



To Swap or Not To Swap?

Should we enable swap for servers with Firebird (Windows and Linux), even if we have a lot of memory?

Alexey Kovyazin, President of Firebird Foundation

To Swap or Not To Swap?

- Linux
 - Why Linux need page file?
 - What RedHat says
 - What we see in the real world statistics?
- Windows
 - Why Windows need page file?
 - What Microsoft says?
 - What we can recommend?

Why Linux needs swap?

- Prevents Crashes
 - Protects from peaks of memory consumption and OOMKiller
- Enables Hibernation
- Improves Performance
 - By moving infrequently used data from RAM to swap, Linux can free up RAM for more active processes, potentially improving overall system performance
- General Performance Consideration:
 - While swap is slower than RAM (since it's on the disk), it's still a crucial component for system stability
 - However, relying too much on swap can lead to performance degradation, a situation known as "thrashing" where the system spends more time swapping than executing tasks.

What RedHat says?

- https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/storage_administration_guide/ch-swapspace

Table 15.1. Recommended System Swap Space

Amount of RAM in the system	Recommended swap space	Recommended swap space if allowing for hibernation
≤ 2 GB	2 times the amount of RAM	3 times the amount of RAM
> 2 GB – 8 GB	Equal to the amount of RAM	2 times the amount of RAM
> 8 GB – 64 GB	At least 4 GB	1.5 times the amount of RAM
> 64 GB	At least 4 GB	Hibernation not recommended



In Linux - command free -h/-m

```
[root@cbd2020 ~]# free -h
```

	total	used	free	shared	buff/cache	available
Mem:	125Gi	39Gi	4.1Gi	24Mi	82Gi	84Gi
Swap:	11Gi	5.9Gi	6.1Gi			

#	Mem Allocated	Swap allocated	Mem used	Swap Used	free mem	free swap	shar ed	buff/ cache	available
1	498688	4096	149504	551	1638	3584	33	347136	347136
2	257024	4096	141312	440	12288	3656	46	101376	109568
3	160768	4096	159744	0	1740	4096	1600	120832	122880
4	128000	37888	12288	124	21504	37764	7	94208	114688
5	128000	11264	39936	6041	4198	6246	24	83968	86016
6	128000	31744	40960	2867	3276	29696	5	82944	84992
7	63488	975	12288	975	29696	0	0	20480	50176
8	63488	975	15360	619	25600	356	0	21504	47104
9	57196	4095	43297	912	543	3183	46	13355	13497
10	46884	4095	34497	384	2350	3711	12	10036	12019

#	Mem Allocated	Swap allocated	Mem used	Swap Used	free mem	free swap	shared	buff/cache	available
11	32684	2048	22528	512	562	1536	29	8806	8908
12	32684	4095	27636	241	168	3854	42	4878	4775
13	32654	4095	22445	259	254	3836	11	9954	9916
14	32652	4095	25787	417	5851	3678	45	1012	6542
15	25402	4095	18241	495	177	3600	8	6983	6916
16	20144	4095	16078	253	2002	3842	44	2063	3818
17	19934	4095	14776	419	327	3676	45	4830	4911
18	19456	31744	5734	597	5734	31744	192	8192	12288
19	19456	8089	14336	1433	148	6656	3	4300	4096
20	17970	4095	14356	543	102	3552	45	3510	3418
21	16768	4095	12367	428	445	3667	11	3954	4207
22	16332	4095	12958	422	2506	3673	44	867	3147
23	11264	8089	9728	1945	112	6144	2	1843	1638
24	8010	8192	2560	252	569	8010	11	5120	5324
25	6963	3788	2867	829	3174	2969	3	943	3788



Analysis of 25 Linux production servers

- Number of servers with
 - Swap Used >0 — 24 of 25, 1 has 0 used
 - Swap used in full:
 - 64Gb Total, Swap allocated 975Mb, used 975Mb
 - Maximum Swap situation
 - 128Gb Total, Swap allocated 11Gb, used 6Gb

	Swap Allocated	Swap Used
Median	4095	495
Average	8085	864
Maximum	37888	6041
Minimum	975	0

- Allocated swap of 1Gb covers 21 of 25 servers
- Allocated swap of 4Gb covers 23 of 25 servers

Recommendation — use 4Gb of swap as initial value, increase it to 8Gb if all swap will be used

Windows Server

- Why Windows need page file?
- What Microsoft says?
- What we can recommend?

How Windows uses swap?

- Crash Dump Setting
 - Swap must be big enough for crash dump
- For peak memory usage
- Infrequently Accessed Pages

What Microsoft says?

- There is not so much information about Windows Server with databases
- In general, they recommend to let OS decide about swap size
- The links below
 - <https://learn.microsoft.com/en-us/answers/questions/458551/page-file-recommendations>
 - <https://learn.microsoft.com/en-us/troubleshoot/windows-client/performance/how-to-determine-the-appropriate-page-file-size-for-64-bit-versions-of-windows>

```
wmic os get SizeStoredInPagingFiles, TotalVisibleMemorySize,  
FreePhysicalMemory, FreeVirtualMemory -format:list
```


```
FreePhysicalMemory=493620
```

```
FreeVirtualMemory=9562220
```

```
SizeStoredInPagingFiles=30408704
```

```
TotalVisibleMemorySize=16497272
```

	FreePhysicalMemory	FreeVirtualMemory	SizeStoredInPagingFiles	TotalVisibleMemorySize
1	201802200	250745636	44040192	536388200
2	49650612	85927348	39845888	268302452
3	45528228	54264888	11534336	77823472
4	33189904	32903492	0	125828596
5	17214212	29695292	33794932	67108404
6	11370892	22974612	21912468	33518868
7	9124524	16313864	9437184	41941996
8	7836860	15774140	9437184	41941996
9	6226076	14616728	8388152	8388152
10	5426088	5730080	655360	10485300
11	4988616	3809752	655360	10485236
12	4703852	4519972	1310720	8387564
13	4593232	8150452	4194304	12582320
14	4264760	4079624	1048576	10485236
15	2690876	3944088	2064412	8387564
16	2562136	3655008	655360	10485236
17	2439504	3793676	5616716	12581868
18	1932308	15265552	13630968	13630968



- Only 1 computer has disabled swap
 - And, coincidence or not, server is problematic!
- Unlike Linux, swap sizes are comparable with RAM

	In Mb
Median	6838
Average	11297
Maximum	43008
Minimum	0

Recommendations:

- 1) Never disable paging file on Windows
- 2) Let Windows decide on the size
- 3) However, you can select special disk for paging file

I need feedback!

- ak@firebirdsql.org