

Why do most people recommend to reduce swappiness to 10-20 ?

Source : <http://askubuntu.com/questions/184217/why-most-people-recommend-to-reduce-swappiness-to-10-20/184221>



I have seen in several site which recommend to reduce swappiness to 10-20 for better performance.

- <http://community.linuxmint.com/tutorial/view/998>
- <https://sites.google.com/site/easylinuxtipsproject/bugs>
- <http://www.howtogeek.com/115797/6-ways-to-speed-up-ubuntu/>
- <http://www.upubuntu.com/2012/06/11-tips-to-speed-up-computers-running.html>

Is it a myth or not? Is this a general rule? I have a laptop with 4GB Ram and 128GB SSD hard, what value do you recommend for my swappiness?

Thanks.



Because most believe that swapping = bad and that if you don't reduce swappiness, the system will swap when it really doesn't need to. Neither of those are really true. People associate swapping with times where their system is getting bogged down - however, it's mostly **swapping because the system is getting bogged down**, not the other way around. It's true that there are certain times when swapping can have a noticeable performance penalty, but the system will have already factored that in to its decision to swap, and decided that not doing so would have a greater overall penalty in system performance or stability which may later become noticeable. Overall the default settings should result in the best overall performance and reliability. I'd recommend leaving it at the default.

I hereby reproduce the answer I've given to [another question here on AskUbuntu](#):

How Linux uses RAM (very simplified) :

Each application can use some of your memory. Linux uses all otherwise unoccupied memory (except for the last few Mb) as "cache". This includes the page cache, inode caches, etc. **This is a good thing** - it helps speed things up heaps. Both writing to disk and reading from disk can be sped up immensely by cache.

Ideally, you have enough memory for all your applications, and you still have several hundred Mb left for cache. In this situation, as long as your applications don't increase their memory use and the system isn't putting too much pressure on the cache, there is no

need for any swap.

So, when someone refers to "free" RAM, this may or may not include cache, since cache will only occupy "free" RAM. The cache will shrink as RAM is occupied by more critical stuff, like applications.

How Linux uses swap (even more simplified) :

Once you have used up enough memory that there is not enough left for a smooth-running cache, Linux may re-allocate some unused application memory from RAM to swap.

It doesn't do this according to a definite cutoff though. **It's not like you reach a certain percentage of allocation then Linux starts swapping.** It has a rather "fuzzy" algorithm. It takes a lot of things into account, which can best be described by "how much pressure is there for memory allocation". If there is a lot of "pressure" to allocate new memory, then it will increase the chances some will be swapped to make more room. If there is less "pressure" then it will decrease these chances.

Your system has a "**swappiness**" setting which helps you tweak how this "pressure" is calculated. **It's normally not recommended to alter this at all, and I would certainly never recommend you alter it.** Swapping is overall a very good thing - any occasional performance penalties are intended to be offset by a gain in overall system responsiveness and stability for a wide range of tasks. If you reduce the swappiness, you let the amount of cache memory shrink a little bit more than it would otherwise, even when it may really be useful. **You therefore risk slowing down your computer in general, because there is less cache, while memory is being taken up by applications that aren't even using it.** Whether this is a good enough trade-off for whatever specific problem you're having with swapping is up to you. If you go further than this and actually disable swapping, you risk system instability in the event that the system runs out of memory for processes.

What is happening when the system is bogged down and swapping heavily?

A lot of the time people will look at their system that is thrashing the disk heavily and using a lot of swap space and blame swapping for it. That's the wrong approach to take. If swapping ever reaches this extreme, it means that your system is extremely low in memory and the swapping is the only thing keeping it from crashing or killing processes randomly. Without swapping, in this situation, processes will crash and die. **The swapping is a symptom of a deeper problem.** In a system with enough memory for all its tasks, swapping only ensures that memory is utilised in efficient ways, dealing out memory to the cache over dormant processes where it sees it will be worth it. In a system where swapping is relentlessly thrashing the disk, it's not swapping's fault.

When choosing what is to be swapped to disk, the system tries to pick memory that is not actually being used - read to or written from. It has a pretty simple algorithm for calculating this that chooses well most of the time.

In a desktop system that is used sporadically, there is a well-known situation in which suddenly trying to use process that has been sitting idle for quite some time results in a delay where the memory for the idle process has to be swapped back from disk to memory. This can result in a disk based bottleneck while it happens that slows down the system in general. It's typical when, for example, you've let your computer do tasks such as backups, torrents etc overnight, then suddenly in the morning you re-open your minimised browser window and try to use it again. This is a result of the system being unable to predict that a process that's been idle for a long period will suddenly be required once more. It's a side effect of the swapping system. Turn off swap and this initial sluggishness after the backup/virus scan may not happen, but the system may run a little bit slower all day long as there is more pressure on the cache size, and you lose the added protection against out-of-memory conditions which can cause processes to crash or be killed.

Note: none of this is a situation or problem that's limited to Linux. All modern operating system will operate this way.

The most effective way to solve sluggishness associated with memory swapping is to increase the physical memory size (or run less memory-hungry applications). Turning off swap or dialling it back is one of the least effective ways to solve it because swap is only a coping mechanism for the low memory situation.

Can/should swap be disabled on a system that has lots of RAM anyway?

If you have a system where you have a huge amount of RAM (at time of writing -July 2012-, 8GB is a huge amount for a typical Linux distro), then you will very rarely ever hit a situation where swap is needed at all. There is no need to explicitly disable swap at all or adjust its swappiness - your system will realise that there is plenty of RAM for both applications and cache and will never make the decision to swap.

You *could* disable swap or decrease swappiness, and it will very probably work just fine, but in all normal situations it will make absolutely zero difference because your system would not have decided to swap anyway due to having plenty of memory. The only situations in which it *would* make a difference would be in the unlikely situation the system finds itself running out of memory and consequently the cache system is getting hampered, and it's in this type of situation where you would *want* swap most. So you can safely leave swap on its normal settings for added peace of mind without it having a negative effect when you have plenty of memory.

But how can swap speed up my system? Doesn't swapping slow things down?

The act of transferring data from RAM to swap is a slow operation, but it's only taken when the kernel is pretty sure the overall benefit will outweigh this. For example, if your application memory has risen to the point that you have almost no cache left and your I/O is very inefficient because of this, you can actually get a lot more speed out of your system by freeing up some memory, even after the initial expense of swapping data in order to free it up.

It's also a last resort should your applications actually request *more* memory than you actually have. In this case, swapping is necessary to prevent an out-of-memory situation which will often result in an application crashing or having to be forcibly killed. There's no getting around the fact that your system will probably become dog-slow when this happens, but console yourself by realising that were it not for swap, your application would likely have crashed instead, potentially resulting in data loss.

Swapping is only *associated* with times where your system is performing poorly because it happens at times when you are running out of usable RAM, which would have slowed your system down (maybe even crashed applications) even if you didn't have swap. So to simplify things, **swapping happens because your system is becoming bogged down, rather than the other way around** - and in some cases it can save the day.

Once data is in swap, when does it come out again?

Transferring data out of swap is (for traditional hard disks, at least) just as time-consuming as putting it in there. So understandably, your kernel will be just as reluctant to remove data from swap, especially if it's not actually being used (ie read from or written to). If you have data in swap and it's not being used, then it's actually a good thing that it remains in swap, since it leaves more memory for other things that *are* being used, potentially speeding up your system.