

Re: page allocation failure

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-01/5007.html>

From: Andrew Morton (akpm_at_osdl.org)

Date: 01/20/04

Date: Mon, 19 Jan 2004 19:38:37 -0800

To: Oliver Kiddle <okiddle@yahoo.co.uk>

Oliver Kiddle <okiddle@yahoo.co.uk> wrote:

>
> *There seems to be a problem with 2.6.1 on my machine. It will be fine
> for a matter of a few days and then this error will appear on the
> console. The message then appears repeatedly and continuously. The
> first I know is that my remote login shell ceases to respond. About the
> only thing I can do is switch between virtual consoles (until I hit the
> reset button).*
>
> */var/log/messages shows:*
> *kernel: cat: page allocation failure. order:0, mode:0x20*
>
> *Then the same for lots of other processes (pdflush, syslogd, klogd,
> kswapd0, nfsd to name a few). I expect that after a point it is unable
> to even log stuff so syslog is quiet after a while.*
>
> *It has happened three times now and on all occasions, I was untarring a
> huge file on an XFS partition. I assume the problem is something to do
> with VM. The machine has 1GB of RAM which should be plenty. For the
> most part it is just serving NFS and NIS (to no more than about 10
> clients).*

Does the machine actually recover, or does it grind to a halt and need resetting?

Is there much network receive happening at the time?

Are you using gig-E with large MTU's?

> *If anyone can suggest any /proc variables I might change to reduce the
> risk of it doing this again, I would appreciate it. I tried increasing
> /proc/sys/vm/min_free_kbytes after the first time this happened. Not
> that I understand what that does: I searched the archives and it was
> mentioned in a vaguely relevant looking post.*

Yup, min_free_kbytes is the right thing to increase. Try it again, perhaps increasing it by more – to 10000 or something like that.

Linux-Kernel: Re: page allocation failure

min_free_kbytes will increase the amount of memory which the VM keeps in reserve to satisfy interrupt-time memory allocation attempts – most especially network receive.

You probably should apply this patch to tell us where the allocation failures are coming from. Make sure that CONFIG_KALLSYMS is enabled in kernel config.

```
diff -puN mm/page_alloc.c~a mm/page_alloc.c
--- 25/mm/page_alloc.c~a Mon Jan 19 19:34:09 2004
+++ 25-akpm/mm/page_alloc.c Mon Jan 19 19:34:21 2004
@@ -674,6 +674,7 @@ nopage:
     printk("%s: page allocation failure."
           " order:%d, mode:0x%x\n",
           p->comm, order, gfp_mask);
+ dump_stack();
     }
     return NULL;
got_pg:
```

–

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>