
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"
PHYS_PAGES="2147483647:2147483647"
VMGUAR_PAGES="2147483647:2147483647"
OOMGUAR_PAGES="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