

BoF: Memory Efficiency on Modern Linux Systems

Mykolas Krupauskas, LPC 2026 October


BoF

- It's an open discussion, propose questions
- Briefly introduce yourself, your organization, and your operational experience
- The outcome we want is some action items
- What can we do to improve this for the community?

Background

- Memory prices vertical (4-5x in the last 18 months)
- Utilization in the overall industry is not high, what can we do to help

DDR4-3200 2x8GB (Average price in USD over last 18 months)



Source: <https://pcpartpicker.com/trends/price/memory>

Discussion questions

1. What pain points do admins have trying to get to a higher memory utilization?
2. How do we better answer the question of how much memory a workload needs?
3. What should we prioritize as a community?
What has the highest ROI?

Example datacenter context

- Footprint many PBs
- Utilization 35% in aggregate (non-available memory)
- Running heterogeneous hardware and container workloads (kubernetes, containers)
- The main challenge is how do we rightsize workloads to memory allocation
- How do we know how much memory a workload needs?
- We plan the capacity we buy and in cloud we can resize VMs

Session notes (key takeaways)

- Memory sizing remains difficult: Understanding how much memory a workload actually needs is still a central challenge.
- Compression helps, but measurement is hard: Meta reported broad zswap deployment. Quantifying benefits and benchmarking remain challenging.
- Reclaim needs improvement: Better LRU policies and proactive reclaim, without pushing workloads too hard. Swap performance improvements.
- Pressure handling must survive pressure: Userspace daemons can stall under heavy pressure. Triggers, profiling before killing, graceful degradation, ebpf oom killer, and protecting high-priority workloads were discussed.

Session notes (key takeaways)

- Data centers, embedded systems, Android, and Chromebooks all want to do more with less memory.
- Historically OS optimized for performance with a fixed memory capacity rather than memory efficiency. In a data center environment you want to rightsize memory capacity to a certain workload performance.
- Memory-aware scheduling and host-guest coordination: swap spikes, central memory visibility with node-level decisions, and VM memory redistribution, with guest cooperation still debated.
- We should continue improving: [LRU performance](#), [swap readahead](#), [DAMON](#).