

Resource Management on Embedded Linux Gaming Devices

Stefan Bossbaly
sboss <at> meta.com

Linux Plumbers Conference 2025



Agenda

1

Why Resource Management is
Important

2

What Mechanisms Linux Gives
Us

3

How Android Handles Resource
Management

4

Expanding on that
implementation with
Orchestrator

5

Our Success

6

Our Plans for the Future

Why is Resource Management Important

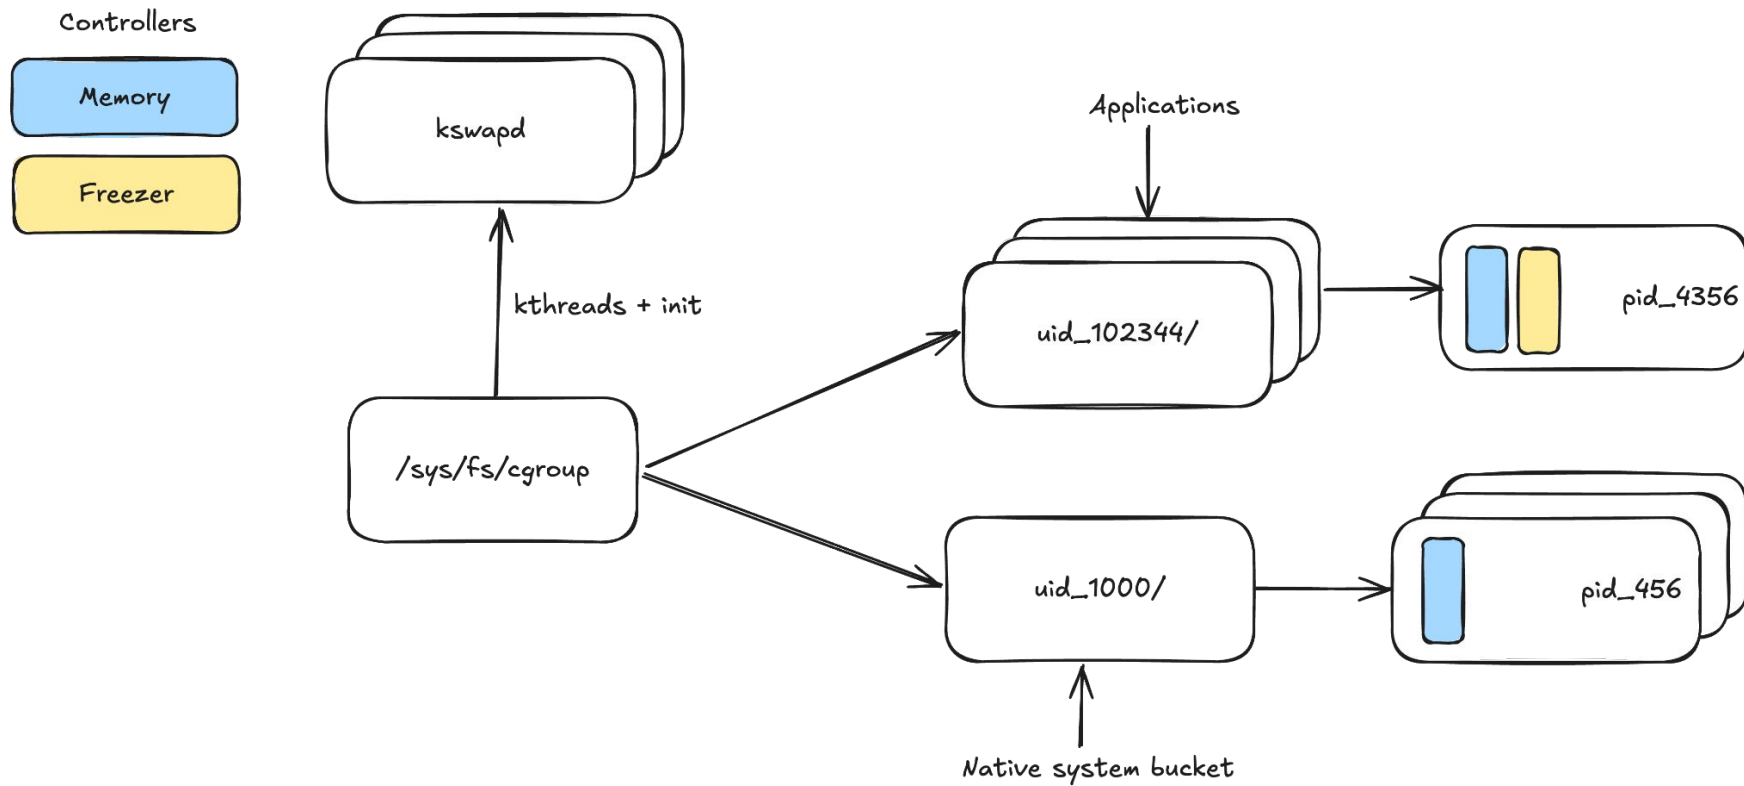
- Embedded devices are limited in their capabilities
 - Hardware
 - Thermals
 - Energy
- Embedded devices are generally multi purpose
 - Noisy Neighbors
 - Mixed criticality workloads
- Cost Optimization
 - Hardware is expensive
 - More devices to more people

What Are Our Goals?

- Accurate collection of resource usage
 - Attribute to a process
 - Low overhead
- Way to control the resource usage per process
- Sole Decision Maker
- Telemetry
 - Log resource usage, detect regressions
 - Log enforcement actions
- Experimentation
 - Trade-offs

Linux Provided Mechanisms

- Cgroupv2
 - Introduced unified hierarchy over cgroupv1
 - Up to the implementation to organise processes
 - Provides control “knobs” and read-only files that are generic across all devices
 - Controllers are hardware agnostic
- BPF
 - Allows the running of user-specified code at known hook points
 - Aligns perfectly with resource monitoring
 - Can be driver specific, adding “non-standard” hook points
 - Up to the implementation to ensure correctness

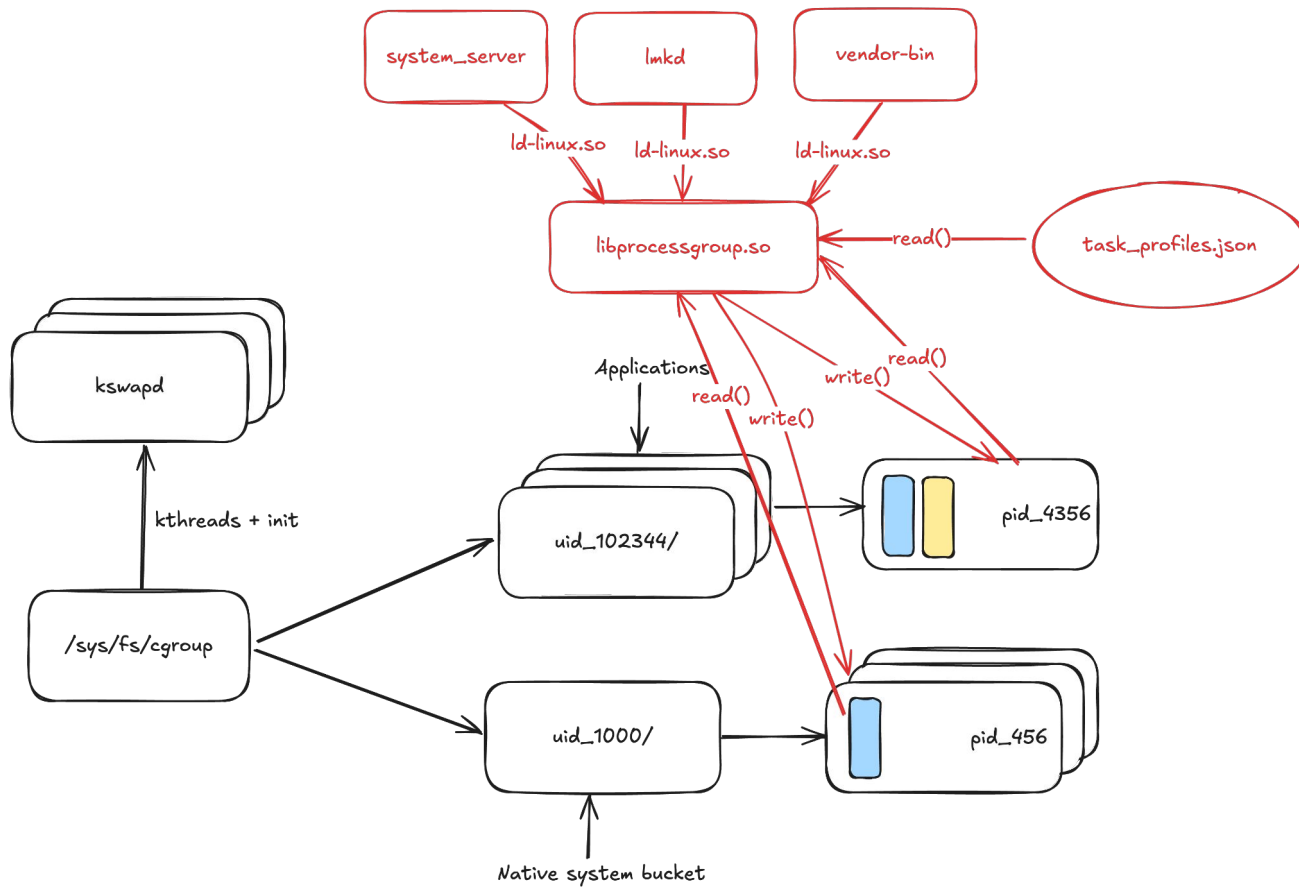


How does Android Controls the Knobs

- Uses libprocessgroup, a cgroup abstraction layer
- Source of truth: cgroups.json & task_profiles.json
- Terminology
 - Attributes
 - Actions
 - Profiles
 - Aggregate Profiles
- Profile and Aggregate Profiles can then be applied to processes
- system_server applies resource for application
- Native Services apply them to themselves

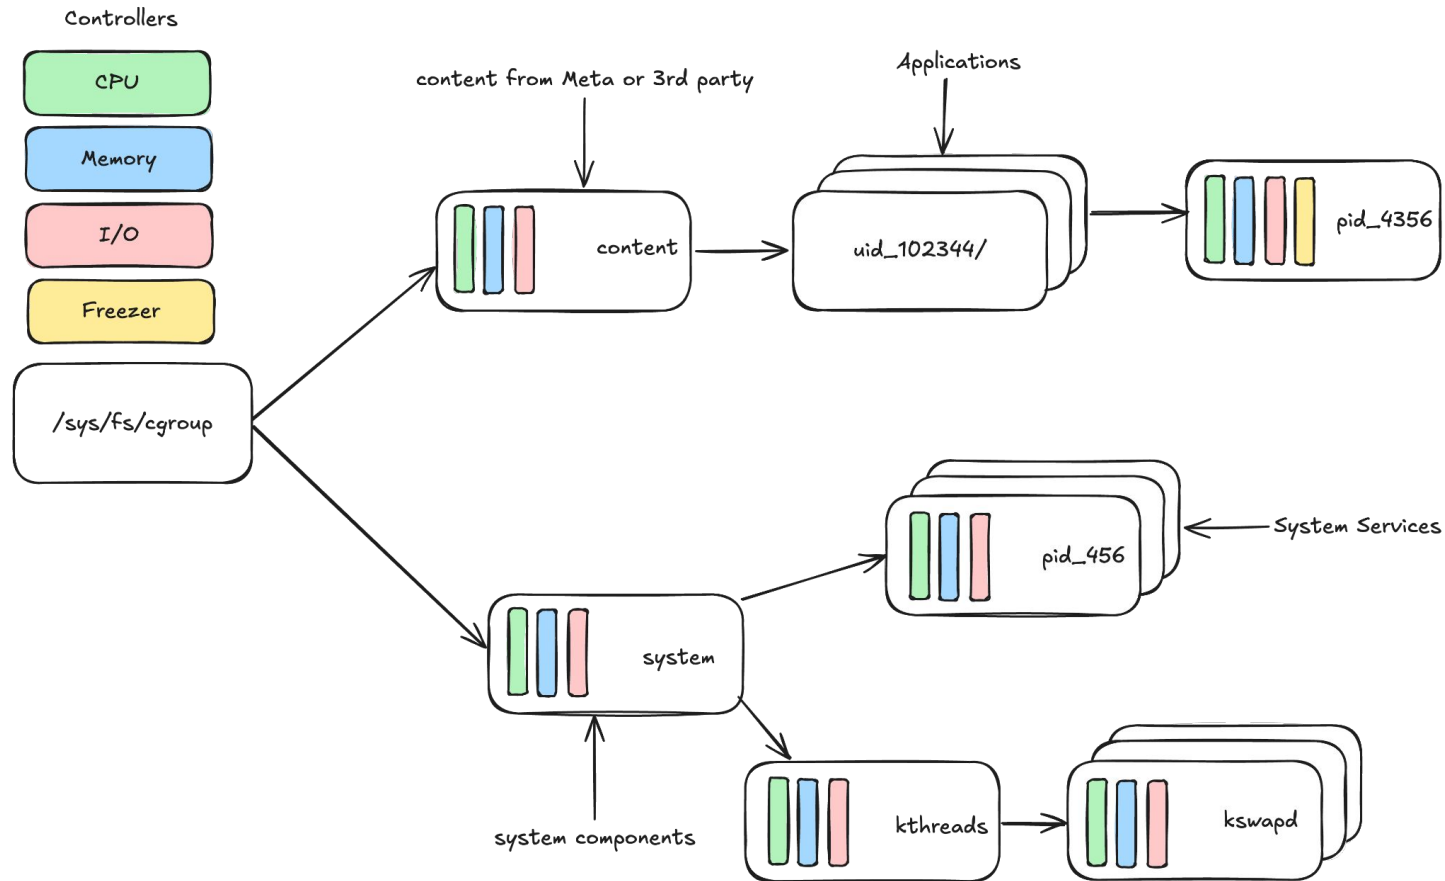
Example task_profiles.json

```
{
  "Attributes": [
    { "Name": "MemcgHigh", "Controller": "memory", "File": "memory.high" },
    { "Name": "MemcgMax", "Controller": "memory", "File": "memory.limit_in_bytes", "FileV2": "memory.max" }
  ],
  "Profiles": [
    {
      "Name": "SCHED_SP_BACKGROUND",
      "Actions": [
        { "Name": "SetAttribute", "Params": { "Name": "MemcgHigh", "Value": "128MB" } },
        { "Name": "SetAttribute", "Params": { "Name": "MemcgMax", "Value": "256MB" } },
        { "Name": "JoinCgroup", "Params": { "Controller": "cpu", "Path": "background" } }
      ]
    }
  ]
}
```

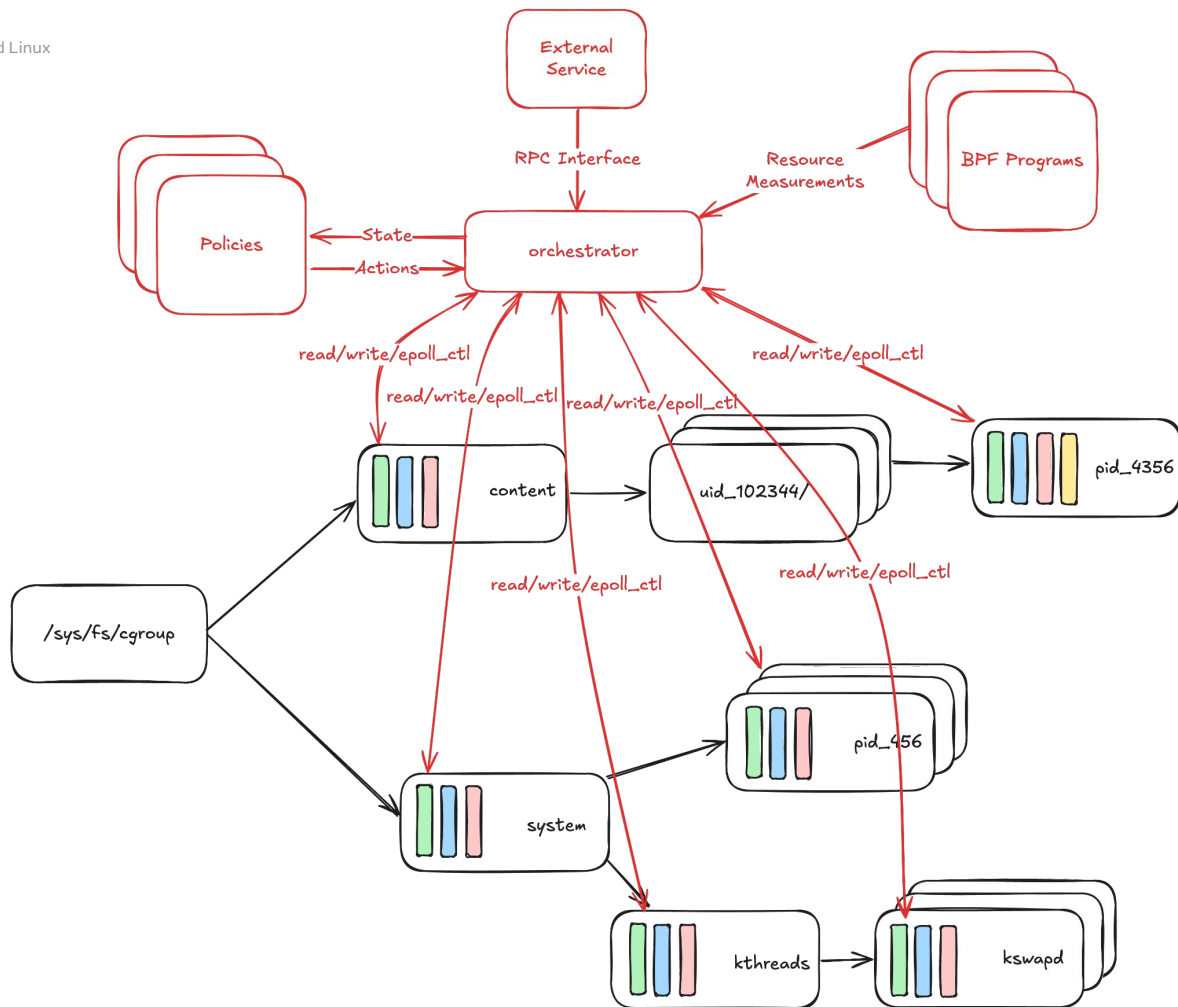
Expanding libprocessgroup

- Benefits
 - Cgroup Abstraction
 - Centralizes Configuration
 - Vendor Overriding
- Limitations
 - Cgroup abstraction is incomplete
 - cgroupv1 cpuctl implementation
- Resource Management as shared library
 - Deconfliction is nearly impossible
- Experimentation
 - A/B Dynamic Configuration



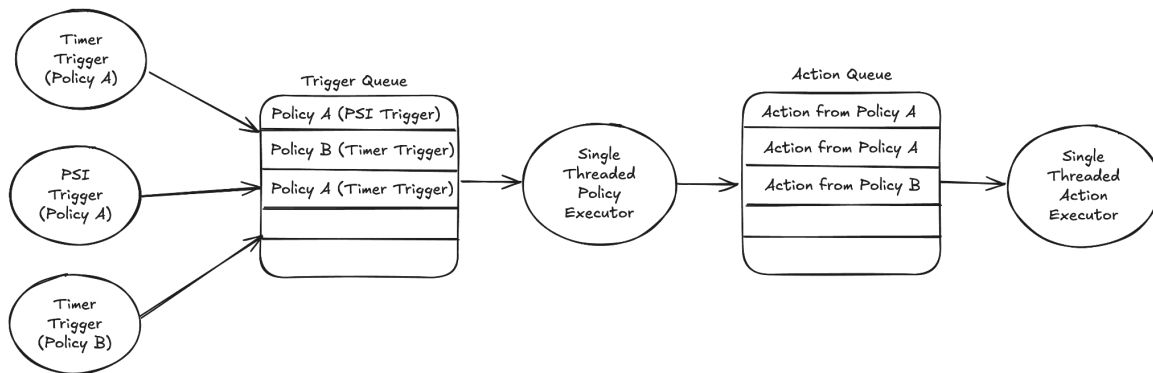
Implementing a Centralized Approach to Resource Control

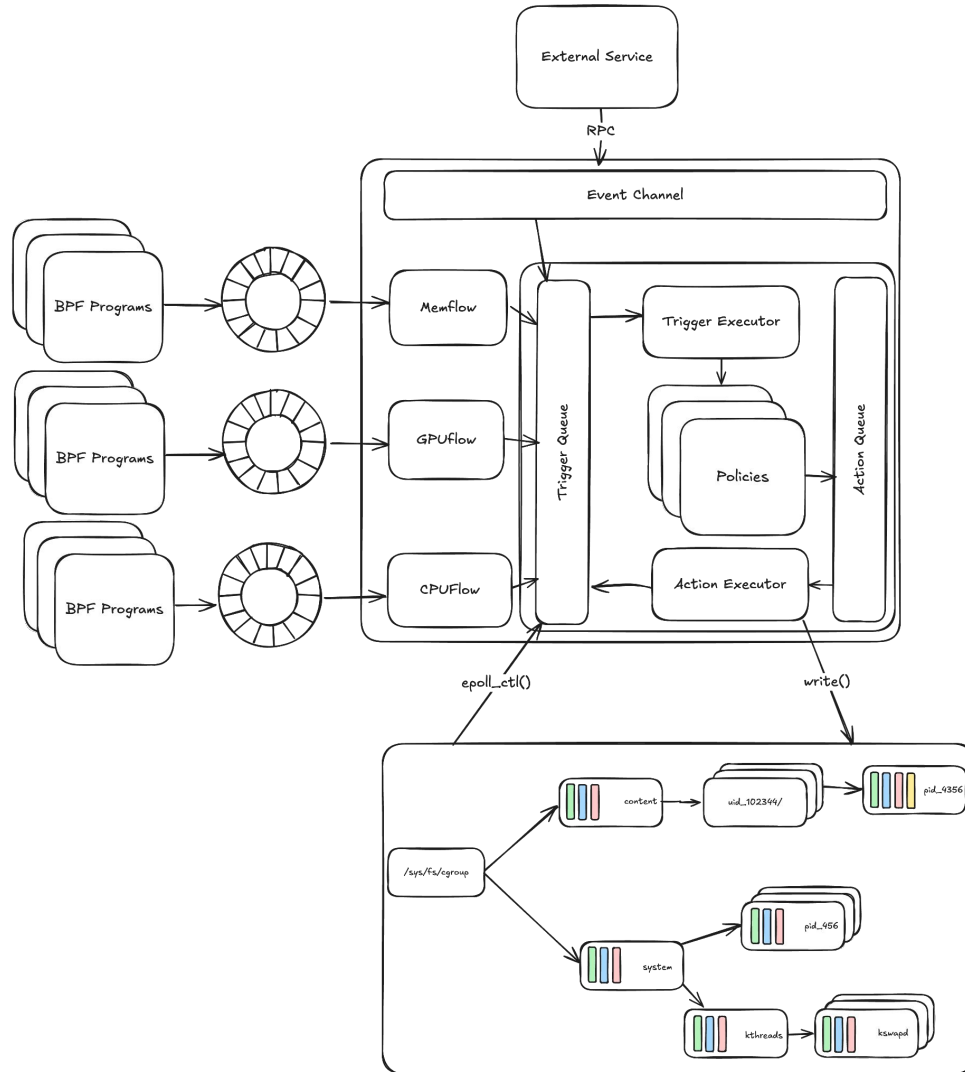
- Rust based policy framework called Orchestrator
- Solely responsible for the management of cgroup hierarchy
- Components
 - Cgroupv2 controller interface
 - Policy Manager and Executor
 - BPF Programs for Metric Collection
 - RPC interface for applications to provide context
- Goals
 - Policy Framework with reusable components
 - Unify the cgroup management into one process
 - Unification of resource monitoring and resource management



Policy Based Engine

- Terminology
 - Policy
 - Policy Factory
 - Trigger
 - Action
- Framework Benefits
 - Real-time guarantee
 - Clear Responsibilities
 - Deconflicting
 - Experimentation






```
1  define_policy! {  
2      MemoryPolicy {  
3          description: "A test policy to show case the define_policy! schemantics",  
4          state: {  
5              count: u32,  
6          },  
7          triggers: {  
8              (PsiHigh: PsiTrigger),  
9              (PsiLow: PsiTrigger),  
10             (Timer: TimerTrigger),  
11         },  
12         queue: Some(Queue::Memory),  
13         execute: |trigger_event| {  
14             match trigger_event {  
15                 Self::PsiHigh(_psi_data) => {  
16                     Some(Action::KillLowPriorityProcesses)  
17                 }  
18                 Self::PsiLow(_psi_data) => {  
19                     self.count += 1;  
20                     None  
21                 }  
22                 Self::Timer(duration) => {  
23                     let current_count = self.count;  
24                     self.count = 0;  
25                     if current_count > 10 {  
26                         Some(Action::PerformMemoryAudit)  
27                     } else {  
28                         None  
29                     }  
30                 }  
31             }  
32         }  
33     }  
34 }
```


Currently Deployed at Meta

- Orchestrator is still in its infancy
- Memory resource is first to be implemented
- Telemetry Policy
 - Provides high level view of memory usage per process
 - Leverages BPF provided data
- Heap Type Telemetry Policy
 - Provides detailed view of heap type allocations per process
- Budget Service Policy
 - Leverages cgroupv2 knobs
 - Tracks memory for all sources
 - Solves the noisy neighbor problem

The Future

- Transparent Memory Offloading Policy
 - Swapping of unused pages
 - Pressure Stall Information
- CPU Policies
 - Budget Policy
- GPU Policies
 - Budget Policy
- Open Sourced
 - [Github](#)



Source: [Engineering at Meta](#)

Thank you.