

Linuxcgroup

Linux

BIOS Kernel MemTotal free RAM Linux

- free
- /proc/meminfo

Linux

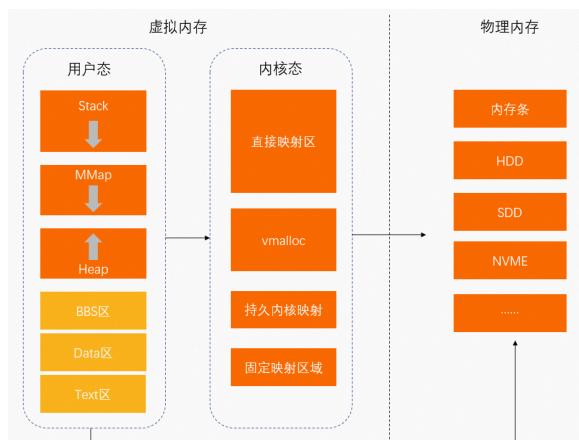
```
# = + +
total = used + free + buff/cache
```

Kernel

- PageCache

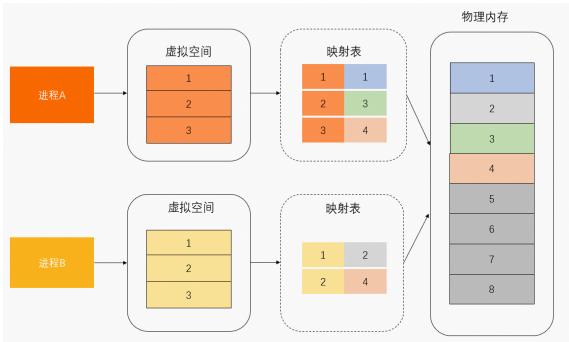
Content Menu

- Linux
 -
 - PageCache
 - top
 - ps
 - smem
 - cgroup
- DockerK8s
 - docker stat
 - kubectl top pod
-



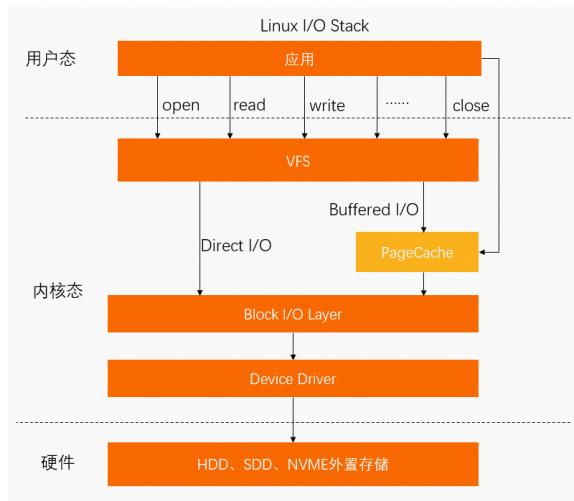
-
- - Stack
 - MMap(Memory Mapping Segment)
 - Heap
 - BBS
 - Data
 - Text
- MMap
 - VMALLOC
 -
 -

13A2B4AB



PageCache

MMapBuffered I/O SyscallPageCachePageCache



- anno_rssmalloc
- file_rssfile_rss
- shmem_rssshmem_rss



RSS(resident set size)

top



KB

```
top - 13:41:18 up 31 days, 3:46, 2 user, load average: 3.36, 2.52, 2.49
Tasks: 793 total, 1 running, 792 sleeping, 0 stopped, 0 zombie
Cpu(s): 4.4% us, 5.2% sy, 8.8% ni, 91.8% id, 0.5% wa, 0.3% hi, 0.4% si
Mem: 32G total, 24.8G used, 7.4G free, 0.4G swap
Swap: 0.0M total, 0.0M used, 0.0M free, 0.0M available swap[0m]
```

top	VIRT(Virtual Set Size)		
	RES(Resident Set Size)	RSS	anno_rss + file_rss + shmem_rss
	SHR(Shared Memory)		file_rss + shmem_rss
	%MEM		RES / MemTotal

ps



KB

```
[root@tcs-10-0-15-12 ~]# ps aux | more
  PID  User  Command          RSS   %CPU   %MEM
 7447 root    /init           0  0.0  684.0K
93523 root    /pause          0  0.0  684.0K
93595 root    /pause          0  0.0  684.0K
100294 root    /pause          0  0.0  684.0K
103300 root    /pause          0  0.0  684.0K
104534 root    /pause          0  0.0  684.0K
105666 root    /pause          0  0.0  684.0K
105188 root    /pause          0  0.0  684.0K
105609 root    /pause          0  0.0  684.0K
107013 root    /pause          0  0.0  684.0K
110834 root    /pause          0  0.0  684.0K
111448 root    /pause          0  0.0  684.0K
115251 root    /pause          0  0.0  684.0K
```

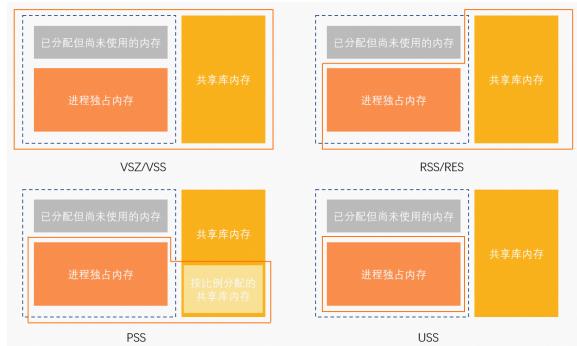
ps	VSZ(Virtual Set Size)		
	RSS(Resident Set Size)	RSS	anno_rss + file_rss + shmem_rss
	%MEM		RSS / MemTotal

smem



```
[root@tcs-10-0-15-12 ~]# smem -k
  PID  User  Command          USS   RSS
 7447 root    /init           0  684.0K
93523 root    /pause          0  684.0K
93595 root    /pause          0  684.0K
100294 root    /pause          0  684.0K
103300 root    /pause          0  684.0K
104534 root    /pause          0  684.0K
105666 root    /pause          0  684.0K
105188 root    /pause          0  684.0K
105609 root    /pause          0  684.0K
107013 root    /pause          0  684.0K
110834 root    /pause          0  684.0K
111448 root    /pause          0  684.0K
115251 root    /pause          0  684.0K
```

smem	USS(Unique Set Size)		anno_rss
	PSS(Proportional Set Size)		anno_rss + file_rss/m + shmem_rss/n
	RSS(Resident Set Size)	RSS	anno_rss + file_rss + shmem_rss

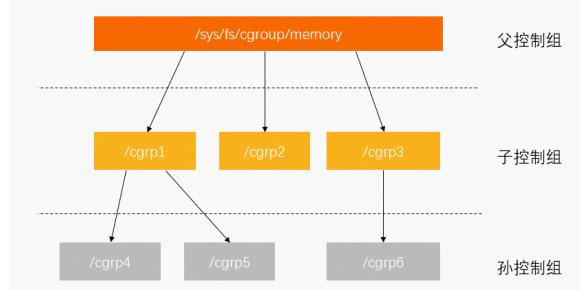




WSS(Memory Working Set Size) Linux Page Reclaim WSS

cgroup

cgroupLinux
cgroupMemory Control Group (memcg)



memory cgroup

```
cgroup.event_control      # eventfd
memory.usage_in_bytes     #
memory.limit_in_bytes     # /
memory.failcnt            #
memory.max_usage_in_bytes #
memory.soft_limit_in_bytes / 
memory.stat                # cgroup
memory.use_hierarchy       # /cgroupcgroup
memory.force_empty         # cgroup
memory.pressure_level     # cgroup.event_control
memory.swappiness          # swappiness
memory.move_charge_at_immigrate # cgroup
memory.oom_control         # /oom controls
memory.numa_stat            # numa
```

3

- memory.limit_in_bytes cgroupk8s docker memory limits
- memory.usage_in_bytes cgroup memory.stat RSS+Cache
- memory.stat cgroup

memory.stat	
cache	PageCache
rss	cgroupanno_rss
mapped_file	cgroupfile_rssshmem_rss
active_anon	LRUleast-recently-usedAnonymousSwap tmpfsshmem bytes
inactive_anon	LRUAnonymousSwap tmpfsshmembytes
active_file	LRUFile-backedbytes
inactive_file	LRUFile-backedbytes
unevictable	bytes

total_cgroup total_rss cgroupRSS

cgroup

- cgroupRSSanno_rssUSScgroupmapped_file+RSSRSS
- PageCachecgrou memcg PageCache

		cgroup(memcg)
RSS	anon_rss + file_rss shmem_rss	anon_rss
mapped_file		file_rss + shmem_rss
cache		PageCache

DockerK8s

DockerK8SLinux memcg

docker stat

```
CONTAINER ID  NAME          CPU %     MEM USAGE / LIMIT      MEM %     NET I/O           BLOCK I/O          PIDS
87c1d1871d2d  amazing_shtern  0.37%    475.7MiB / 1.939GiB  23.96%   2.26MiB / 860kB  4.99MiB / 77MiB  36
```

- LIMITmemory.limit_in_bytes
- MEM USAGEmemory.usage_in_bytes - memory.stat[total_cache]

 docker stat

kubectl top pod

kubectl topMetric-serverHeapsterCadvisorworking_setPodPauseMetrics-serverPod

```
func decodeMemory(target *resource.Quantity, memStats *stats.MemoryStats) error {
    if memStats == nil || memStats.WorkingSetBytes == nil {
        return fmt.Errorf("missing memory usage metric")
    }

    *target = *uint64Quantity(*memStats.WorkingSetBytes, 0)
    target.Format = resource.BinarySI

    return nil
}
```

Cadvisorworkingset

```

func setMemoryStats(s *cgroups.Stats, ret *info.ContainerStats) {
    ret.Memory.Usage = s.MemoryStats.Usage.Usage
    ret.Memory.MaxUsage = s.MemoryStats.Usage.MaxUsage
    ret.Memory.Failcnt = s.MemoryStats.Usage.Failcnt

    if s.MemoryStats.UseHierarchy {
        ret.Memory.Cache = s.MemoryStats.Stats["total_cache"]
        ret.Memory.RSS = s.MemoryStats.Stats["total_rss"]
        ret.Memory.Swap = s.MemoryStats.Stats["total_swap"]
        ret.Memory.MappedFile = s.MemoryStats.Stats["total_mapped_file"]
    } else {
        ret.Memory.Cache = s.MemoryStats.Stats["cache"]
        ret.Memory.RSS = s.MemoryStats.Stats["rss"]
        ret.Memory.Swap = s.MemoryStats.Stats["swap"]
        ret.Memory.MappedFile = s.MemoryStats.Stats["mapped_file"]
    }

    if v, ok := s.MemoryStats.Stats["pgfault"]; ok {
        ret.Memory.ContainerData.Pgfault = v
        ret.Memory.HierarchicalData.Pgfault = v
    }

    if v, ok := s.MemoryStats.Stats["pgmajfault"]; ok {
        ret.Memory.ContainerData.Pgmajfault = v
        ret.Memory.HierarchicalData.Pgmajfault = v
    }

    workingSet := ret.Memory.Usage
    if v, ok := s.MemoryStats.Stats["total_inactive_file"]; ok {
        if workingSet < v {
            workingSet = 0
        } else {
            workingSet -= v
        }
    }
    ret.Memory.WorkingSet = workingSet
}

```

```
kubectl top podMemory Usage = Memory WorkingSet = memory.usage_in_bytes -
memory.stat[total_inactive_file]
```

		Memory Usage
docker stat	Docker	memory.usage_in_bytes - memory.stat[total_cache]
kubectl top pod	K8s	memory.usage_in_bytes - memory.stat[total_inactive_file]

toppsMemory Usagetoppss

Memcg	rss + cacheactive cache + inactive cache
Docker	rss
K8s	rss + active cache

- <https://www.alibabacloud.com/help/zh/arms/application-monitoring/memory-metrics>
- <http://hustcat.github.io/memory-usage-in-process-and-cgroup>
- <https://www.51cto.com/article/692936.html>
- <https://itnext.io/from-rss-to-wss-navigating-the-depths-of-kubernetes-memory-metrics-4d7d77d8fdcb>

