

# CAP\_BPF and CAP\_TRACING

Alexei Starovoitov, Song Liu

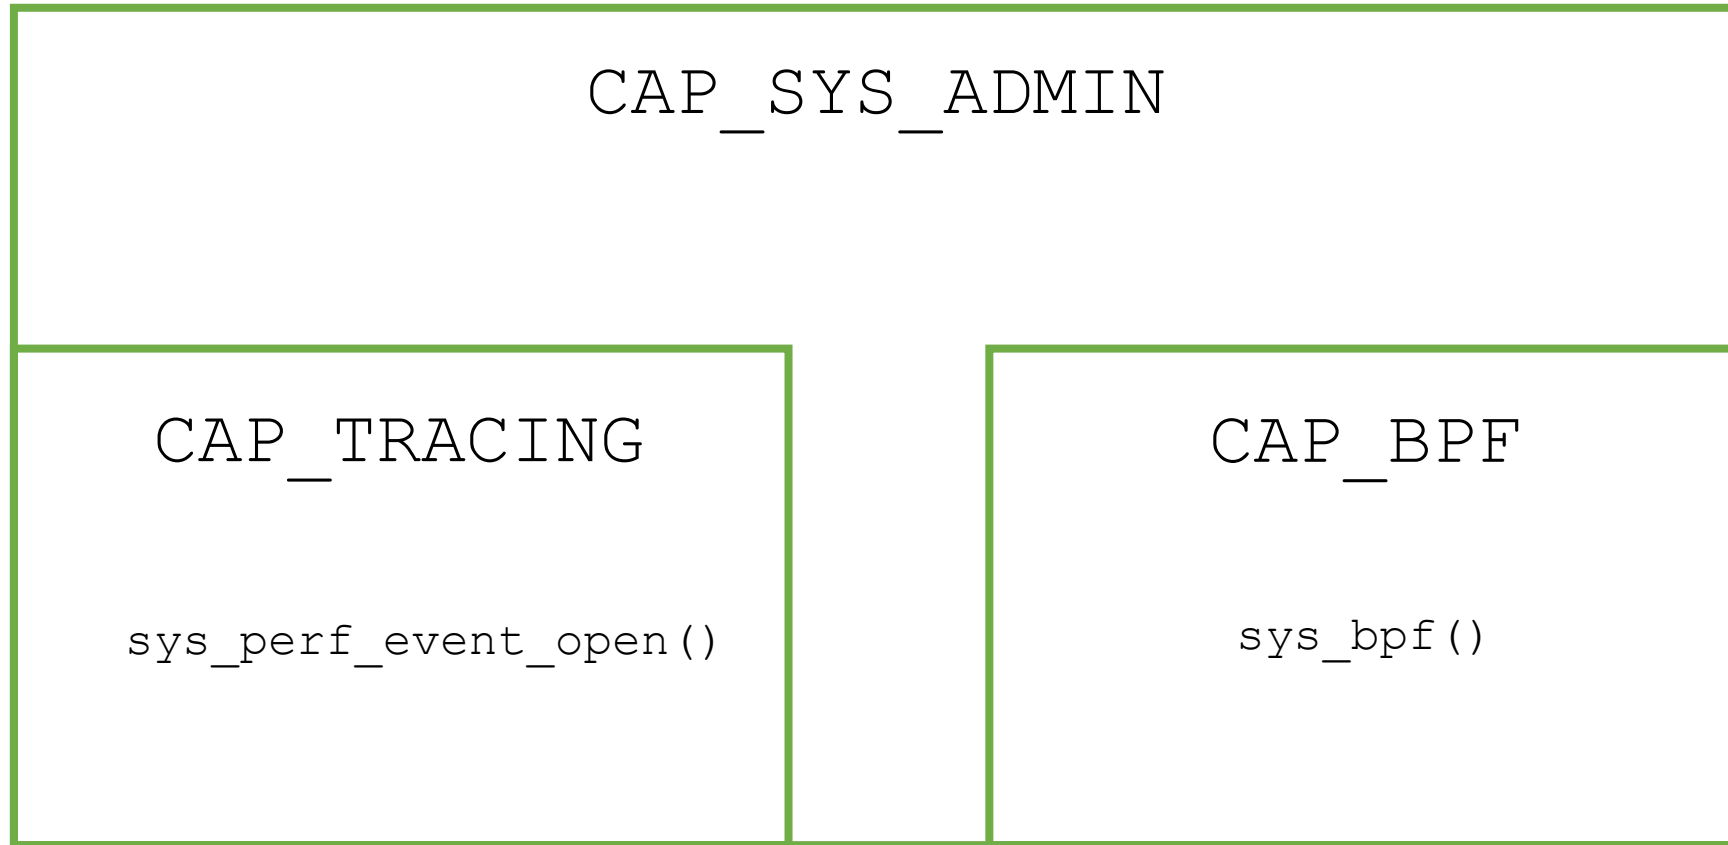
# Overview

`CAP_SYS_ADMIN`

`sys_perf_event_open()`

`sys_bpf()`

# Overview



# Use Cases

- BPF tracing users
  - `CAP_BPF + CAP_TRACING`
  
- BPF networking users
  - `CAP_BPF + CAP_NET_ADMIN`

# sys\_bpf() map

only CAP_BPF	+ CAP_TRACING	+ CAP_NET_ADMIN	+ CAP_SYS_ADMIN
BPF_MAP_CREATE <b>except</b>	<b>create</b> stackmap	<b>create</b> devmap	<b>create</b> cpumap
BPF_PROG_LOAD	BPF_RAW_TRACEPOINT_OPEN	BPF_PROG_ATTACH	<b>hardware offload</b>
BPF_BTF_LOAD	BPF_TASK_FD_QUERY	BPF_PROG_DETACH	
BPF_*_GET_NEXT_ID		BPF_PROG_QUERY	
BPF_*_GET_FD_BY_ID		BPF_PROG_TEST_RUN	
BPF_OBJ_GET_INFO_BY_FD			

**with map FD**

BPF\_MAP\_\*

**with access to bpfes**

BPF\_OBJ\_PIN

BPF\_OBJ\_GET

Do **NOT** `setcap cap_bpf bpftool`

- **NOR**

- `setcap cap_bpf bpftrace`
- `setcap cap_tracing perf`

- Existing tools are too powerful for untrusted users
- `CAP_BPF` and `CAP_TRACING` make it possible to build `bpf/perf/tracing` tools for untrusted users
  - Modified `funcount.py`