

Re: How to set up a Linux machine that occupies the minimum memory footprint ?

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.misc/2004-11/2681.html>

From: Penang (penang_at_myrealbox.com)

Date: 11/30/04

Date: 29 Nov 2004 22:49:03 -0800

First of all, I need to thank you for your very helpful reply. To all, sorry for the long, long quotes.

Abdullah Ramazanoglu <abdullah@ramazanoglu.tr> wrote in message news:<310h9mF3287peU1@uni-berlin.de>...

> >>>> *Anyway, is there not some limiting device in Linux to put a cap on the
> >>>> mount of memory each user is allowed to use?*

> >>> *Sure. But how do you know how much he is using? Plenty of it will be
> >>> shared – is he using it or only 1/n'th of it? And then what about
> >>> memory he has been promised but will never ask for? Do you count memory
> >>> that is copy-on-write but has not been written to?*

> >> *Agreed in general. However there must be some limitation of a 'user'
> >> program that obtains then uses all available memory until the swap
> >> partition is full. sort of imagine a program that calls gobs of
> >> Mallocs, writes them with zeroes and keeps going until the kernel is
> >> virtually paralysed.*

> >> *If the user in question runs his rendering software in 'user' mode
> >> with say a limit of 3.8G out of 4G RAM, then it can optimise its
> >> functions around this rather than fondly imagining it has (say) 8G
> >> memory and hence causes the kernel to 'thrash' thus slowing down the
> >> application.*

> > *I read the above with great interest, because it addresses the
> > potential traps about the very task that I am about to commit.
> > However, I must apologize for my own stupidity, because I do not fully
> > comprehend what is stated above.*

> *Parent poster didn't say a system with 8GB of RAM would run slower than
> one with 4GB. What he said was that, with 8GB of *virtual* memory (i.e.
> 4GB physical RAM + 4GB swap area) the application will think that it can
> use all of it freely, and thus it will use up all the RAM plus all the
> swap area, and thus it will cause heavy swapping, and thus the performance*

comp.os.linux.misc: Re: How to set up a Linux machine that occupies the minimum memory footprint ?

- > *will suffer. Of course, if you use 8GB of RAM, or more precisely if you*
- > *use a swap area that is no more than 1/10 of the physical RAM, then you*
- > *can't be swapping heavily. You will get out of virtual memory and perhaps*
- > *your app will crash, but it will not be able to get a bigger footprint*
- > *than 110% of physical RAM. So, the more RAM is the better. Period.*

As I stated above, I am a dumbass. Please accept my apology for my following dumb question:

Are you saying that if I have a 4GB RAM memory, and I set up a 400MB (10%) of swap space, then Linux will somehow "know" that it doesn't have too much to play on, and it will "follow the line", so to speak, and not thrashing the disk for no apparent reason ?

If so, what if my application suddenly needs 4.8 GB of memory, and since Linux and all the other overheads already occupied 100–200MB of memory, the actual Virtual Memory Space has only 4.2GB left. And if my application needs 4.8 GB of memory, my application will crash.

Am I right ?

Hmmm.... in that case, I think I better go the 8GB route, for I don't want a crash, any crash at all, in the middle of an intense rendering job.

- > > *The software in question is a beast – originally run under Solaris –*
- > > *and although we've managed to ported it to run under Linux, we haven't*
- > > *have the chance to throw everything on it, yet.*

- > *Let us get it straight. You've said in another message that your app is*
- > *not 64-bit aware. But it works under 64-bit Sparc/Solaris with more than*
- > *4GB of RAM, right? If so it will also run under Opteron/Linux with the*
- > *same amount of RAM.*

Originally it knew 64Bit, but when we ported back to Linux – that time AMD's 64bit thing didn't look very promising – we aimed for the IA32 architecture – so now this beast only "see" 32Bit space.

I dunno how tough it'd be to "change" the specs again for the beast to work in 64Bit environment, we'd try, but I dunno if that's going to have any negative effect or not. Very tight timeframe we're talking about.

- > > *True, we have successfully get it to render snippets of*
- > > *not-so-very-detailed test runs, and because we haven't really tapped*
- > > *out it's full potential, I can't really tell you how hungry the beast*
- > > *can be, if we are to render something like what we are about to do.*

Re: How to set up a Linux machine that occupies the minimum memory footprint ?

comp.os.linux.misc: Re: How to set up a Linux machine that occupies the minimum memory footprint ?

- > > *The 4GB figure is just a rough calculation – it could be more ! – and*
- > > *that's why I am worried about whether the entire thing could be*
- > > *achieved via the PC platform.*

- > *But you don't need to calculate! You have your previous experiences with*
- > *Solaris, shedding light on how much of RAM you would need under Linux.*

- > *I have a vague idea that you have not applied the same stressed conditions*
- > *to Solaris as yet. May be you have been working on Solaris with relatively*
- > *smaller projects, and now you've got a much bigger one and you're trying*
- > *to figure out whether a Sparc/Solaris upgrade or a new Linux setup would*
- > *be better? There are certain ends in your case that I can't meet, so I*
- > *have to speculate. Advice based on speculation would be dubious at best.*
- > *Could you clear up things a bit more please?*

You hit the bull's eye ! Essentially we are trying to get a huge thing to run through a little hole – and while moving from Solaris to Linux may seem to be a "downgrade" but in reality it's an upgrade.

Because of the task at hand, we are going to run this beast to its limits – and we are talking about rendering a feature movie, which no one has ever used this proprietary beast to do.

The Boss's original idea was to follow George Lucas' idea of using Maya / 3DMax on Windoze, we somehow managed to convince them to go the Unices way, and when we evaluate the (old) Solaris setup, we knew that it just didn't have the oomph for completing the job. That's why we "convince" the moneybag (our boss) to go the Linux way – and ported the thing into Linux.

Right now what I am facing is the next to the final stage – setting up a machine as a FULL SCALE TEST RUN – and if this is successful, we may see a farm setup.

If everything goes well, it would be good for all parties – Linux will have another credit to its fame, and we wouldn't be stucked with lots and lots of Windoze machines.

- > > *About optimization – we thought about it, but unfortunately we neither*
- > > *have the talent nor the time, budget, etc to achieve anything*
- > > *meaningful at that front, so right now what I am doing is to bite the*
- > > *bullet and set up a working system and start throw everything at it.*

- > > *After reading the above quoted message, my main worry is that if I set*
- > > *up a system with 8GB, would it hinder the rendering job MORE than a*
- > > *mere 4GB configuration ?*

- > *No.*

- > > *Or is there any special switch that I need to know about – either*
- > > *compilation, or configuration, or both – that can minimize the*

Re: How to set up a Linux machine that occupies the minimum memory footprint ?

comp.os.linux.misc: Re: How to set up a Linux machine that occupies the minimum memory footprint ?

> > *headache ?*

> *What I would suggest is more or less the compilation of what has already
> been advised by others:*

> *1. Go for 64 bit, preferably Athlon64 or, if you can't find an Athlon64
> motherboard which can handle more than 4GB, then go for a Uniprocessor
> Opteron. If you still can't find a >4GB mobo then go for dual Opteron.
> With dual Opteron I would be alert to get a shared memory system. I'm not
> very familiar with details of Numa architecture, but I gather that most of
> dual Opterons are Numa, and it may either be not possible or incur some
> overhead to use the whole memory from a single processor (thus from a
> single process). So a uniprocessor mobo with >4G capability is preferable
> for me (guaranteed to work) over a dual CPU mobo. I'm on thin ice here,
> though. So take it with a grain of salt. BTW, I would appreciate (and
> learn) if others point out possible errors I may have made here.*

> *But there's one thing I'm certain about: If your app can't be divided into
> separate processes (as I assume) then dual CPU will bring you only a
> marginal benefit. You will be able to run Linux and everything on one
> processor (mostly idling), which can help memory footprint a little in
> Numa, but your app will still have to run on a single CPU.*

> *2. For the short term, tweak your app so that, from the most urgent to the
> least:*

> *2/i) it will not break when it gets out of virtual memory, but it will
> gracefully adapt to the situation, perhaps at the cost of some performance
> degradation due to serialization of work,*

I'd find out if this is feasible.

> *2/ii) better, it will accept a parameter on how much memory to
> use,*

I'd be looking into it

> *2/iii) still better, it will be able to divide the work into smaller
> chunks.*

We are looking for ways to do this. Perhaps as little as a 2-second segment at a time. I know it's gonna to be very painful, but heck, if that's what it takes, that's what it takes.

> *The first item is crucial. If you can't tweak the program to do any one of
> these, then you have no other choice but use whatever RAM it gets to do
> the job.*

> *3. You can limit memory utilization to a certain extent with "ulimit" but*

Re: How to set up a Linux machine that occupies the minimum memory footprint ?

comp.os.linux.misc: Re: How to set up a Linux machine that occupies the minimum memory footprint ?

- > *please see the below thread for caveats. It's only better than nothing.*
- > <http://groups.google.com/groups?&threadm=8c526b62.0407291436.2955f4f2@posting.google.com>

- > *Also, to be able to limit your app's memory utilization externally, you*
- > *need to do the changes above in "2/i". As an alternative to ulimit, you*
- > *can simply use very small swap area (e.g. 1/16 to 1/64 of your RAM) and*
- > *unleash your app without ulimit. It will eventually hit the wall when it*
- > *depletes the whole virtual memory (RAM+swap). This is a surer method than*
- > *ulimit, but I'm not sure that it is better. When your vmem is maxed out*
- > *you won't be able to login as root to see what's going on! (open a root*
- > *session before hand). Please see the same thread above for this too.*

Many thanks for the insights !

- > *4. Returning back to the swap size, it depends on the locality of your*
- > *app. If your app, for instance, divides the work up into chunks of data to*
- > *be rendered separately (even if all within the virtual memory) and then*
- > *starts working on those chunks serially (i.e. one data chunk at a time =*
- > *you have high locality of reference), then having a huge swap area is*
- > *effectively equivalent to having huge data chunks waiting to be processed*
- > *on disk. Then you can use as much swap as needed without heavy swap*
- > *activity (only once in a while to swap out the finished chunk and swap in*
- > *the raw chunk). But if your app accesses all over the working data set at*
- > *the same time (i.e. one full-pass at a time = you have poor locality of*
- > *reference), then you have to keep all the data in RAM at all times, and*
- > *your app wouldn't tolerate even a low swap usage, so you should use a very*
- > *tiny swap space (but not zero, because there are still unused dead pages*
- > *here and there and you would want them get swapped out).*

- > *All in all, the most straightforward way seems to be finding a*
- > *uniprocessor Athlon64 or Opteron board with more than 4GB RAM capability.*

I will. And if anyone reads this message, and has info on which mobo is best suitable, I'd be more than grateful for any suggestion.

All in all, many, MANY THANKS for your kind input.