Linux Memory Management or 'Why is there no free RAM?'

View unanswered posts
View posts from last 24 hours
Goto page 1, 2, 3, 4 Next

postreply Gentoo Forums Forum Index → Documentation, Tips & Tricks

View previous topic :: View next topic

Author | Message

sapphirecat
Guru

Joined: 14 Jan 2003
Posts: 376
Location: Jamestown NY (amazingly enough)

Posted: Wed May 19, 2004 8:16 pm Post subject: Linux Memory Management or 'Why is there no free RAM?' quote

Linux Memory Management or 'Why is there no free RAM?'

Revision 2.3

Sections

1. Overview of memory management

Traditional Unix tools like 'top' often report a surprisingly small amount of free memory after a system has been running for a while. For instance, after about 3 hours of uptime, the machine I'm writing this on reports under 60 MB of free memory, even though I have 512 MB of RAM on the system. Where does it all go?

The biggest place it's being used is in the disk cache, which is currently over 290 MB. This is reported by top as "cached". Cached memory is essentially free, in that it can be replaced quickly if a running (or newly starting) program needs the memory.

The reason Linux uses so much memory for disk cache is because the RAM is

wasted if it isn't used. Keeping the cache means that if something needs the same data again, there's a good chance it will still be in the cache in memory. Fetching the information from there is around 1,000 times quicker than getting it from the hard disk. If it's not found in the cache, the hard disk needs to be read anyway, but in that case nothing has been lost in time.

To see a better estimation of how much memory is really free for applications to use, run the command:

Code:

```
free -m
```

The -m option stands for megabytes, and the output will look something like this:

Code:

```
             total       used       free     shared    buffers     cached
Mem:           503        451         52          0         14        293
-/+ buffers/cache:        143        360
Swap:         1027          0       1027
```

The -/+ buffers/cache line shows how much memory is used and free from the perspective of the applications. Generally speaking, if little swap is being used, memory usage isn't impacting performance at all.

Notice that I have 512 MB of memory in my machine, but only 503 is listed as available by free. This is mainly because the kernel can't be swapped out, so the memory it occupies could never be freed. There may also be regions of memory reserved for/by the hardware for other purposes as well, depending on the system architecture.

2. The mysterious 880 MB limit on x86

By default, the Linux kernel runs in and manages only low memory. This makes managing the page tables slightly easier, which in turn makes memory accesses slightly faster. The downside is that it can't use all of the memory once the amount of total RAM reaches the neighborhood of 880 MB. This has historically not been a problem, especially for desktop machines.

To be able to use all the RAM on a 1GB machine or better, the kernel needs recompiled. Go into 'make menuconfig' (or whichever config is preferred) and set the following option:

Code:

```
Processor Type and Features ---->
High Memory Support ---->
(X) 4GB
```

This applies both to 2.4 and 2.6 kernels. Turning on high memory support theoretically slows down accesses slightly, but according to Joseph_sys and log, there is no practical difference.

3. The difference among VIRT, RES, and SHR in top output

VIRT stands for the virtual size of a process, which is the sum of memory it is actually using, memory it has mapped into itself (for instance the video card's RAM for the X server), files on disk that have been mapped into it (most notably shared libraries), and memory shared with other processes. VIRT represents how much memory the program is able to access at the present moment.

RES stands for the resident size, which is an accurate representation of how much actual physical memory a process is consuming. (This also corresponds directly to the %MEM column.) This will virtually always be less than the VIRT size, since most programs depend on the C library.

SHR indicates how much of the VIRT size is actually sharable (memory or libraries). In the case of libraries, it does not necessarily mean that the entire library is resident. For example, if a program only uses a few functions in a library, the whole library is mapped and will be counted in VIRT and SHR, but only the parts of the library file containing the functions being used will actually be loaded in and be counted under RES.

4. The difference between buffers and cache

Buffers are associated with a specific block device, and cover caching of filesystem metadata as well as tracking in-flight pages. The cache only contains parked file data. That is, the buffers remember what's in directories, what file permissions are, and keep track of what memory is being written from or read to for a particular block device. The cache only contains the contents of the files themselves.

Corrections and additions to this section welcome; I've done a bit of guesswork based on tracing how /proc/meminfo is produced to arrive at these conclusions.

5. Swappiness (2.6 kernels)

Since 2.6, there has been a way to tune how much Linux favors swapping out to disk compared to shrinking the caches when memory gets full.

ghoti adds:

When an application needs memory and all the RAM is fully occupied, the kernel has two ways to free some memory at its disposal: it can either reduce the disk cache in the RAM by eliminating the oldest data or it may swap some less used portions (pages) of programs out to the swap partition on disk.
It is not easy to predict which method would be more efficient.
The kernel makes a choice by roughly guessing the effectiveness of the two

methods at a given instant, based on the recent history of activity.

Before the 2.6 kernels, the user had no possible means to influence the calculations and there could happen situations where the kernel often made the wrong choice, leading to thrashing and slow performance. The addition of swappiness in 2.6 changes this.
Thanks, ghoti!

Swappiness takes a value between 0 and 100 to change the balance between swapping applications and freeing cache. At 100, the kernel will always prefer to find inactive pages and swap them out; in other cases, whether a swapout occurs depends on how much application memory is in use and how poorly the cache is doing at finding and releasing inactive items.

The default swappiness is 60. A value of 0 gives something close to the old behavior where applications that wanted memory could shrink the cache to a tiny fraction of RAM. For laptops which would prefer to let their disk spin down, a value of 20 or less is recommended.

As a sysctl, the swappiness can be set at runtime with either of the following commands:

Code:

```
# sysctl -w vm.swappiness=30
# echo 30 >/proc/sys/vm/swappiness
```

The default when Gentoo boots can also be set in /etc/sysctl.conf:

Code:

```
# Control how much the kernel should favor swapping out applications (0-100)
vm.swappiness = 30
```

Some patchsets allow the kernel to auto-tune the swappiness level as it sees fit; they may not keep a user-set value.

.....

I promised to write about this as soon as I figured out the last thing I was curious about myself. If there are other topics I should cover, please send me a PM.

_____

Whatever background or technological wizardry we possess, if we do not have kindness and empathy, we are lost. (seen on e2)

Last edited by sapphirecat on Mon Jun 28, 2004 12:06 pm; edited 4 times in total

Back to top profile pm

andrewy
Guru

Posted: Wed May 19, 2004 8:22 pm Post subject: quote

Thanks, I'm sure this will help many people.
Why Linux seemed to use so much memory is one thing that I never understood when I first started using Linux, it took a long time for me to learn why, so I'm

Joined: 06 Apr 2004
Posts: 595

glad someone has taken the trouble to explain it for the people that might be newer to Linux.

Back to top

profile pm

beastmaster
Apprentice

Joined: 23 May 2004
Posts: 230

Posted: Sun May 30, 2004 3:25 pm Post subject:

quote

Does that mean that I should get a 1 GB+ DRAM just to take advantage of linux 2.6's memory management system?
I'm currently having 512 , and by looking at the "free", there are only 77 free mem left

Code:

```
             total       used       free     shared    buffers     cached
Mem:           502        425         77          0         10        210
-/+ buffers/cache:        204        297
Swap:         1153          0       1153
```

So, by getting a 1 GB+ I can enjoy even some better locality performance?

Back to top

profile pm

stahlsau
Guru

Joined: 09 Jan 2004
Posts: 513
Location: WildWestwoods

Posted: Sun May 30, 2004 3:36 pm Post subject:

quote

mmh..as i understood, you have too much memory ´cause 70mb are free

_____
Join the adopt an unanswered post initiative today

Back to top

profile pm

beastmaster
Apprentice

Joined: 23 May 2004
Posts: 230

Posted: Sun May 30, 2004 3:45 pm Post subject:

quote

Quote:

> mmh..as i understood, you have too much memory ´cause 70mb are free

, I think 512MB RAM is not linux thirsty enough...

Back to top

profile pm

grzewho
l33t

Joined: 31 Dec 2002
Posts: 614
Location: /home/g

Back to top

Posted: Sun May 30, 2004 3:59 pm Post subject: quote

640kb ought to be enough for everyone

_____

Code:

```
USE="freedom -software_patents" emerge --deep --update world
```

profile pm

sapphirecat
Guru

Joined: 14 Jan 2003
Posts: 376
Location: Jamestown NY (amazingly enough)

Posted: Mon May 31, 2004 8:04 pm Post subject: quote

beastmaster wrote:

> **Does that mean that I should get a 1 GB+ DRAM just to take advantage of linux 2.6's memory management system?**
> **I'm currently having 512, and by looking at the "free", there are only 77 free mem left**
>
> **Code:**
>
> ```
>              total       used       free     shared    buffers     cached
> Mem:           502        425         77          0         10        210
> -/+ buffers/cache:        204        297
> Swap:         1153          0       1153
> ```
>
> **So, by getting a 1 GB+ I can enjoy even some better locality performance?**

Probably not; the cache has a quickly diminishing performance boost. The only reason Linux lets it get so big is because the memory is available and a marginal gain is still a gain.

Besides, if you look at the buffers/cache line, over half your memory is still truly available.

_____

Whatever background or technological wizardry we possess, if we do not have kindness and empathy, we are lost. (seen on e2)

Back to top

profile pm

beastmaster
Apprentice

Joined: 23 May 2004
Posts: 230

Posted: Tue Jun 01, 2004 9:31 pm Post subject: quote

Back to top

profile pm

dalek
l33t

Posted: Tue Jun 01, 2004 11:20 pm Post subject: quote

Joined: 19 Sep 2003
Posts: 978
Location: Mississippi USA

I'm bookmarking this thread. I am in a lot of forums and people freak out because Linux really USES the memory to make the system faster.

Finally a guru who explained it. Would have myself but nobody would listen anyway.

Later

_____
My rig: ABIT NF7 mobo with sound
AMD 2500+ || Volcano 12+ Cooler
1Gb Kingston Ram || fx 5200 NVIDIA 128Mb video card
Maxtor 80 GB & Western Digital 80GB HD.

BAD TYPER AND SPELLER. Don't bug me about it.

Back to top | profile | pm

neenee
Veteran
Veteran

Joined: 20 Jul 2003
Posts: 1783

Posted: Wed Jun 02, 2004 1:56 am Post subject: quote

hm.. perhaps this can be made a sticky or something - since many new users ask the question about why gentoo/linux/etc 'uses up' most of their ram.

_____
proud to be a scout and a chronic penguin hugger
Legion of Lore - site

Back to top | profile | pm

Souperman
Guru

Joined: 14 Jul 2003
Posts: 449
Location: Cape Town, South Africa

Posted: Fri Jun 04, 2004 5:55 am Post subject: quote

Quite informative, thanks. I figured out how to read meminfo a while ago, but something that I didn't know (not that I really researched it, to be honest) is what exactly buffers and cache are.

_____
moo

Back to top | profile | pm

gentoofailure
n00b

Joined: 18 Apr 2004
Posts: 30

Posted: Fri Jun 04, 2004 12:42 pm Post subject: quote

Great post this is what I was just looking for and has put my mind at ease

Back to top | profile | pm

DooMi
n00b

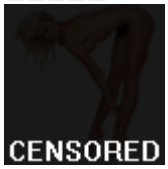

Joined: 03 May 2004
Posts: 74
Location: /dev/null

Back to top

Posted: Fri Jun 04, 2004 12:58 pm Post subject: quote

very nice explenation.
thanks

profile pm

dalek
l33t

Joined: 19 Sep 2003
Posts: 978
Location: Mississippi USA

Posted: Fri Jun 04, 2004 3:31 pm Post subject: quote

I need sound. How do you do a whistle on this thing? Love that avatar, DooMi.

Later

edit: I need to learn to type better.

_____
My rig: ABIT NF7 mobo with sound
AMD 2500+ || Volcano 12+ Cooler
1Gb Kingston Ram || fx 5200 NVIDIA 128Mb video card
Maxtor 80 GB & Western Digital 80GB HD.

BAD TYPER AND SPELLER. Don't bug me about it.

Back to top

profile pm

ghoti
Veteran
Veteran

Joined: 30 Dec 2002
Posts: 2074
Location: Belgium

Posted: Fri Jun 04, 2004 6:19 pm Post subject: Re: Linux Memory Management or 'Why is there no free RAM?' quote

sapphirecat wrote:

> **Linux Memory Management or 'Why is there no free RAM?'**
> *Revision 2*

Excellent !
It is indeed a FAQ, even in french, so I decided to translate your text and post it on the french forum .
I only hope not to have betrayed your mind !

Back to top

profile pm

sapphirecat
Guru

Posted: Sat Jun 05, 2004 11:53 am Post subject: Re: Linux Memory Management or 'Why is there no free RAM?' quote

Joined: 14 Jan 2003
Posts: 376
Location: Jamestown NY (amazingly enough)

ghoti wrote:

> Excellent !
> It is indeed a FAQ, even in french, so I decided to translate your text and post it on the french forum .
> I only hope not to have betrayed your mind !

Great!

I can't read French, but that didn't stop me from trying. AFAICT, Leander256 expanded on the swappiness section? I would appreciate having that in English to merge in here. Feel free to PM me about that.

_____
Whatever background or technological wizardry we possess, if we do not have kindness and empathy, we are lost. (seen on e2)

Back to top

profile pm

ghoti
Veteran
Veteran

Joined: 30 Dec 2002
Posts: 2074
Location: Belgium

Posted: Sat Jun 05, 2004 7:52 pm Post subject: Re: Linux Memory Management or 'Why is there no free RAM?' quote

sapphirecat wrote:

> AFAICT, Leander256 expanded on the swappiness section? I would appreciate having that in English to merge in here. Feel free to PM me about that.

In fact, Leander256 was afraid a french speaking n00b could not understand the english word "swappiness", nor my poor translation of the word. So I've added a paragraph to explain the concepts even more.

I've tried to translate it in english but please be lenient and feel free to adapt this to a less "frenchie english"

Quote:

> When an application needs memory and all the RAM is fully occupied, the kernel has two "relief tanks" at disposal :
> it can either reduce the disk cache in the RAM by eliminating the oldest data or it may swap some less used portions (pages) of programs out to the swap partition on disk.
> It is not easy to predict which method would be more efficient.
> Almost because of the Murphy's Law (AKA "law of the buttered slice of bread") which says that the last portion you've just eliminated will at once be required the next second after !
> The kernel makes a choice by roughly guessing the effectiveness of the two methods at a given instant.
> For instance, it can see that, in the last minutes, all the data in the cache where manipulated by a program but that another program was sleeping in the meantime.
> One may rightfully suppose that this situation will not change in the next seconds. So, it would be better to swap out the "sleeping" program and to keep the "alive" datas in the faster RAM.
> On the contrary, it can happen that all the programs are very active but that they manipulate almost no data. It would then be more efficient to keep the programs in memory and eliminate those old datas that seem not used anymore.
>
> Before the 2.6 kernels, the user had no possible mean to influence the

> calculations and there could happen situations where the kernel often made the bad choice. But the consequence of a bad choice is a great activity from the disk which results in a considerable reduction of performances.
>
> Since the 2.6, the user can assign a variable to force on every instant the degree of preference the kernel must use to make his choices. This is the "swappiness" variable (awkwardly translated in french as "permutabilité") that refers to the degree with which the kernel would tend to swap the pages to the disk rather than to reduce the disk cache in RAM.

Back to top profile pm

sapphirecat
Guru

Joined: 14 Jan 2003
Posts: 376
Location: Jamestown NY (amazingly enough)

Posted: Sun Jun 06, 2004 11:31 am Post subject: Re: Linux Memory Management or 'Why is there no free RAM?' quote

> ghoti wrote:
>
> I've tried to translate it in english but please be lenient and feel free to adapt this to a less "frenchie english"

Alright, revision 2.1 done. I hope you don't mind me taking the liberty of condensing it quite a bit.

_____
Whatever background or technological wizardry we possess, if we do not have kindness and empathy, we are lost. (seen on e2)

Back to top profile pm

revertex
l33t

Joined: 23 Apr 2003
Posts: 788

Posted: Mon Jun 07, 2004 1:40 am Post subject: quote

Hi, i have 512Mb ram, and i'm trying to use every mega as possible, but changing the swappiness priority seems does't work well in gentoo.
I'm trying change

> Code:
>
> /proc/sys/vm/swappiness

to zero to prevent things to be swapped to disk but ever time i reboot it comes to 83.
I had set swap priority in fstab to zero and edited
/etc/sysctl.conf. if it matters.

http://kerneltrap.org/node/view/3000

Back to top profile pm

beastmaster
Apprentice

Joined: 23 May 2004
Posts: 230

Posted: Mon Jun 07, 2004 3:42 am Post subject: quote

hi revertex,

After a semester of "machine arch" course,
I remember my prof said about it's good to have big size L1, L2, and L3 caches, just so it can greatly redude cache misses (capacity misses in this case)...
The idea is that not needing to touch the swap space as much as possible (since it's on-disk, so the move to transfering data into ram is like... I think it's said

100x ~1000x slower), since kernel 2.6 will make fully use of the RAMs.

so in your case (maybe a lot of mem intensive programs hogging), if you are running out of physical mem, just buy another 512 MB RAM, and once again every bite is worth the money

edit: I may be wrong, so correct me if there is

Back to top profile pm

sapphirecat
Guru

Joined: 14 Jan 2003
Posts: 376
Location: Jamestown NY (amazingly enough)

Posted: Mon Jun 07, 2004 9:36 am Post subject: quote

> **revertex wrote:**
>
> I'm trying change
>
> Code:
>
> ```
> /proc/sys/vm/swappiness
> ```
>
> to zero to prevent things to be swapped to disk but ever time i reboot it comes to 83.

Is it one of the kernels which auto-tunes swappiness? Does the value change after you run it a while (say, while emerging something)?

_____
Whatever background or technological wizardry we possess, if we do not have kindness and empathy, we are lost. (seen on e2)

Back to top profile pm

revertex
l33t

R T F M

Joined: 23 Apr 2003
Posts: 788

Posted: Tue Jun 08, 2004 11:33 pm Post subject: quote

sory for delay, now my swap file is untouched!
my box stay long time whitout reboot, maybe i forgot it
thanks all for reply, i just noted that my emerge sync now take ages, dunno why.
my kernel is gentoo-dev-sources.

Last edited by revertex on Mon Jun 21, 2004 6:23 pm; edited 1 time in total

Back to top profile pm

dalek
l33t

Joined: 19 Sep 2003
Posts: 978
Location: Mississippi USA

Posted: Wed Jun 09, 2004 12:40 am Post subject: quote

I went to sleep waiting on mine to sync up yesterday. I did notice that portage had a update. That may have had something to do with it.

I though it would never finish.

Later

_____
My rig: ABIT NF7 mobo with sound

AMD 2500+ || Volcano 12+ Cooler
1Gb Kingston Ram || fx 5200 NVIDIA 128Mb video card
Maxtor 80 GB & Western Digital 80GB HD.

BAD TYPER AND SPELLER. Don't bug me about it.

Back to top profile pm

revertex
l33t
R T F M
Joined: 23 Apr 2003
Posts: 788

Posted: Wed Jun 09, 2004 6:04 am Post subject: quote

It's happen again!
I'm working with some svg files, not so big, then my machine start to use swap, after this swappiness goes to 94!

Code:
```
cat /proc/sys/vm/swappiness 94
```

i will take a closer look in my kernel config later, thanks about the auto-tunes swappiness info.
beastmaster, i'm running fluxbox with some "diet" apps, 512 isn't enough?
therefore another 512 is a matter of time (read money)

Back to top profile pm

Evangelion
Veteran
Veteran

Joined: 31 May 2002
Posts: 1087
Location: Helsinki, Finland

Posted: Wed Jun 09, 2004 6:48 am Post subject: Re: Linux Memory Management or 'Why is there no free RAM?' quote

sapphirecat wrote:
> 2. The mysterious 800 MB limit on x86

related to this: do you need to enable high memory-support on AMD64-machines to take advantage of lots of RAM?

_____
My tech-blog | My other blog

Back to top profile pm

Display posts from previous: All Posts Oldest First Go

postreply Gentoo Forums Forum Index → Documentation, Tips & Tricks

All times are GMT - 5 Hours
Goto page 1, 2, 3, 4 Next

Page 1 of 4

Jump to: Documentation, Tips & Tricks Go

You cannot post new topics in this forum
You cannot reply to topics in this forum
You cannot edit your posts in this forum
You cannot delete your posts in this forum
You cannot vote in polls in this forum

Links: forums.gentoo.org | www.gentoo.org | bugs.gentoo.org | forum-mods@gentoo.org