# Live Kernel Patching vs BPF

How to play along

# Problem statement

Live kernel patching uses ftrace to “hijack” the function call to call a updated function. In doing so, it sets the IP_MODIFY flag because it modifies where the ftrace trampoline will return to which can only be done by one user attached to a specific function.

BPF uses direct calls which must also set the IP_MODIFY flag as the return of the ftrace trampoline will go to the BPF direct trampoline. This prevents BPF and Live Kernel patching from operating on the same function.

# BPF direct trampoline

```
<kernel_func>:
nop
[..]
```

# BPF direct trampoline (solo)

# BPF direct trampoline (shared)

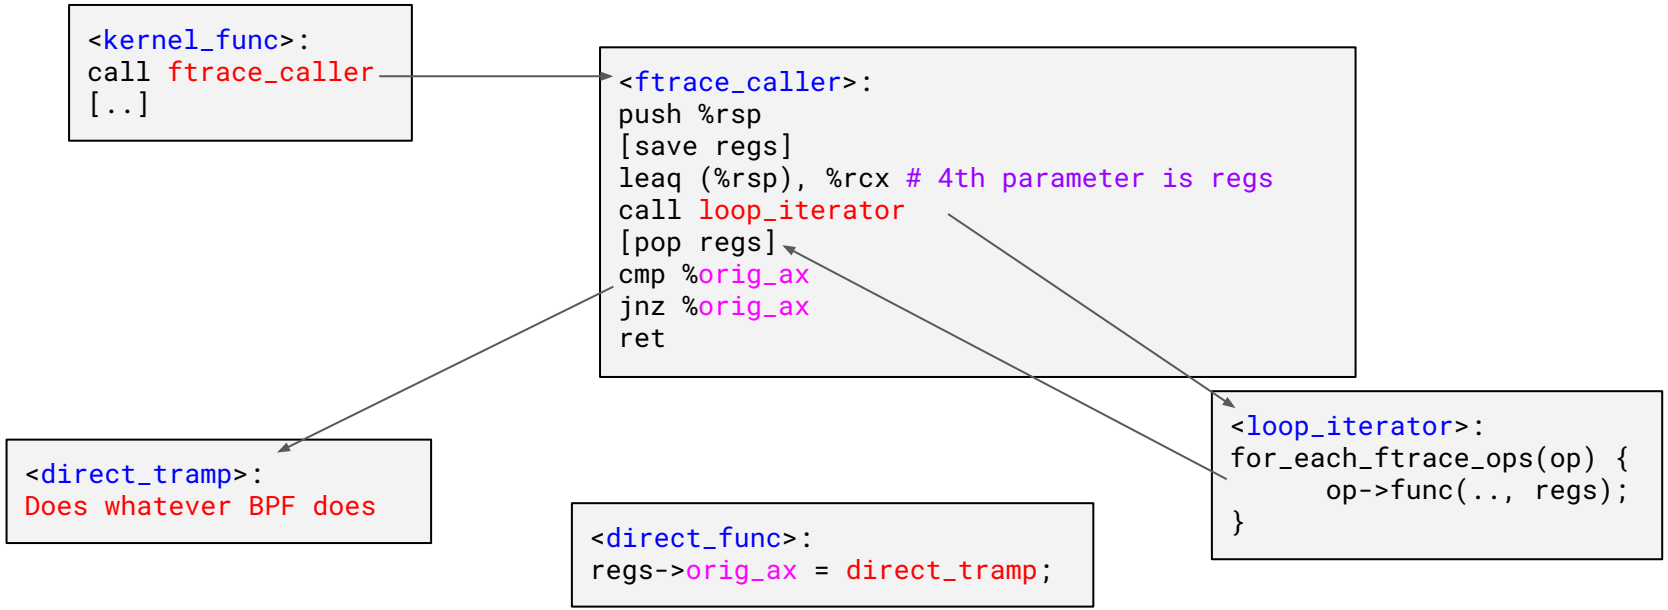

# Live kernel patching

```
<kernel_func>:
nop
[..]
```

# Live kernel patching

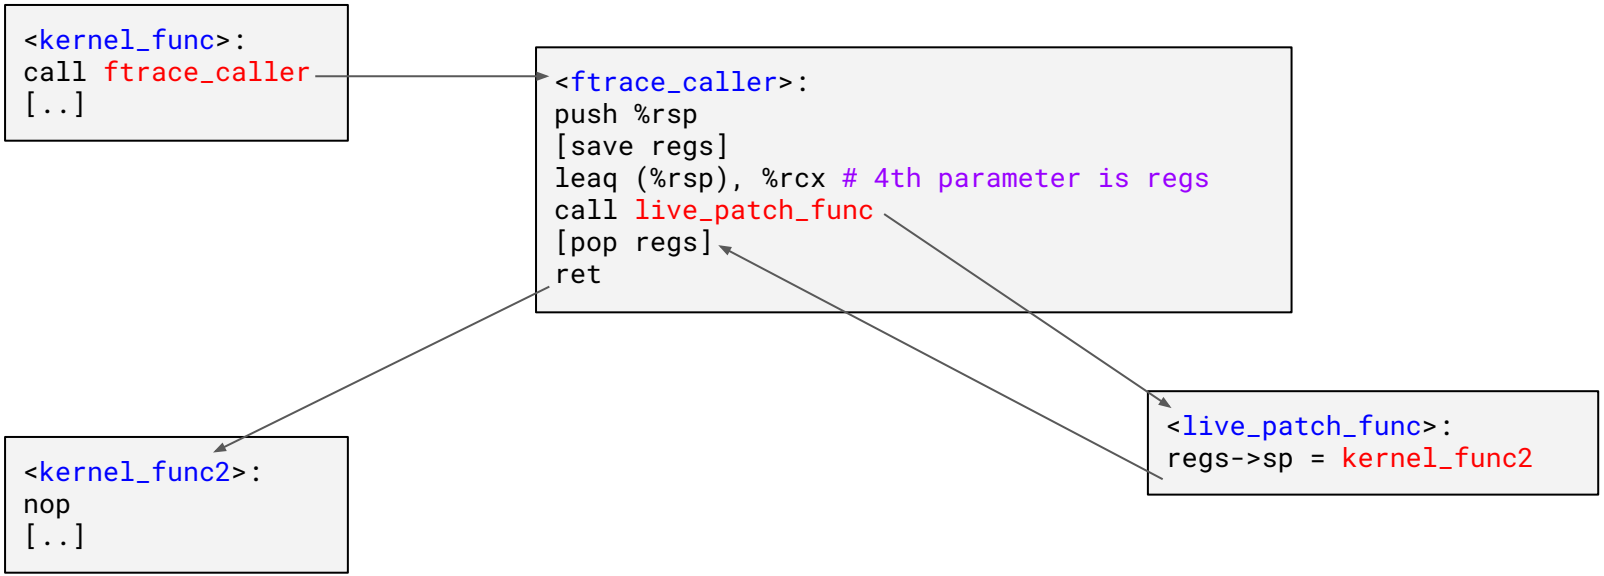

# Live kernel patching with BPF direct trampolines

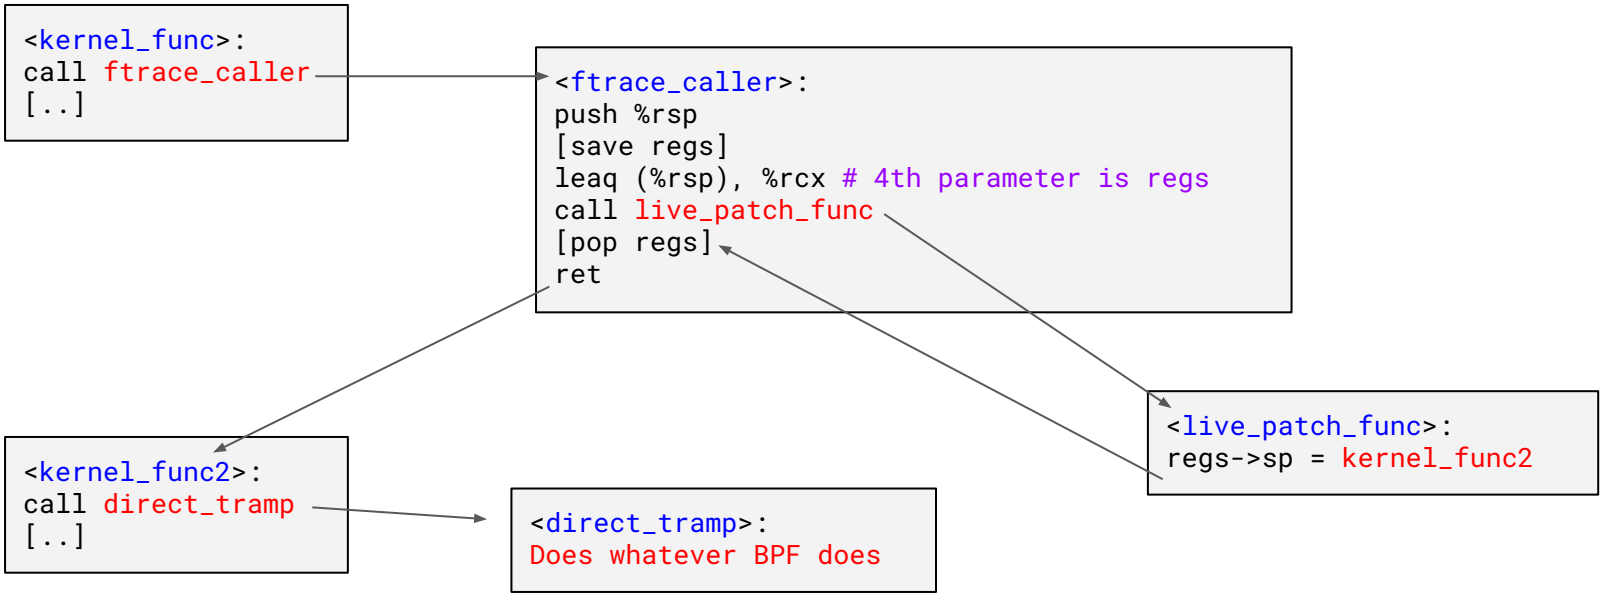

# Live kernel patching with BPF direct trampolines

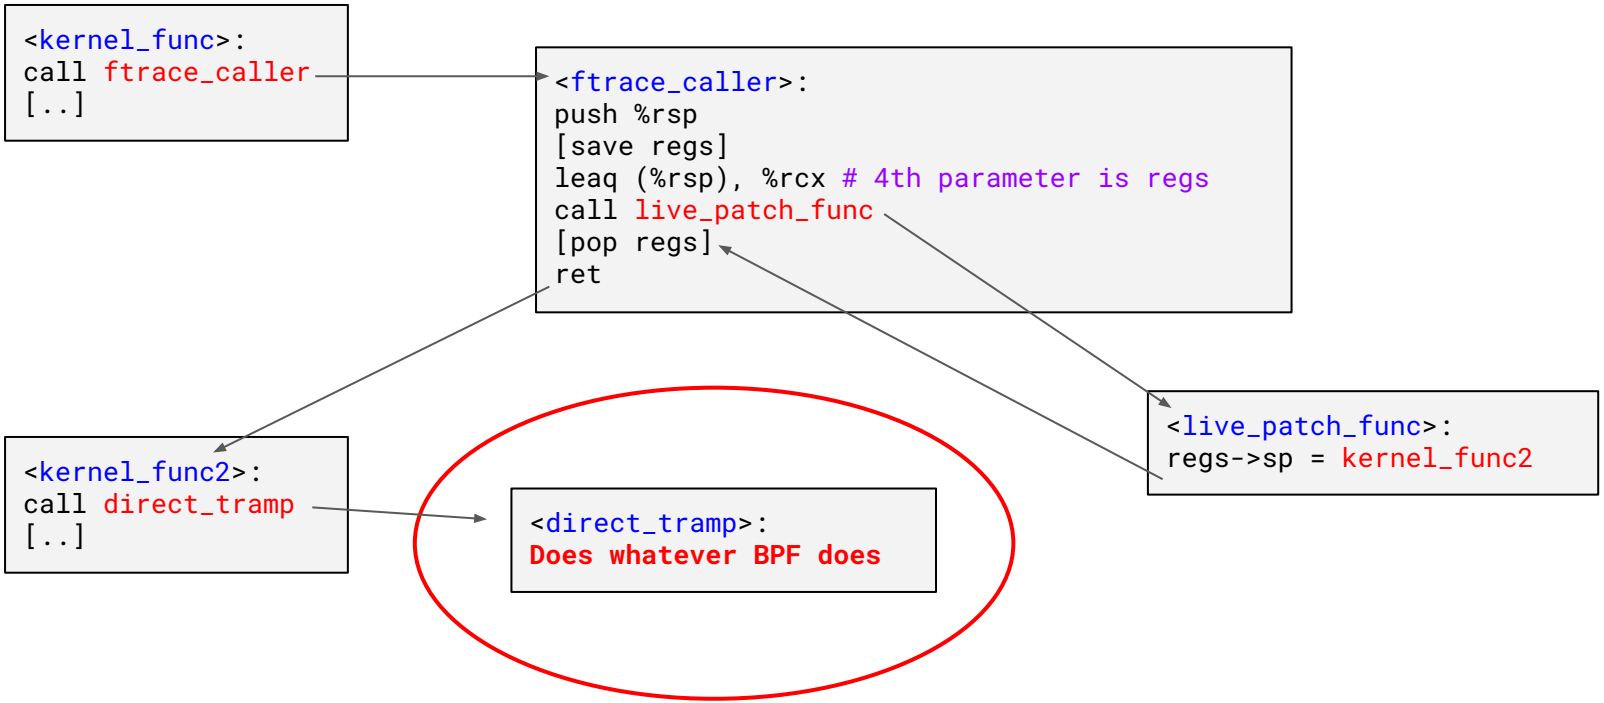

# Live kernel patching with BPF direct trampolines

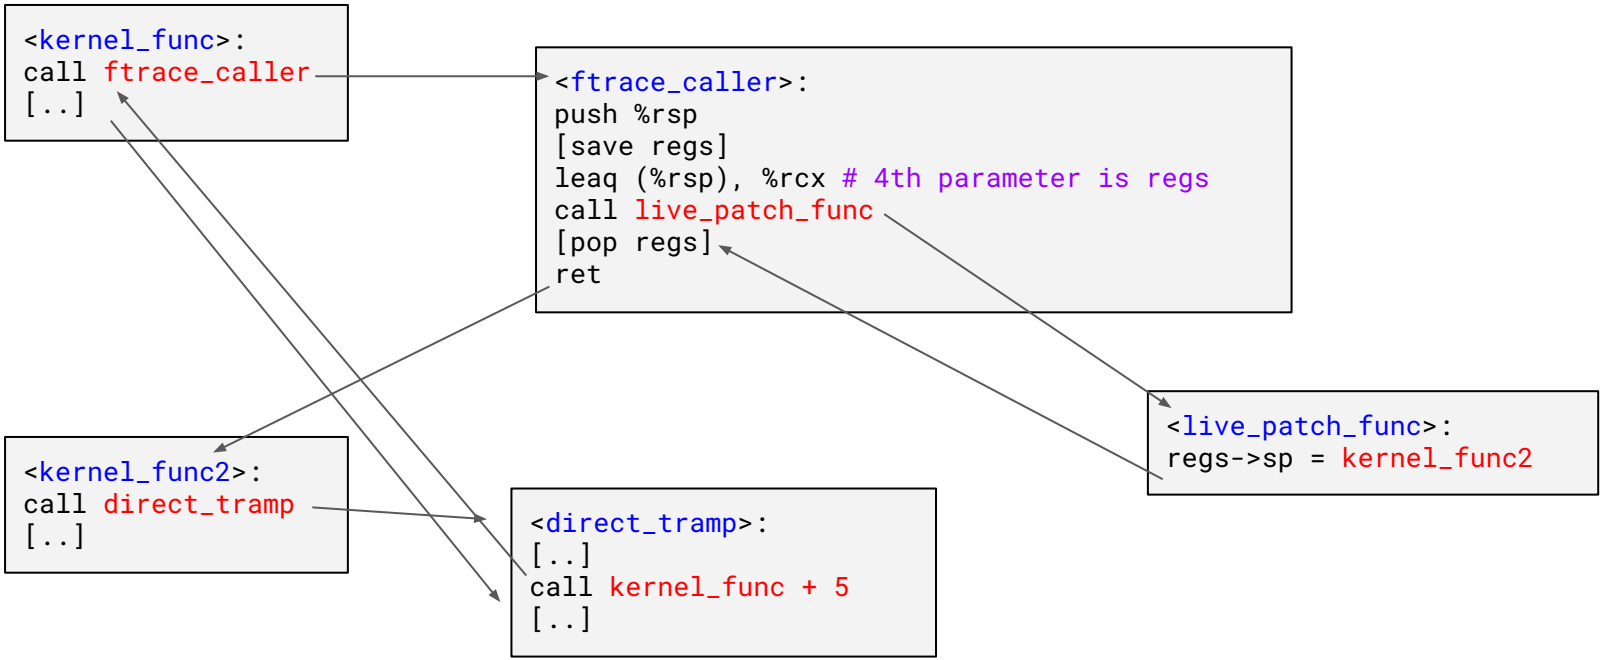

# Live kernel patching with BPF direct trampolines