bytes:2912992402 (2.7 GiB)  
Base address:0xdc00 Memory:fcfc0000-fcfe0000

lo     Link encap:Local Loopback  
inet addr:127.0.0.1 Mask:255.0.0.0  
inet6 addr: ::1/128 Scope:Host  
UP LOOPBACK RUNNING MTU:16436 Metric:1  
RX packets:8 errors:0 dropped:0 overruns:0 frame:0  
TX packets:8 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:704 (704.0 b) TX bytes:704 (704.0 b)

---

---

Posted by [maratrus](#) on Thu, 03 Jul 2008 13:37:29 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

---

---

Posted by [umask](#) on Thu, 03 Jul 2008 13:59:09 GMT

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

---

```
./httperf --server 192.168.1.1 --port 80 --uri=/index.html --num-conns 1000000 --rate 15000
```

--timeout 0.1

httpperf --timeout=0.1 --client=0/1 --server=192.168.1.1 --port=80 --uri=/index.html --rate=15000  
--send-buffer=4096 --recv-buffer=16384 --num-conns=1000000 --num-calls=1  
Maximum connect burst length: 30

Total: connections 1000000 requests 1000000 replies 1000000 test-duration 66.760 s

Connection rate: 14979.0 conn/s (0.1 ms/conn, <=269 concurrent connections)

Connection time [ms]: min 0.5 avg 3.1 max 17.8 median 2.5 stddev 1.9

Connection time [ms]: connect 0.4

Connection length [replies/conn]: 1.000

Request rate: 14979.0 req/s (0.1 ms/req)

Request size [B]: 75.0

Reply rate [replies/s]: min 14972.3 avg 14979.1 max 14987.6 stddev 4.9 (13 samples)

Reply time [ms]: response 2.7 transfer 0.0

Reply size [B]: header 209.0 content 4.0 footer 0.0 (total 213.0)

Reply status: 1xx=0 2xx=1000000 3xx=0 4xx=0 5xx=0

CPU time [s]: user 14.09 system 51.83 (user 21.1% system 77.6% total 98.7%)

Net I/O: 4212.8 KB/s (34.5\*10<sup>6</sup> bps)

Errors: total 0 client-timo 0 socket-timo 0 connrefused 0 connreset 0

Errors: fd-unavail 0 addrunavail 0 ftab-full 0 other 0

--cpuunits num

CPU weight for a VE. Argument is positive non-zero number, which passed to and used in kernel fair scheduler.

The larger the number is, the more CPU time this VE get. Maximum value is 500000, minimal is 8. Number is

relative to weights of all the other running VEs. If cpuunits not specified default value 1000 is used.

You can set CPU weight for VE0 (hardware node itself) as well (use `vzctl set 0 --cpuunits num`). Usually,

OpenVZ initscript (/etc/init.d/vz) takes care of setting this.

---

---

Posted by [maratrus](#) on Thu, 03 Jul 2008 14:16:44 GMT

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

---

---

---

Posted by [maratrus](#) on Thu, 03 Jul 2008 14:49:23 GMT

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

---

veth?

---

---

---

Posted by [umask](#) on Thu, 03 Jul 2008 15:03:11 GMT

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

---

veth?

venet?

---

---

Posted by [maratrus](#) on Thu, 03 Jul 2008 15:26:36 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quote:

Quote:

Quote:

---

---

Posted by [umask](#) on Thu, 03 Jul 2008 18:39:45 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

---

---

Posted by [maratrus](#) on Fri, 04 Jul 2008 08:22:16 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

---

---

Posted by [umask](#) on Fri, 04 Jul 2008 20:34:21 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

\*1\*

netstat -s c (1)

Ip:

- 3613925 total packets received
- 4 with invalid headers
- 134 forwarded
- 0 incoming packets discarded
- 3613787 incoming packets delivered
- 3613790 requests sent out

Icmp:

- 2 ICMP messages received
- 0 input ICMP message failed.
- ICMP input histogram:
  - echo requests: 2
- 9 ICMP messages sent
- 0 ICMP messages failed
- ICMP output histogram:
  - destination unreachable: 3
  - time exceeded: 4
  - echo replies: 2

Tcp:

- 602193 active connections openings
- 6 passive connection openings
- 1 failed connection attempts
- 1 connection resets received
- 2 connections established
- 3613655 segments received
- 3613536 segments send out
- 28 segments retransmitted
- 0 bad segments received.
- 10 resets sent

Udp:

- 82 packets received
- 3 packets to unknown port received.
- 0 packet receive errors
- 83 packets sent

TcpExt:

- 1 resets received for embryonic SYN\_RECV sockets
- 3 TCP sockets finished time wait in fast timer
- 7 delayed acks sent
- 81 packets directly queued to recvmsg prequeue.
- 24 packets directly received from prequeue
- 602278 packets header predicted
- 1204446 acknowledgments not containing data received



- 602423 predicted acknowledgments
- 1 times recovered from packet loss due to SACK data
- 10 congestion windows recovered after partial ack
- 1 TCP data loss events
- 1 fast retransmits
- 27 other TCP timeouts
- 1 DSACKs sent for old packets
- 2 DSACKs received
- 1 connections reset due to unexpected data
- 1 connections reset due to early user close

netstat -s c (2)

Ip:

- 7227192 total packets received
- 7226305 forwarded
- 0 incoming packets discarded
- 854 incoming packets delivered
- 7226920 requests sent out

Icmp:

- 9 ICMP messages received
- 0 input ICMP message failed.
- ICMP input histogram:
  - echo requests: 9
- 9 ICMP messages sent
- 0 ICMP messages failed
- ICMP output histogram:
  - echo replies: 9

Tcp:

- 9 active connections openings
- 9 passive connection openings
- 0 failed connection attempts
- 0 connection resets received
- 3 connections established
- 798 segments received
- 548 segments send out
- 9 segments retransmitted
- 0 bad segments received.
- 0 resets sent

Udp:

- 77 packets received
- 0 packets to unknown port received.
- 0 packet receive errors
- 79 packets sent

TcpExt:

- 10 TCP sockets finished time wait in fast timer
- 601948 TCP sockets finished time wait in slow timer

12 delayed acks sent  
65 packets directly queued to recvmmsg prequeue.  
17 packets directly received from prequeue  
179 packets header predicted  
59 acknowledgments not containing data received  
347 predicted acknowledgments  
4 congestion windows recovered after partial ack  
0 TCP data loss events  
8 other TCP timeouts  
2 DSACKs sent for old packets  
5 DSACKs received

netstat -s c (3)

Ip:

3613147 total packets received  
0 forwarded  
0 incoming packets discarded  
3613147 incoming packets delivered  
3613158 requests sent out

Icmp:

4 ICMP messages received  
0 input ICMP message failed.  
ICMP input histogram:  
    timeout in transit: 4  
4 ICMP messages sent  
0 ICMP messages failed  
ICMP output histogram:  
    destination unreachable: 4

Tcp:

9 active connections openings  
602198 passive connection openings  
4 failed connection attempts  
0 connection resets received  
0 connections established  
3613169 segments received  
3613184 segments send out  
0 segments retransmitted  
0 bad segments received.  
0 resets sent

Udp:

0 packets received  
4 packets to unknown port received.  
0 packet receive errors  
0 packets sent

TcpExt:

1 delayed acks sent

602191 packets header predicted  
602191 acknowledgments not containing data received  
0 TCP data loss events

\*2\*

NIC statistics:

rx\_packets: 3615960  
tx\_packets: 3613901  
rx\_bytes: 411792822  
tx\_bytes: 289840057  
rx\_errors: 0  
tx\_errors: 0  
tx\_dropped: 0  
multicast: 2  
collisions: 0  
rx\_length\_errors: 0  
rx\_over\_errors: 0  
rx\_crc\_errors: 0  
rx\_frame\_errors: 0  
rx\_no\_buffer\_count: 0  
rx\_missed\_errors: 0  
tx\_aborted\_errors: 0  
tx\_carrier\_errors: 0  
tx\_fifo\_errors: 0  
tx\_heartbeat\_errors: 0  
tx\_window\_errors: 0  
tx\_abort\_late\_coll: 0  
tx\_deferred\_ok: 0  
tx\_single\_coll\_ok: 0  
tx\_multi\_coll\_ok: 0  
tx\_timeout\_count: 0  
rx\_long\_length\_errors: 0  
rx\_short\_length\_errors: 0  
rx\_align\_errors: 0  
tx\_tcp\_seg\_good: 20  
tx\_tcp\_seg\_failed: 0  
rx\_flow\_control\_xon: 0  
rx\_flow\_control\_xoff: 0  
tx\_flow\_control\_xon: 0  
tx\_flow\_control\_xoff: 0  
rx\_long\_byte\_count: 411792822  
rx\_csum\_offload\_good: 3613986  
rx\_csum\_offload\_errors: 0  
rx\_header\_split: 0  
alloc\_rx\_buff\_failed: 0

NIC statistics:

rx\_packets: 3616218  
tx\_packets: 3613918  
rx\_bytes: 290005297  
tx\_bytes: 411677781  
rx\_broadcast: 2006  
tx\_broadcast: 7  
rx\_multicast: 2  
tx\_multicast: 6  
rx\_errors: 0  
tx\_errors: 0  
tx\_dropped: 0  
multicast: 2  
collisions: 0  
rx\_length\_errors: 0  
rx\_over\_errors: 0  
rx\_crc\_errors: 0  
rx\_frame\_errors: 0  
rx\_no\_buffer\_count: 0  
rx\_missed\_errors: 0  
tx\_aborted\_errors: 0  
tx\_carrier\_errors: 0  
tx\_fifo\_errors: 0  
tx\_heartbeat\_errors: 0  
tx\_window\_errors: 0  
tx\_abort\_late\_coll: 0  
tx\_deferred\_ok: 0  
tx\_single\_coll\_ok: 0  
tx\_multi\_coll\_ok: 0  
tx\_timeout\_count: 0  
tx\_restart\_queue: 0  
rx\_long\_length\_errors: 0  
rx\_short\_length\_errors: 0  
rx\_align\_errors: 0  
tx\_tcp\_seg\_good: 1  
tx\_tcp\_seg\_failed: 0  
rx\_flow\_control\_xon: 0  
rx\_flow\_control\_xoff: 0  
tx\_flow\_control\_xon: 0  
tx\_flow\_control\_xoff: 0  
rx\_long\_byte\_count: 290005297  
rx\_csum\_offload\_good: 3614184  
rx\_csum\_offload\_errors: 0  
rx\_header\_split: 0

alloc\_rx\_buff\_failed: 0  
tx\_smbus: 0  
rx\_smbus: 0  
dropped\_smbus: 0

\*3\*

```
while : ; do php dummy.php ; echo ; done | grep -e "[1-9]\.[0-9]"  
3.0010950565338  
3.0016250610352  
3.0018901824951  
3.00084400177  
3.0018458366394  
3.0011219978333  
3.0014960765839  
3.001119852066
```

[http://people.mobiledirect.ru/people/umask/public/ovz\\_proble ms/](http://people.mobiledirect.ru/people/umask/public/ovz_proble ms/)

HN\_where\_test\_running.out[.gz] - (1)

---

---

Posted by [umask](#) on Sat, 05 Jul 2008 05:07:31 GMT

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

---

veth.

---

---

Posted by [umask](#) on Tue, 08 Jul 2008 08:47:22 GMT

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

---

---

---

Posted by [maratrus](#) on Tue, 08 Jul 2008 10:24:47 GMT

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

---

---

---

Posted by [umask](#) on Tue, 08 Jul 2008 11:00:03 GMT

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

---

---

---

Posted by [hiddenman](#) on Fri, 20 Jul 2012 08:03:07 GMT

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

---

