
Subject: Re: too many of orphaned sockets
Posted by [Vasily Tarasov](#) on Wed, 30 Aug 2006 05:50:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

So you're using 2.6.16 series...

Look at the code:

```
static inline int ub_too_many_orphans(struct sock *sk, int count)
{
#define CONFIG_USER_RESOURCE
    if (__ub_too_many_orphans(sk, count))                      # MAY BE WE HAVE 1 HERE?
        return 1;
#endif
    return (ub_get_orphan_count(sk) > sysctl_tcp_max_orphans ||
            (sk->sk_wmem_queued > SOCK_MIN_SNDBUF &&
             atomic_read(&tcp_memory_allocated) > sysctl_tcp_mem[2]));
}
```

So, what we have in __ub_too_many_orphans(sk, count):

```
int __ub_too_many_orphans(struct sock *sk, int count)
{
    struct user_beancounter *ub;

    if (sock_has_ubb(sk)) {
        for (ub = sock_bc(sk)->ub; ub->parent != NULL; ub = ub->parent);
        if (count >= ub->ubb_parms[UB_NUMTCPSOCK].barrier >> 2)          # IT HOLDS
TRUE
        return 1;
    }
    return 0;
}
```

So the number of orphaned sockets (count) is greater, then (barrier of NUMTCPSOCK parameter)/4. Thus, if the reason is that, you can increase the barrier (not limit!) of numtcpsock parameter.

HTH.
