

DragonFlyBSD - Bug #2868

tmpfs problems: incorrect size, no auto-size

12/18/2015 01:16 PM - arcade@b1t.name

Status: Closed	Start date: 12/18/2015
Priority: Low	Due date:
Assignee:	% Done: 0%
Category: VFS subsystem	Estimated time: 0.00 hour
Target version:	
Description	
tmpfs version in DragonFly is limited when comparing to FreeBSD:	
1. When tmpfs is mounted automatically it's size by default (per man page) should be equal to all free mem + all free swap. Instead tmpfs filesystems are created with a fairly small size - like 200k on host with 1G mem.	
2. tmpfs size should be auto adjusted when other processes are eating mem or when swap is added. Instead as swap is initialized after mounting /tmp swap addition doesn't make tmpfs grow. On the other hand I never saw tmpfs shrinking when all available memory is used.	

History

#1 - 12/18/2015 01:23 PM - arcade@b1t.name

- Category changed from Kernel to VFS subsystem

#2 - 12/20/2015 12:40 AM - marino

"When tmpfs is mounted automatically it's size by default (per man page) should be equal to all free mem + all free swap"

I think you misinterpreted the man page

"Instead tmpfs filesystems are created with a fairly small size - like 200k on host with 1G mem. "

I get 30G on a system with 8G ram and a swappsize of 56G.

I get 108G on a system with 32G ram and a swappsize of 200G

The default tmpfs size is affected what the sytem has, but as not as you present. If you create a second tmpfs mount, it has the same capacity as the first one. (e.g. 108 + 108 > 208)

"tmpfs size should be auto adjusted when other processes are eating mem or when swap is added. "
Adjusted how?

"Instead as swap is initialized after mounting /tmp swap addition doesn't make tmpfs grow."
what do you mean, "grow"? Do you mean the potential capacity?

"On the other hand I never saw tmpfs shrinking when all available memory is used."
again, I assume you are talking about capacity.

I don't see a problem. If you don't want some "random" default tmpfs size, then specify the size with -s. I doubt anybody wants a tmpfs mount that fluxuates in capacity depending on what else the system is doing. I'm not sure how FreeBSD implementation is "superior" in any way.

#3 - 12/20/2015 12:47 AM - marino

208 should have been 232 but then it would be 108 + 108 + 108 > 232.

The point being dragonfly doesn't try to resize X tmpfs mounts based on available memory. If you have 3 mounts, do you shrink them proportionally? by the same amount? only 1 and leave the X-1 others alone?

Again, just defining the max size would solve all the uncertainty.

#4 - 12/20/2015 12:25 PM - arcade@b1t.name

Actually `mount_tmpfs` (8) says:

-s size Specifies the total file system size. If zero is given (the default), the available amount of memory (including main memory and swap space) will be used. Note that four megabytes are always reserved for the system and cannot be assigned to the file system.

There's nothing in here limiting tmpfs size down to 200k on a system with 1G mem or down to 30G on a system with 56G swap.

On FreeBSD... In FreeBSD each tmpfs size equals 'used space' + 'available_space' so that capacity of mounted tmpfs by default depletes when swap and memory are fully used up. Dunno whether this is hard to achieve but it makes single simple line in fstab useful in most conditions.

PS: And mentioning that FreeBSD version can be superior... For example in 2013 they replaced linked lists for directory structures with RB-Tree.

PPS: I'm not talking about blindly copying everything from FreeBSD. I'm just pointing out "weird" defaults not clarified by man pages. From my point of view tmpfs is a vital addition to hammer.

#5 - 12/22/2015 02:23 AM - marino

"the available amount of memory (including main memory and swap space)"

I don't read that to mean available-memory = main + swap.

It says the algorithm that determines how much memory is available to tmpfs will pull from ram and swap, but doesn't say what proportion.

"In FreeBSD each tmpfs size equals 'used space' + 'available_space' so that capacity of mounted tmpfs by default "

It's the same for DragonFly except the amount is what the algorithm determines.

#6 - 02/22/2018 01:42 AM - arcade@b1t.name

- *Status changed from New to Closed*