

Here's a tentative development plan on resource management from our side.
Would request others to comment/add their plans as part of this discussion.

This plan, alongwith appropriate usage scenarios, is being requested for kernel-summit discussion. Most likely, the resource management plans will be presented as part of container discussion.

1. Infrastructure

Userland interface for task-grouping, specifying group usage limits etc

Paul Menage's container patches will be the basis for this infrastructure.

2. cpu controller

- Atleast 1-level deep hierarchical proportional fair-share scheduling
- Soft and Hard limit cpu usage of group
- Fine grained cpu control (ex: parts per 1000?)
- SCHED_FIFO policy for groups
 - Groups having this policy will recv all cpu cycles as long as they have any task runnable. This is at the cost of starving other non-SCHED_FIFO groups.
- Sleeper fairness - Will help bursty workloads
- Accounting - group cpu usage

3. memory controller

- Basic rss controller
- user space oom handler
- Improve shared page accounting
- Handling of soft limits
- Provide support for accounting kernel/user memory, slabs, page tables, dcache entries, vma's, etc
- Improved statistics (folding)
- Add support for better tuning of containers
 - i. Low and high water marks per container
 - ii. Per container swappiness support
- Per container swap file(s)

4. Disk io controller

[I don't have sufficient details at the moment on disk io controller requirements. Will add after I talk to relevant groups here]

--

Regards,
vatsa

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
