
Subject: Re: [PATCH v6 06/10] tcp buffer limitation: per-cgroup limit
Posted by KAMEZAWA Hiroyuki on Mon, 28 Nov 2011 03:24:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

some comments.

On Fri, 25 Nov 2011 15:38:12 -0200
Glauber Costa <glommer@parallels.com> wrote:

> This patch uses the "tcp.limit_in_bytes" field of the kmem_cgroup to
> effectively control the amount of kernel memory pinned by a cgroup.
>
> This value is ignored in the root cgroup, and in all others,
> caps the value specified by the admin in the net namespaces'
> view of tcp_sysctl_mem.
>
> If namespaces are being used, the admin is allowed to set a
> value bigger than cgroup's maximum, the same way it is allowed
> to set pretty much unlimited values in a real box.
>
> Signed-off-by: Glauber Costa <glommer@parallels.com>
> CC: David S. Miller <davem@davemloft.net>
> CC: Hiroyuki Kamezawa <kamezawa.hiroyu@jp.fujitsu.com>
> CC: Eric W. Biederman <ebiederm@xmission.com>

<snip>

```
> EXPORT_SYMBOL(tcp_destroy_cgroup);
> +
> +int tcp_update_limit(struct mem_cgroup *memcg, u64 val)
> +{
> + struct net *net = current->nsproxy->net_ns;
> + struct tcp_memcontrol *tcp;
> + struct cg_proto *cg_proto;
> + int i;
> + int ret;
> +
> + cg_proto = tcp_prot.proto_cgroup(memcg);
> + if (!cg_proto)
> + return -EINVAL;
> +
> + tcp = tcp_from_cgproto(cg_proto);
> +
> + ret = res_counter_set_limit(&tcp->tcp_memory_allocated, val);
```

Here, you changed the limit.

> + if (ret)

```
> + return ret;  
> +  
> + val >>= PAGE_SHIFT;
```

Here, you modifies 'val'

```
> +  
> + for (i = 0; i < 3; i++)  
> +   tcp->tcp_prot_mem[i] = min_t(long, val,  
> +     net->ipv4.sysctl_tcp_mem[i]);  
> +  
> + if (val == RESOURCE_MAX)  
> +   jump_label_dec(&memcg_socket_limit_enabled);
```

the 'val' never be RESOUCE_MAX.

```
> + else {  
> +   u64 old_lim;  
> +   old_lim = res_counter_read_u64(&tcp->tcp_memory_allocated,  
> +     RES_LIMIT);
```

old_lim is not already overwritten ?

```
> + if (old_lim == RESOURCE_MAX)  
> +   jump_label_inc(&memcg_socket_limit_enabled);  
> + }  
> + return 0;  
> +}  
> +
```

Thanks,
-Kame
