

---

Posted by [kir](#) on Mon, 09 Jan 2006 08:27:54 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Most parameters provide both accounting of some system resource and allow controlling its consumption. The exceptions are `physpages` (accounting only) and `vmguarpages` (no accounting, control only).

`vmguarpages`: memory allocation guarantee.

This parameter controls how much memory is available to the Virtual Private Server (i.e. how much memory its applications can allocate by `malloc(3)` or other standard Linux memory allocation mechanisms). The more clients are served or the more "heavy" the application is, the more memory it needs.

The amount of memory that Virtual Private Server's applications are guaranteed to be able to allocate is specified as the barrier of `vmguarpages` parameter. The current amount of allocated memory space is accounted into `privvmpages` parameter, and `vmguarpages` parameter does not have its own accounting. The barrier and the limit of `privvmpages` parameter impose an upper limit on the memory allocations. The meaning of the limit for the `vmguarpages` parameter is unspecified in the current version and should be set to the maximal allowed value (2147483647).

If the current amount of allocated memory space does not exceed the guaranteed amount (the barrier of `vmguarpages`), memory allocations of Virtual Private Server's applications always succeed. If the current amount of allocated memory space exceeds the guarantee but below the barrier of `privvmpages`, allocations may or may not succeed, depending on the total amount of available memory in the system.

Starting from the barrier of `privvmpages`, normal priority allocations and, starting from the limit of `privvmpages`, all memory allocations made by the applications fail.

The memory allocation guarantee (`vmguarpages`) is a primary tool for controlling the memory available to Virtual Private Servers, because it allows administrators to provide Service Level Agreements -- agreements guaranteeing certain quality of service, certain amount of resources

and general availability of the service. The unit of measurement of `vmguarpages` values is memory pages (4KB) on 32-bit Intel-family processes).

The total memory allocation guarantees given to Virtual Private Servers are limited by the physical resources of the computer -- the size of RAM and the swap space.

`privvmpages`: memory allocation limit.

`Privvmpages` parameter allows controlling the amount of memory allocated by applications.

The barrier and the limit of `privvmpages` parameter control the upper boundary of the total size of allocated memory. Note that this upper boundary doesn't guarantee that the Virtual Private Server will be able to allocate that much memory, neither does it guarantee that other Virtual Private Servers will be able to allocate their fair share of memory. The primary mechanism to control memory allocation is the `vmguarpages` guarantee.

`Privvmpages` parameter accounts allocated (but, possibly, not used yet) memory. The accounted value is an estimation how much memory will be really consumed when the Virtual Private Server's applications start to use the allocated memory. Consumed memory is accounted into `oomguarpages` parameter.

Since the memory accounted into `privvmpages` may not be actually used, the sum of current `privvmpages` values for all Virtual Private Servers may exceed the RAM and swap size of the computer.

There should be a safety gap between the barrier and the limit for `privvmpages` parameter to reduce the number of memory allocation failures that the application is unable to handle. This gap will be used for "high-priority" memory allocations, such as process stack expansion. Normal priority allocations will fail when the barrier if `privvmpages` is reached.

Total `privvmpages` should match the physical resources of the computer. Also, it is important not to allow any Virtual Private Server to allocate a significant portion of all system RAM to avoid serious service level degradation for other VPSs.

---