

Week10: My VM is Lighter (and Safer) than your Container

- What are the advantages and disadvantages of Tinyx over unikernels?
- How does LightVM balance isolation and performance compared to containers?
- The paper mentions that the XenStore introduces significant overhead as the number of VMs increases. What alternative design principles could further reduce these bottlenecks in hypervisor communication?

Week9: Operating System Transactions and RCU

- How does TxOS handle conflicts between transactional and non-transactional threads while maintaining system responsiveness?
- What are the disadvantages of the asymmetric conflict resolution scheme used in TxOS?
- TxOS uses lazy versioning instead of eager logging. What are the key benefits of this choice, and in which scenarios might it fail?

Week8: F2FS: A new File System for Flash Storage

- What are the design choices made by F2FS that are specific to flash storage?
- What is the wandering tree problem and how does F2FS prevent it in its design? What is the design philosophy behind it?
- Explain why LFS uses the greedy victim selection policy for its foreground cleaning? Why does this policy minimize the latency visible to applications?
- How does F2FS's adaptive logging scheme help mitigate the write amplification problem commonly faced in flash storage systems?

Week7: The Design and Implementation of a Log-structured File System

- What are the design tradeoffs made by LFS that can pose challenges in the broader adaptation of the LFS? (e.g. types of workload, the storage medium, how the space is reutilized)
- How does LFS preserve consistency in case of a crash? Is there a way to accelerate the recovery process of a LFS?
- How to decide the buffer size to maximize the bandwidth? What can be the drawback of using a large buffer?
- What are the advantages and disadvantages of a LFS in an SSD? Would a LFS be more efficient than a traditional file system in an SSD?

Week6: TMO: Transparent Memory Offloading in Data Centers

- When using SSDs as a memory tier, how does TMO prevent them from being worn out when used for offloading memory?
- When considering offloading, what are the aspects that TMO has considered and what aspects TMO can do better? Considering the following aspects: the time to offload, the amount of data to offload, what data to offload.

- How does TMO perform memory offloading to compressed memory versus SSD, and under what conditions does it choose between the two?
- What are the disadvantages of implementing Senpai in userspace? Can you think of any scenarios in which a kernel space implementation of Senpai would be more performant?

Week5: Practical, transparent operating system support for superpages

- It is known that x86 has a limited number of superpage TLB entries in the per-core TLB structure. Suppose superpages were used pervasively throughout an application, what would be the downside of such an approach, and for what kind of applications? Think in terms of reduced entries available in the TLB if only superpages are used, and how the working set of an application interacts with TLB usage.
- Consider the scenario in which a superpage is read from disk and only a small fraction of it is modified. Propose a system that would be able to only write the necessary base pages to disk without having to demote the superpage.
- Consider a case where the system dram is not enough and you have to swap. What is the performance impact of superpages on swap? And the impact on space efficiency? How would you mitigate these issues caused by superpages?

Week4: ghOST: Fast & Flexible User-Space Delegation of Linux Scheduling

- List some disadvantages of ghOST's delegated scheduling model (i.e., offloading decisions to the user space ghOST agent). Explain how these limitations can be addressed by implementing all of the scheduling logic within BPF programs instead.
- How does ghOST work help minimize system downtime during scheduler updates and rollbacks? How does ghOST ensure fault tolerance when an agent fails or a policy malfunctions?
- Why is eBPF not sufficient to implement a flexible scheduling policy here? How does ghOST overcome the limitations while still leveraging the eBPF? In redesigning or modifying eBPF, what changes will you perform to offload the whole ghOST logic into the kernel and what can be the tradeoffs?
- How is isolation achieved between different applications using ghOST to write scheduling policies?

Week 3: The Linux Scheduler: A Decade of Wasted Cores

- Issue: Choosing a process that is asleep for a long time
Example: A is a CPU-bound process, B is a IO-bound process and is asleep for 10 seconds. When B wakes up, its vruntime is 10 seconds behind A's vruntime
B may end up monopolizing the CPU for the next 10 seconds while it catches up, effectively starving A. Q. Will this situation occur?

- Comment on whether CFS is considered a work-conserving scheduler. Explain an example scenario of thread scheduling where the behavior of a work-conserving and non-work-conserving scheduler will differ.
- **Suppose there is a scheduler which enabled full preemption like Linux and another scheduler with partial preemption only (user threads cannot preempt each other). Is the preemptable kernel always better than the non-preemptable one? If not, explain a case where it is actually not (The workload consist of a load injector and workers that handle requests. Load injector got preempted)**
- Explain one of the bugs described in the paper and why it affects the overall performance of the Linux scheduler. Can you think of any potential issues with the proposed fix for the bug? If not, explain why the solution is optimal.

Week 2: Light-Weight Contexts: An OS Abstraction for Safety and Performance

- What is the difference between lWC and process? Why is creating lWC much cheaper than forking a process? How would you compare the cost of creating lWC against creating other light weight asynchronous execution units, e.g., C++ coroutines?
- Why does switching between lWCs faster than switching between kernel threads?
- How do snapshots and rollbacks work for lWC? Is it possible to reuse the original lWC instead of recreating the snapshot every time to reduce the overhead?