



What does swappiness do and how does it affect swap_tendency?

🔒 SOLUTION VERIFIED - Updated June 10 2021 at 7:16 AM - English ▾

Environment

- Red Hat Enterprise Linux 4
- Red Hat Enterprise Linux 5

Issue

- What does swappiness do and how does it affect swap_tendency?

Resolution

Note : *The swapping algorithm is overhauled in Red Hat Enterprise Linux 6 and up to consider active and inactive memory for both file-backed and anonymous (non-file-backed) memory. While swappiness still exerts a considerable shift in behaviour on swap activity on a system, swappiness no longer uses the algorithm described below. As such, swap_tendency, distress, and mapped_ratio no longer exist in Red Hat Enterprise Linux 6 and above.*

- Swappiness is a property for the Linux kernel that changes the balance between swapping out runtime memory, as opposed to dropping pages from the system page cache. Swappiness can be set to values between 0 and 100, inclusive. A low value means the kernel will try to avoid swapping as much as possible where a higher value instead will make the kernel aggressively try to use swap space.

- As this parameter sets the kernel's balance between reclaiming pages from the page cache and reclaiming pages by swapping out process memory, a discussion of page reclaim overall is warranted. The reclaim code works (in a very simplified way) by calculating a few numbers:
- The `distress` value is a measure of how much trouble the kernel is having freeing memory. The first time the kernel decides it needs to start reclaiming pages, `distress` will be zero; if more attempts are required, that value goes up, approaching a max value of 100.
- The `mapped_ratio` value is an approximate percentage of how much of the system's total memory is mapped (i.e. is part of a process's address space) within a given memory zone.
- And `vm_swappiness` is the swappiness parameter, set to 60 by default.

With these numbers in hand, the kernel calculates its "swap tendency":

```
swap_tendency = mapped_ratio/2 + distress + vm_swappiness;
```

- If `swap_tendency` is below 100, the kernel will only reclaim page cache pages. Once it goes above that value, however, pages which are part of some process's address space will also be considered for reclaim. So, if the system's `distress` value is low and `swappiness` is set to 60, the system will not swap process memory until 80% of the total RAM in the system is allocated. Users who would like to never see application memory swapped out can set `swappiness` to a low value, say 5 or 10, causing the kernel to ignore process memory until the `distress` value gets quite high.
- Overall, increasing this value will make the system more inclined to utilize swap space, leaving more memory free for caches. Decreasing this value will make the system less inclined to swap, and may improve application responsiveness.
- Tuning `vm.swappiness` incorrectly may hurt performance or may have a different impact between light and heavy workloads. Changes to this parameter should be made in small increments and should be tested under the same conditions that the system normally operates.