

Sysctl: Тюним сетевые настройки Linux

Вариант 1

Источник: https://github.com/telemt/telemt/blob/main/docs/Advanced_settings/HIGH_LOAD.ru.md

При разворачивании сервера, работающего под высокой нагрузкой (десятки и сотни тысяч одновременных подключений), стандартные ограничения сетевого стека ОС могут приводить к потерям пакетов, переключениям контекста CPU и отказам в соединениях. В данном руководстве описана настройка ядра Linux, системных лимитов и аппаратной конфигурации для работы в подобных сценариях.

Добавим в файл `/etc/sysctl.conf`:

```
fs.file-max = 2097152
fs.nr_open = 2097152
```

Создадим файл `/etc/sysctl.d/99-network-highload.conf`:

```
net.core.somaxconn = 65535
net.core.netdev_max_backlog = 65535
net.ipv4.tcp_max_syn_backlog = 65535
net.ipv4.tcp_syncookies = 1

net.ipv4.ip_local_port_range = 10000 65535
net.ipv4.tcp_fin_timeout = 15
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_max_tw_buckets = 2000000

net.ipv4.tcp_keepalive_time = 300
net.ipv4.tcp_keepalive_intvl = 30
net.ipv4.tcp_keepalive_probes = 5

# Размеры буферов ядра (по умолчанию и макс)
net.core.rmem_default = 262144
net.core.wmem_default = 262144
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216

# Специфичные TCP буферы (min, default, max)
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216

# Включение BBR
net.core.default_qdisc = fq
net.ipv4.tcp_congestion_control = bbr
```

После внесения правок нужно применить изменения командой:

```
sysctl -p
sysctl -p /etc/sysctl.d/99-network-highload.conf
```

Вариант 2

Источник: <https://www.xodice.org/2025/11/06/modern-linux-distribution-defaults-are-stuck-in-2009/>

Предупреждение: не добавляйте все кучей в свою систему, ситуации могут быть разными!

```
# cat /etc/sysctl.d/99-sysctl.conf
# AMD Ryzen 9 7950X, 128 GB DDR5 ECC, dual 1.92 TB U.2 NVMe, full dual-stack
IPv4/IPv6
#####
# IPv4: security + routing
#####
net.ipv4.icmp_timestamps = 0
net.ipv4.icmp_echo_ignore_broadcasts = 1
net.ipv4.icmp_ignore_bogus_error_responses = 1
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.all.secure_redirects = 0
net.ipv4.conf.default.secure_redirects = 0
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.default.send_redirects = 0

# VPN-friendly reverse path filtering
net.ipv4.conf.all.rp_filter = 2
net.ipv4.conf.default.rp_filter = 2

# TCP & SYN flood aggressive tuning
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_max_syn_backlog = 65536
net.ipv4.tcp_rfc1337 = 1
net.ipv4.tcp_mtu_probing = 1
net.ipv4.tcp_keepalive_time = 300
net.ipv4.tcp_keepalive_intvl = 30
net.ipv4.tcp_keepalive_probes = 5
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
net.ipv4.tcp_max_tw_buckets = 100000
net.ipv4.tcp_max_orphans = 65536

net.ipv4.conf.all.log_martians = 1
net.ipv4.conf.default.log_martians = 1

# WireGuard: enable IPv4 forwarding
```

```
net.ipv4.ip_forward = 1

#####
# IPv6: security + routing
#####
net.ipv6.conf.all.accept_source_route = 0
net.ipv6.conf.default.accept_source_route = 0
net.ipv6.conf.all.accept_redirects = 0
net.ipv6.conf.default.accept_redirects = 0
net.ipv6.conf.all.accept_ra = 0
net.ipv6.conf.default.accept_ra = 0
net.ipv6.conf.all.autoconf = 0
net.ipv6.conf.default.autoconf = 0
net.ipv6.conf.all.forwarding = 1
net.ipv6.conf.default.forwarding = 1
net.ipv6.conf.enp7s0.forwarding = 1
net.ipv6.conf.enp7s0.accept_ra = 0
net.ipv6.conf.enp7s0.autoconf = 0

#####
# General kernel hardening
#####
kernel.randomize_va_space = 2
kernel.kptr_restrict = 2
kernel.dmesg_restrict = 1
kernel.perf_event_paranoid = 3
kernel.kexec_load_disabled = 1
kernel.yama.ptrace_scope = 1

# Filesystem protections
fs.protected_symlinks = 1
fs.protected_hardlinks = 1
fs.protected_fifos = 2
fs.protected_regular = 2
fs.suid_dumpable = 0

# eBPF restrictions
kernel.unprivileged_bpf_disabled = 1
net.core.bpf_jit_harden = 2

# Disable SysRq entirely
kernel.sysrq = 0

#####
# Networking performance
#####
net.core.default_qdisc = fq
net.ipv4.tcp_congestion_control = bbr
net.core.somaxconn = 65536
net.core.netdev_max_backlog = 16384
net.ipv4.ip_local_port_range = 10240 65535
```

```
net.ipv4.tcp_fastopen = 3

#####
# Memory / VM tuning
#####
vm.swappiness = 5
vm.dirty_ratio = 20
vm.dirty_background_ratio = 5
vm.dirty_expire_centisecs = 2000
vm.dirty_writeback_centisecs = 200
vm.vfs_cache_pressure = 20
vm.overcommit_memory = 0
fs.file-max = 4194304

# Huge number of inotify watchers
fs.inotify.max_user_watches = 2097152
fs.inotify.max_user_instances = 16384

#####
# Optional: DoS / connection safety
#####
net.ipv4.icmp_echo_ignore_all = 0
```

[linux](#), [sysctl](#), [network](#), [network tuning](#)

From:

<https://wiki.rtzra.ru/> - RTzRa's hive

Permanent link:

<https://wiki.rtzra.ru/devops/net/sysctl-net-performance>

Last update: **2026/04/18 09:54**

