

more robust
bpf memory allocator

bpf memory allocator

- An any-context allocator for bpf program
 - Process context
 - IRQ context
 - NMI context
 - Nested context (Process + IRQ + NMI context)

immediate reuse

- Desc
 - The freed element may be reuse immediately
 - hash table lookup may return incorrect result
 - bpf_cpumask/qp-trie
- Possible solution
 - BPF_MA_REUSE_AFTER_GP: reuse after one RCU GP and free after another tasks-trace RCU GP
 - increased memory usag
 - degraded performance

bpf global ma

- Desc
 - allocation from root cg
 - any need to limit its total usage ?
 - no destruction and only one pending free callback
 - no free call on CPU X, the pending free elements will not be freed
- Possible solution
 - shrinker ?
 - timer for checking the freeable elements ?

CPU hotplug

- Desc
 - slub support CPU hotplug, do it for bpf ma as well
 - does irq work support CPU hotplug ?
 - refill all bpf ma again when CPU is online again ?
- Possible Solution
 - ?