Subject: Re: Memory & SWAP

Posted by xemul on Tue, 23 Jan 2007 09:17:52 GMT

View Forum Message <> Reply to Message

When you fork your new task inherits all memory from its parent. Thus if your task has allocated 128M of private anonymous mappings and forks new system can potentially eat 256M of memory - if this exceeds privvmpages we don't allow forking.

SWAP isn't limited now - it's just accounted into oomguarpages. When system becomes low of memory it starts to oom kill tasks taking oomguarpages into account.

When VE eats all provvmpages no swapout starts - only when the whole system runs out of mem.

Also note, that freeing memory consists of 3 stages:

- 1. shrink kernel caches
- 2. sysncing file caches to disk
- 3. swapping anonymous pages out

so if kernel caches were shrinked no user pages go away.