

LinuxCon NA 2014
(Aug. 2014)

Kpatch Without Stop Machine

The Next Step of Kernel Live Patching

Masami Hiramatsu
<masami.hiramatsu.pt@hitachi.com>
Linux Technology Research Center
Yokohama Research Lab. Hitachi Ltd.,

Yokohama Research Lab.
Linux Technology Center



- Masami Hiramatsu
 - A researcher, working for Hitachi
 - Researching many RAS features
 - A linux kprobes-related maintainer
 - Ftrace dynamic kernel event (a.k.a. kprobe-tracer)
 - Perf probe (a tool to set up the dynamic events)
 - X86 instruction decoder (in kernel)

Background Story

Kpatch internal

Kpatch without stop_machine

Conclusion and Discussion

Note: this presentation is only focusing on the kernel-module side of kpatch.

More generic design and implementation, please attend to -

kpatch: Have Your Security And Eat It Too! – Josh Poimboeuf

Aug. 22. pm2:30

Background Story

What is kpatch?

Live patching requirements

Major updates and Minor updates

Kpatch internal

Kpatch Overview

Active Safeness check

Stop_machine

Kpatch without stop_machine

Live Patching Rules

Kpatch Reference Counter

Safeness check without stop_machine


Conclusion and Discussion

- Kpatch is a LIVE patching function for kernel
 - This applies a binary patch to kernel on-line
 - Patching is done without shutdown

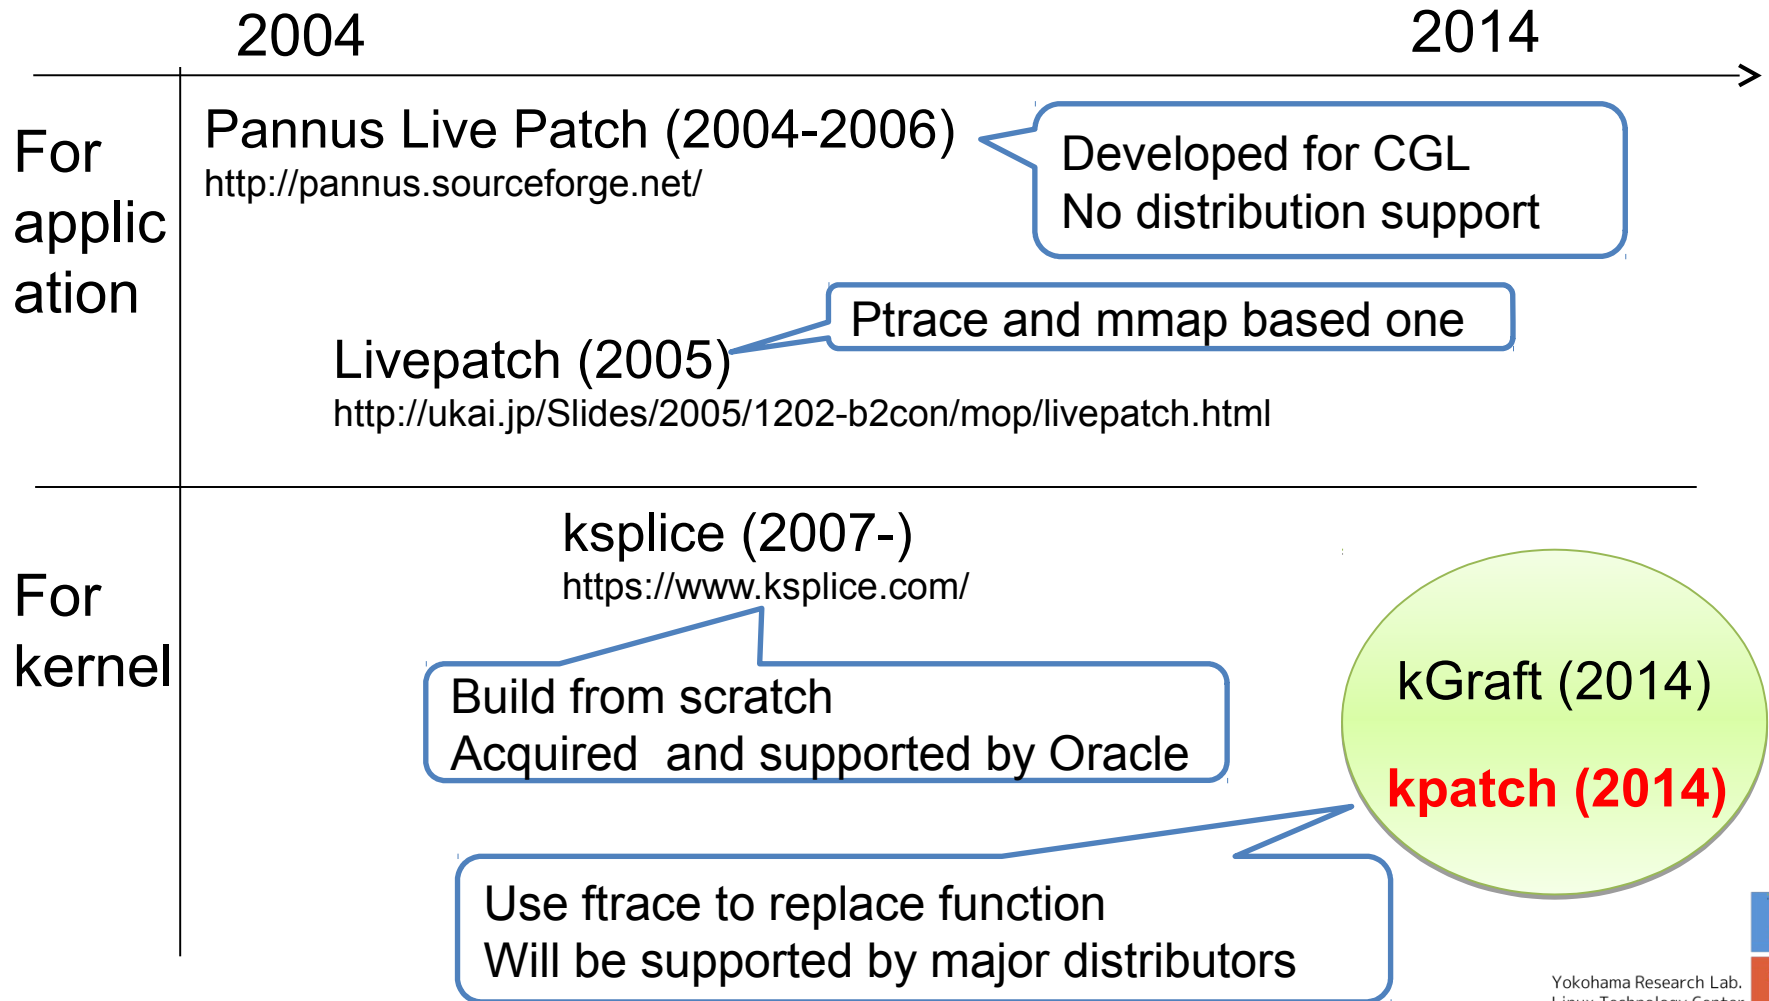
- Only for a small and critical issues
 - Not for major kernel update

- Live patching is important for appliances for mission critical systems
 - Some embedded appliances are hard to maintain frequently
 - Those are distributed widely in country side
 - Not in the big data center!
 - Some appliances can't accept 10ms downtime
 - Factory control system etc.

- M.C. systems have periodic maintenance
 - Major fixes can be applied and rebooted
 - In between the maintenance, live patching will be used
- Live patching and major update are complement each other
 - Live patching temporarily fixes small critical incidents
 - Major update permanently fixes all bugs



- Live patching is not new



Background Story

What is kpatch?

Live patching requirements

Major updates and Minor updates

Kpatch internal

Kpatch Overview

Active Safeness check

Stop_machine

Kpatch without stop_machine

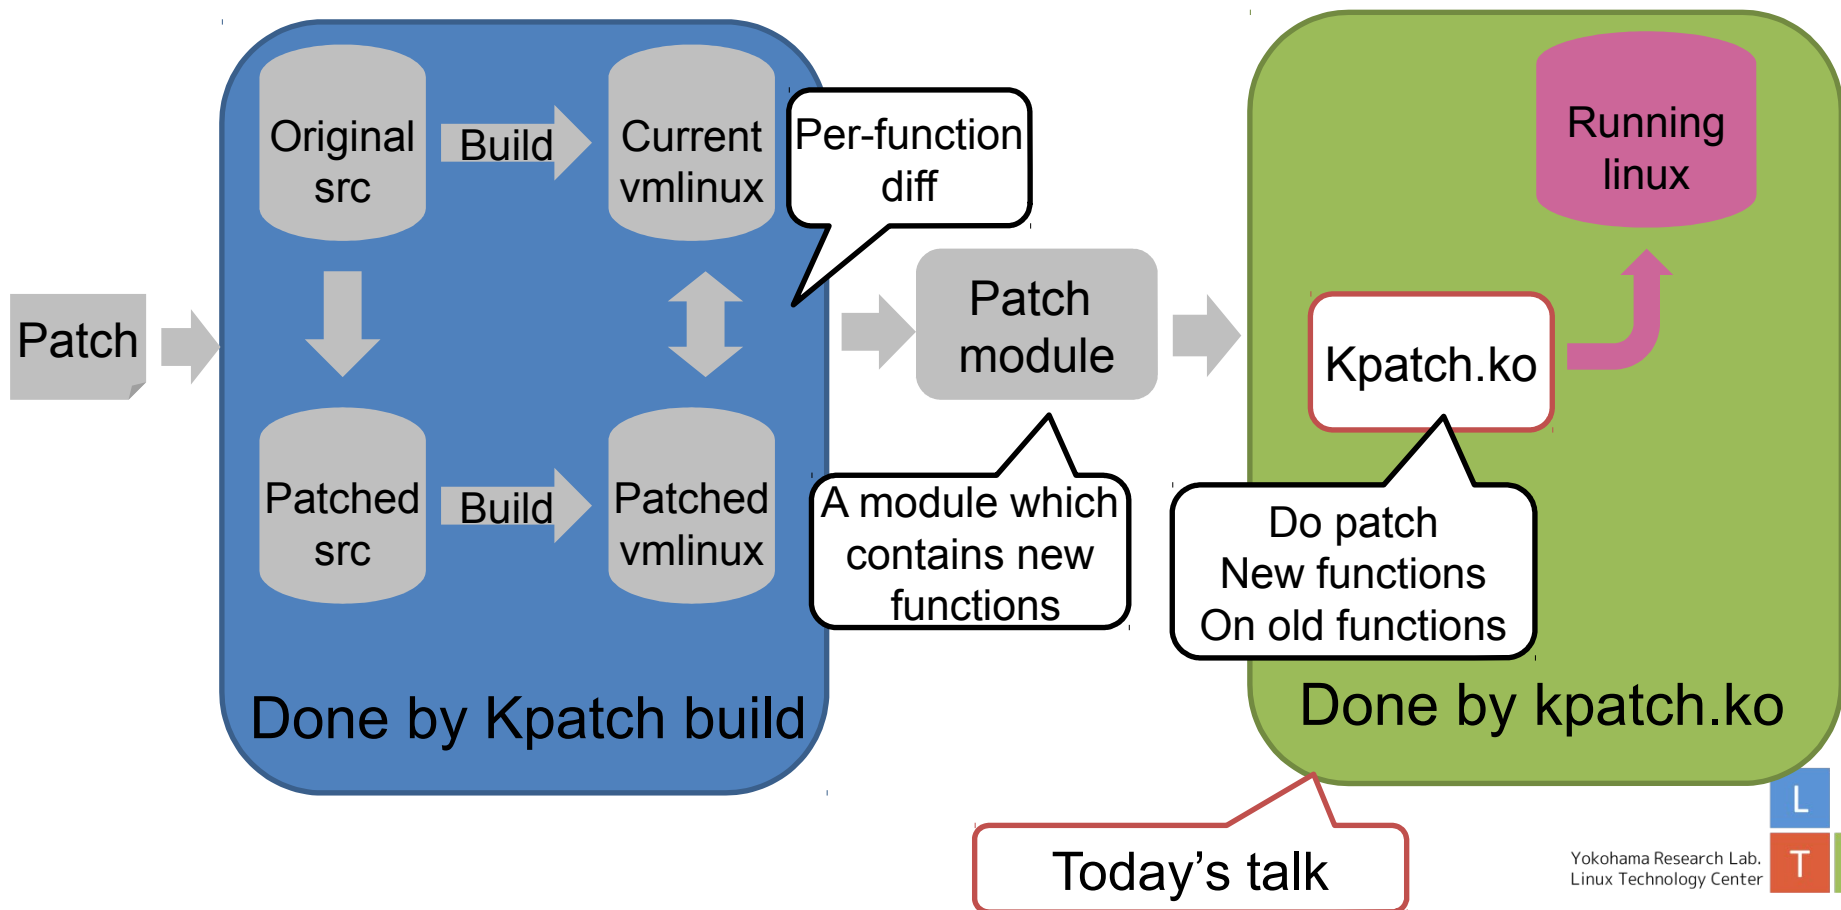
Live Patching Rules

Kpatch Reference Counter

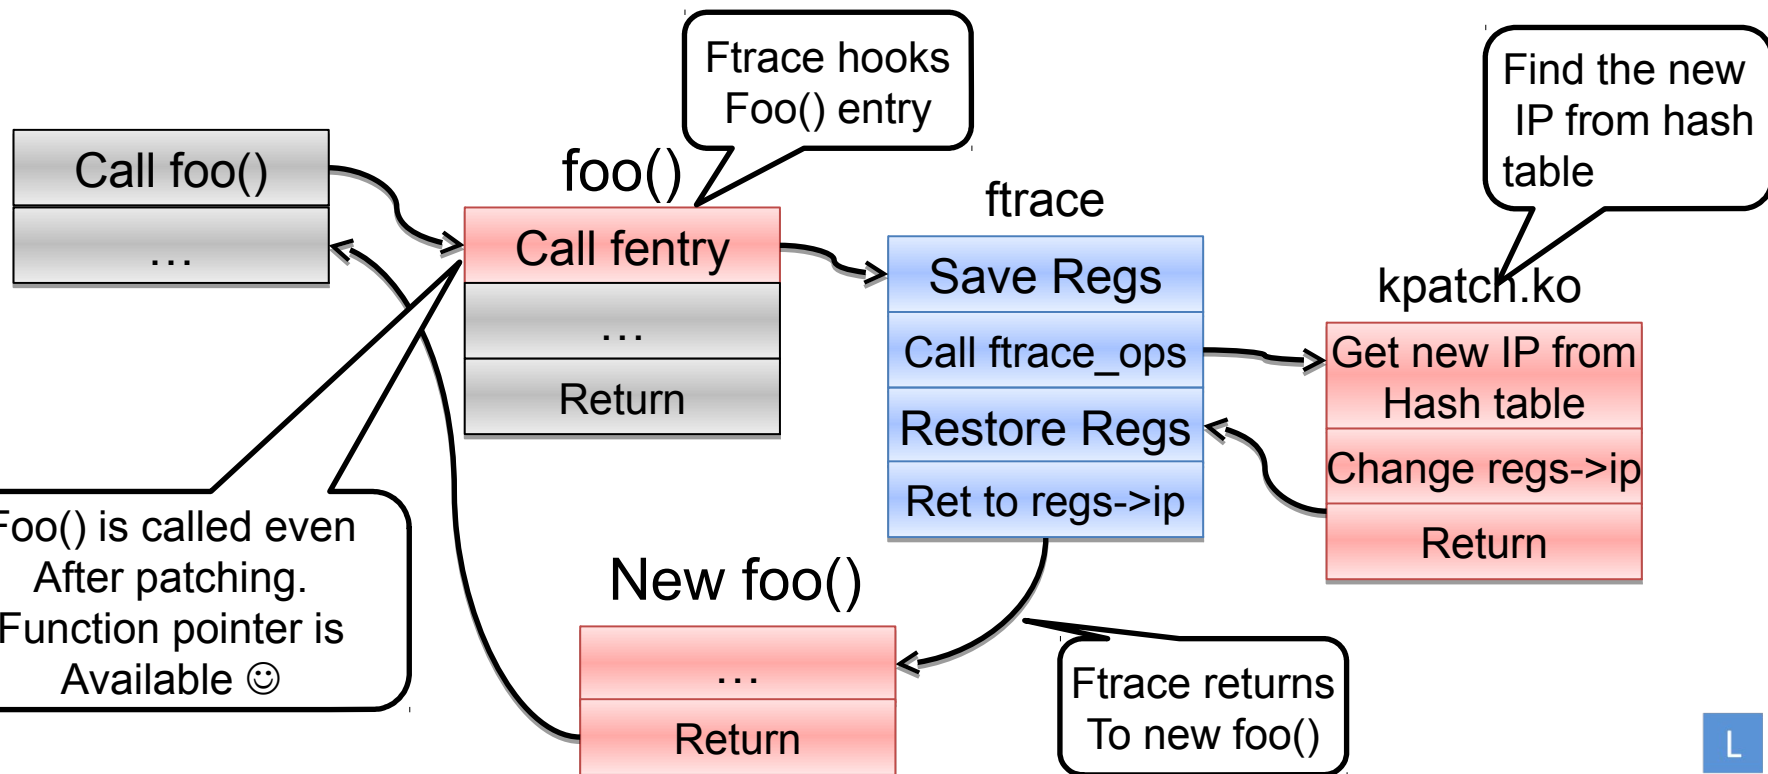
Safeness check without stop_machine

Conclusion and Discussion

- Kpatch has 2 components
 - Kpatch build: Build a binary patch module
 - Kpatch.ko: The kernel module of Kpatch



- Kpatch uses Ftrace to patch
 - Hook the target function entry with registers
 - Change regs->ip to new function (change the flow)




- Kpatch will update the execution path of a function
 - Q: What happen if the patched function is under executed?
 - A: Old and new functions are executed at the same time

!!This should not happen!!

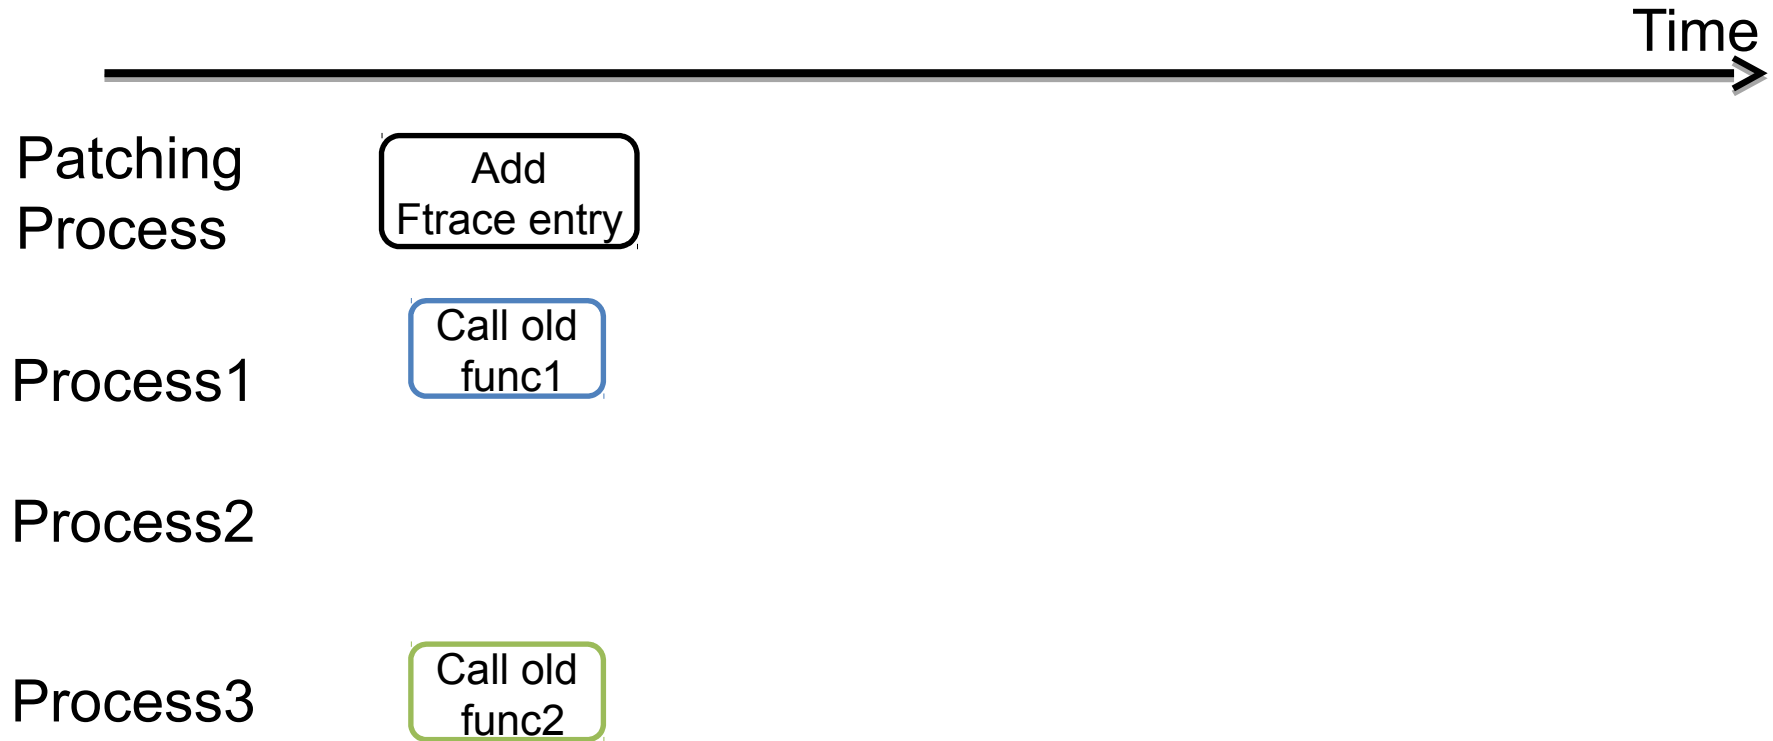
- Kpatch ensures the old functions are not executed when patching
 - “Active Safeness Check”

- Executing functions are on the stack
 - And IP register points current function too

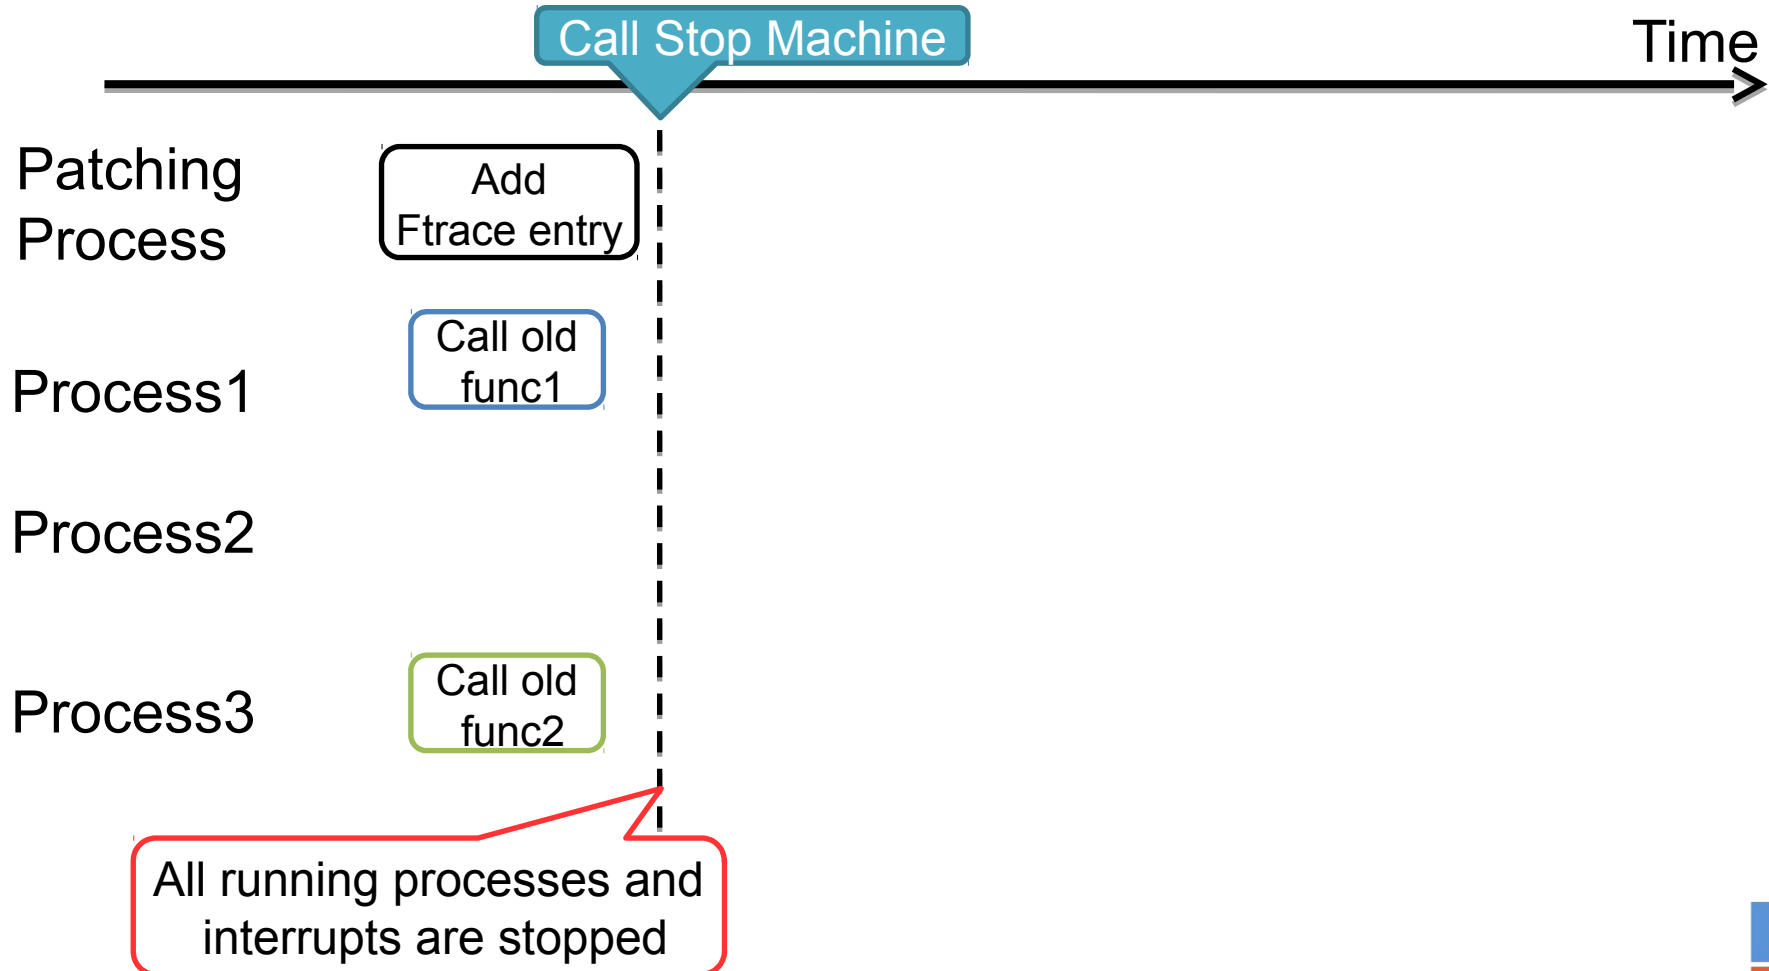


- Active Safeness Check
 - Do stack dump to check the target functions are not executed, for each thread.
 - Need to be done when the process is stopped.
 - stop_machine is used

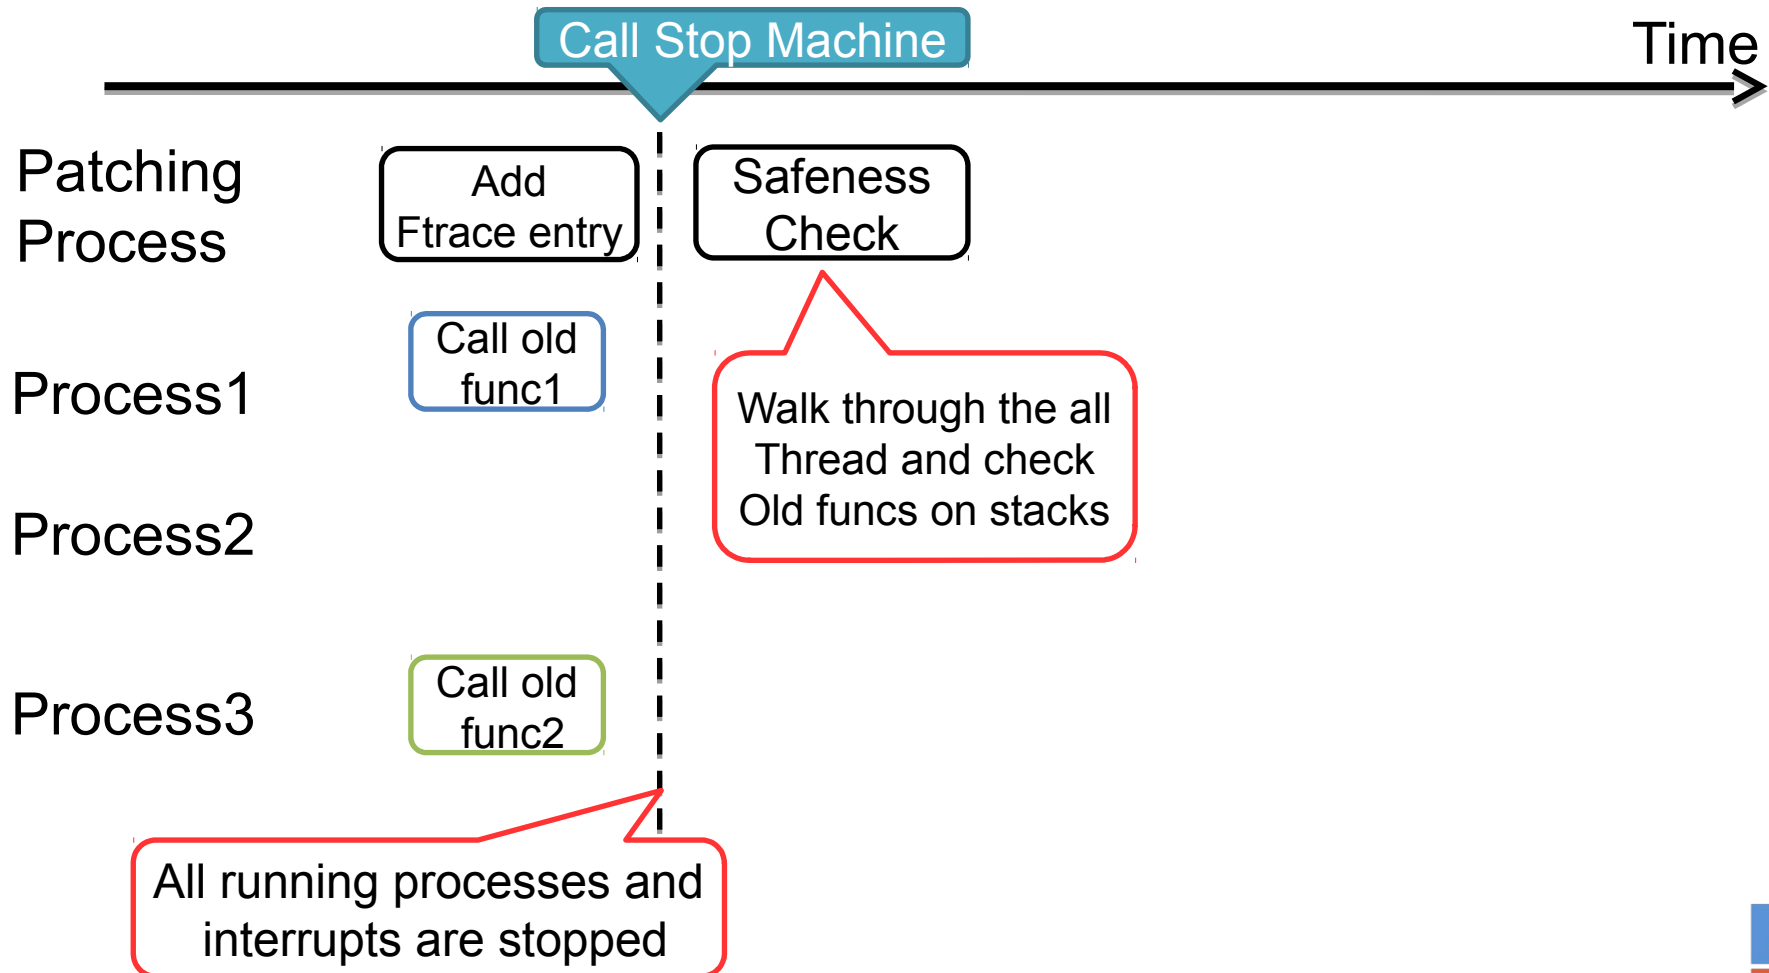
- Kpatch uses stop_machine to check stacks



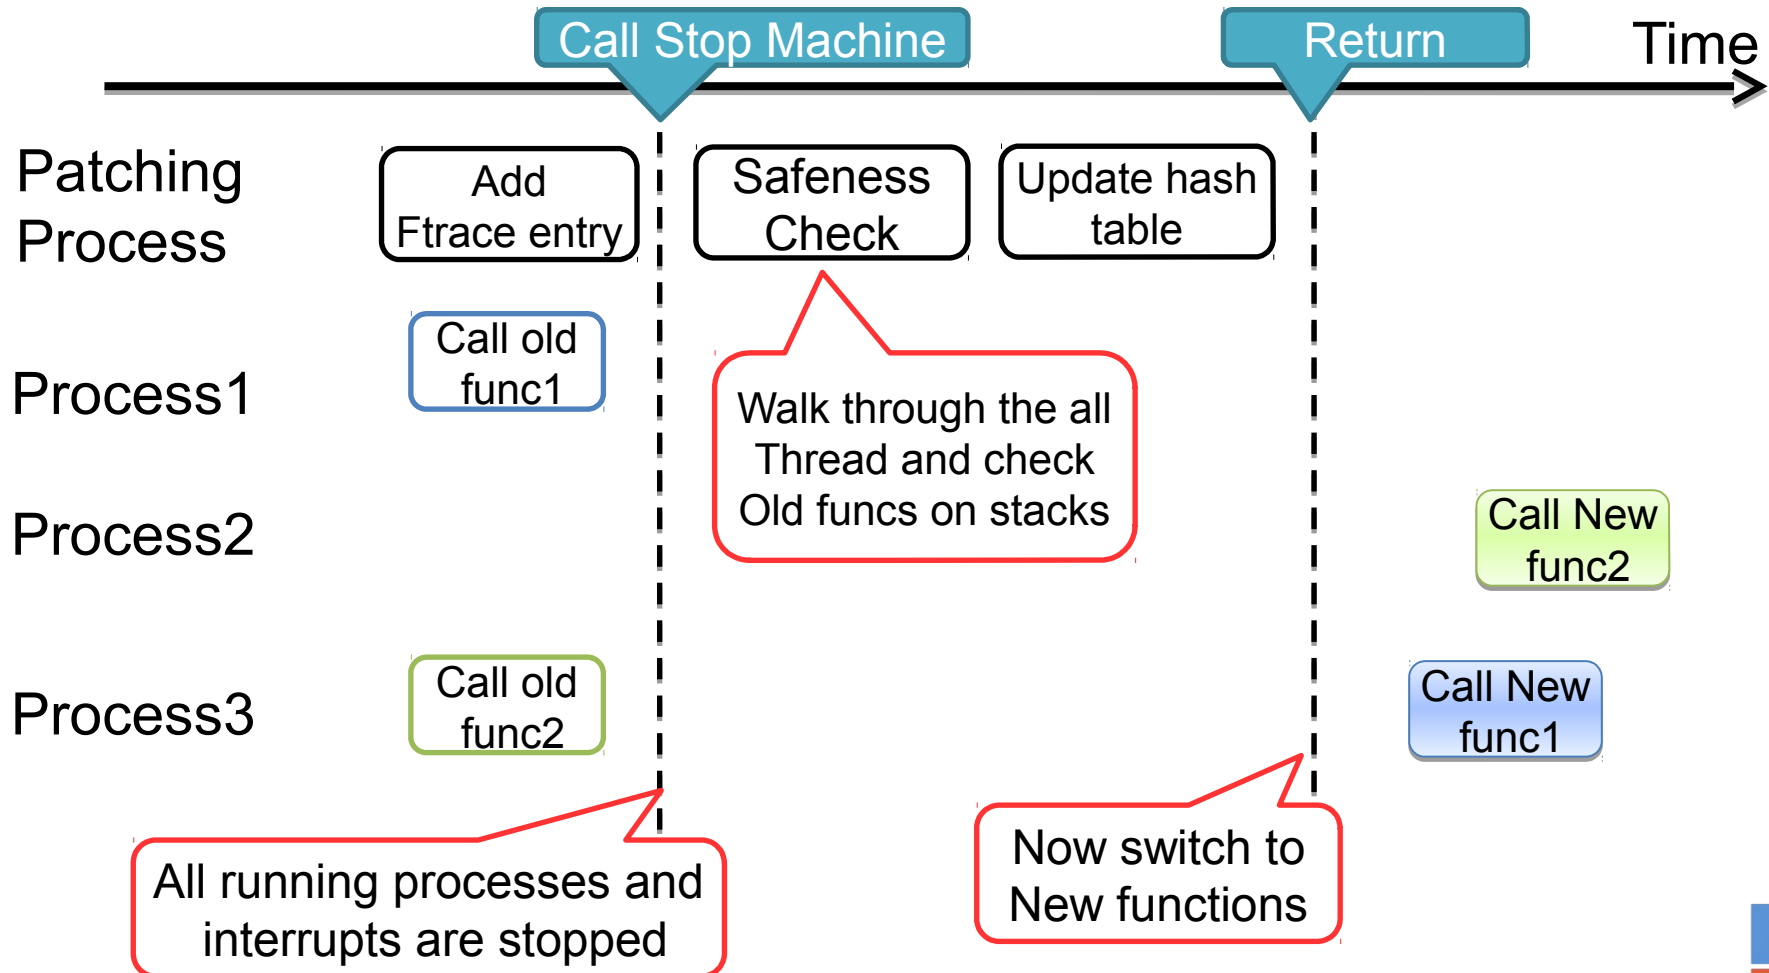
- Kpatch uses stop_machine to check stacks



- Kpatch uses stop_machine to check stacks



- Kpatch uses stop_machine to check stacks



- Pros
 - Safe, simple and easy to review, Good for the 1st version
- Cons
 - Stop_machine stops all processes a while
 - It is critical for control/network appliances
 - In virtual environment, this takes longer time
 - We need to wait all VCPUs are scheduled on the host machine

Background Story

What is kpatch?

Live patching requirements

Major updates and Minor updates

Kpatch internal

Kpatch and Ftrace

Kpatch and Safeness check

Stop_machine

Kpatch without stop_machine

Live Patching Rules

Kpatch Reference Counter

Safeness check without stop_machine

Conclusion and Discussion




- Live patching must follow the rules
 1. All the new functions in a patch must be applied at once
 - We need an atomic operation
 2. After switching new function, the old function must not be executed
 - We have to ensure no threads runs on old functions
 - And no threads sleeps on them

1. Introduce an atomic reference counter
2. Active safeness check at the context switch

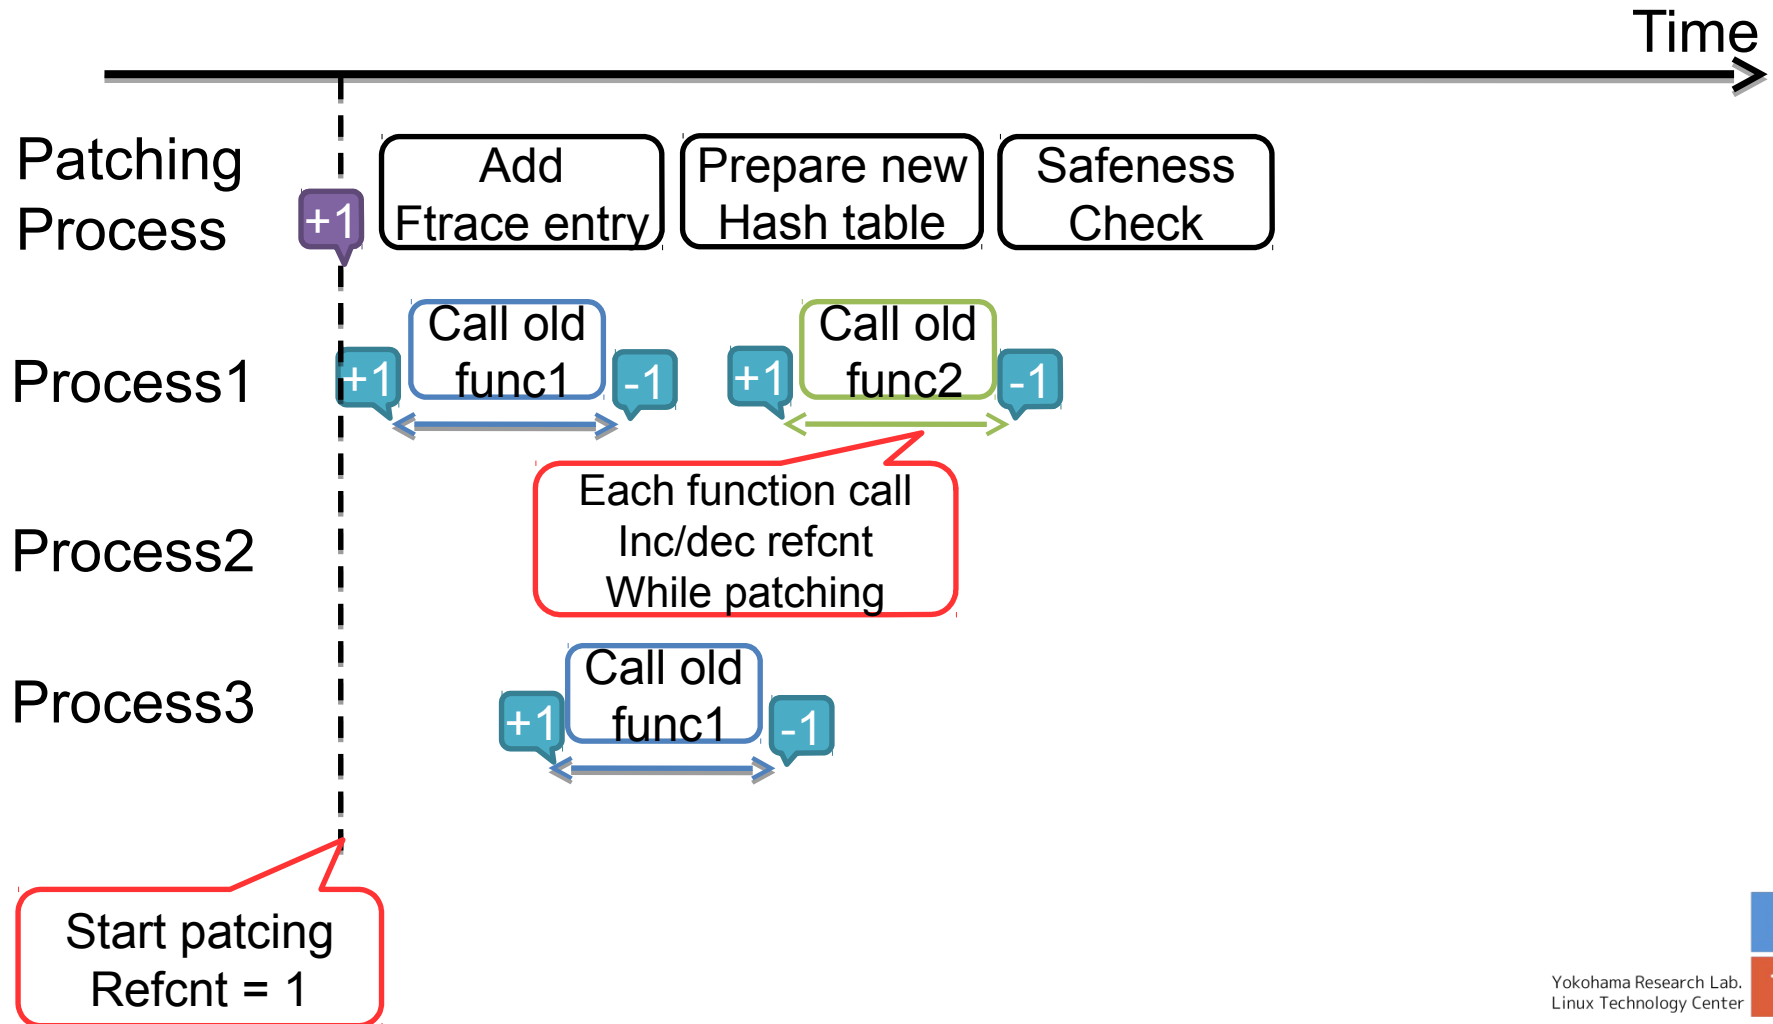
1. Introduce an atomic reference counter

- Without `stop_machine`, functions can be called while patching
 - Ensure no one actually runs functions -> refcounter
 - Increment the refcounter at entry
 - Decrement the refcounter at exit
- If refcounter is 0, update ALL function paths
 - We are sure there is no users

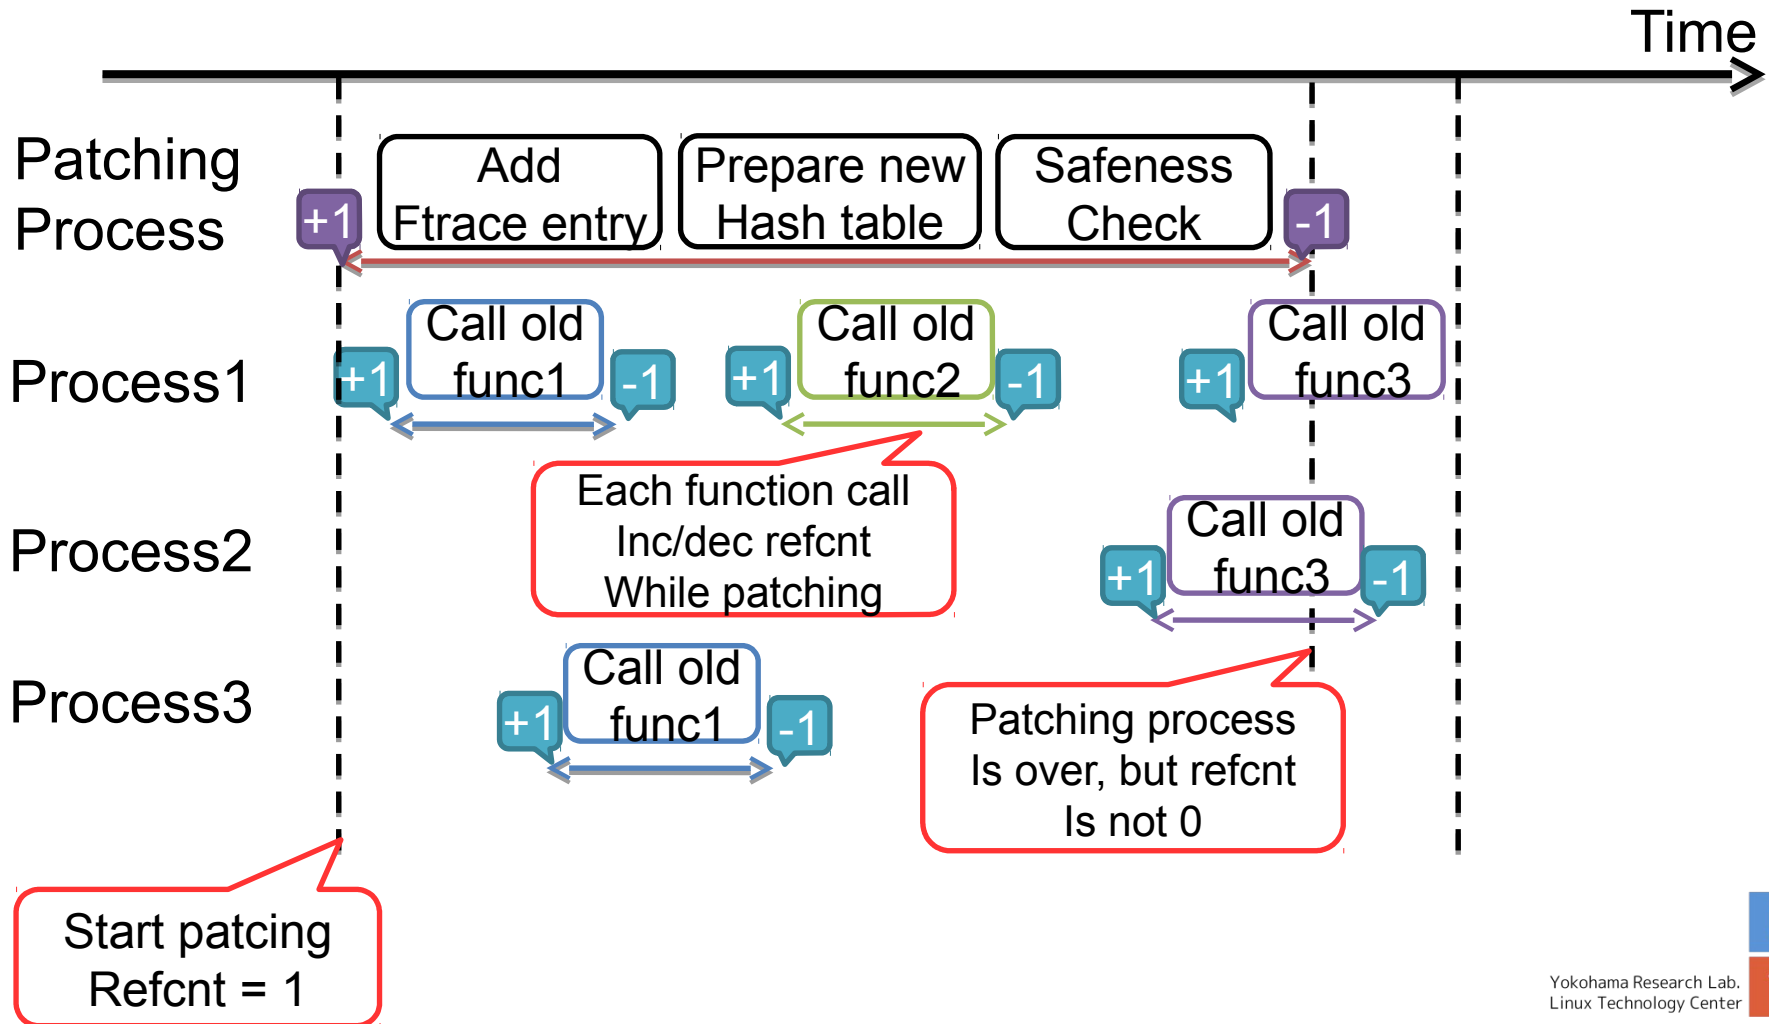
- Patching (switching) controlled by refcount



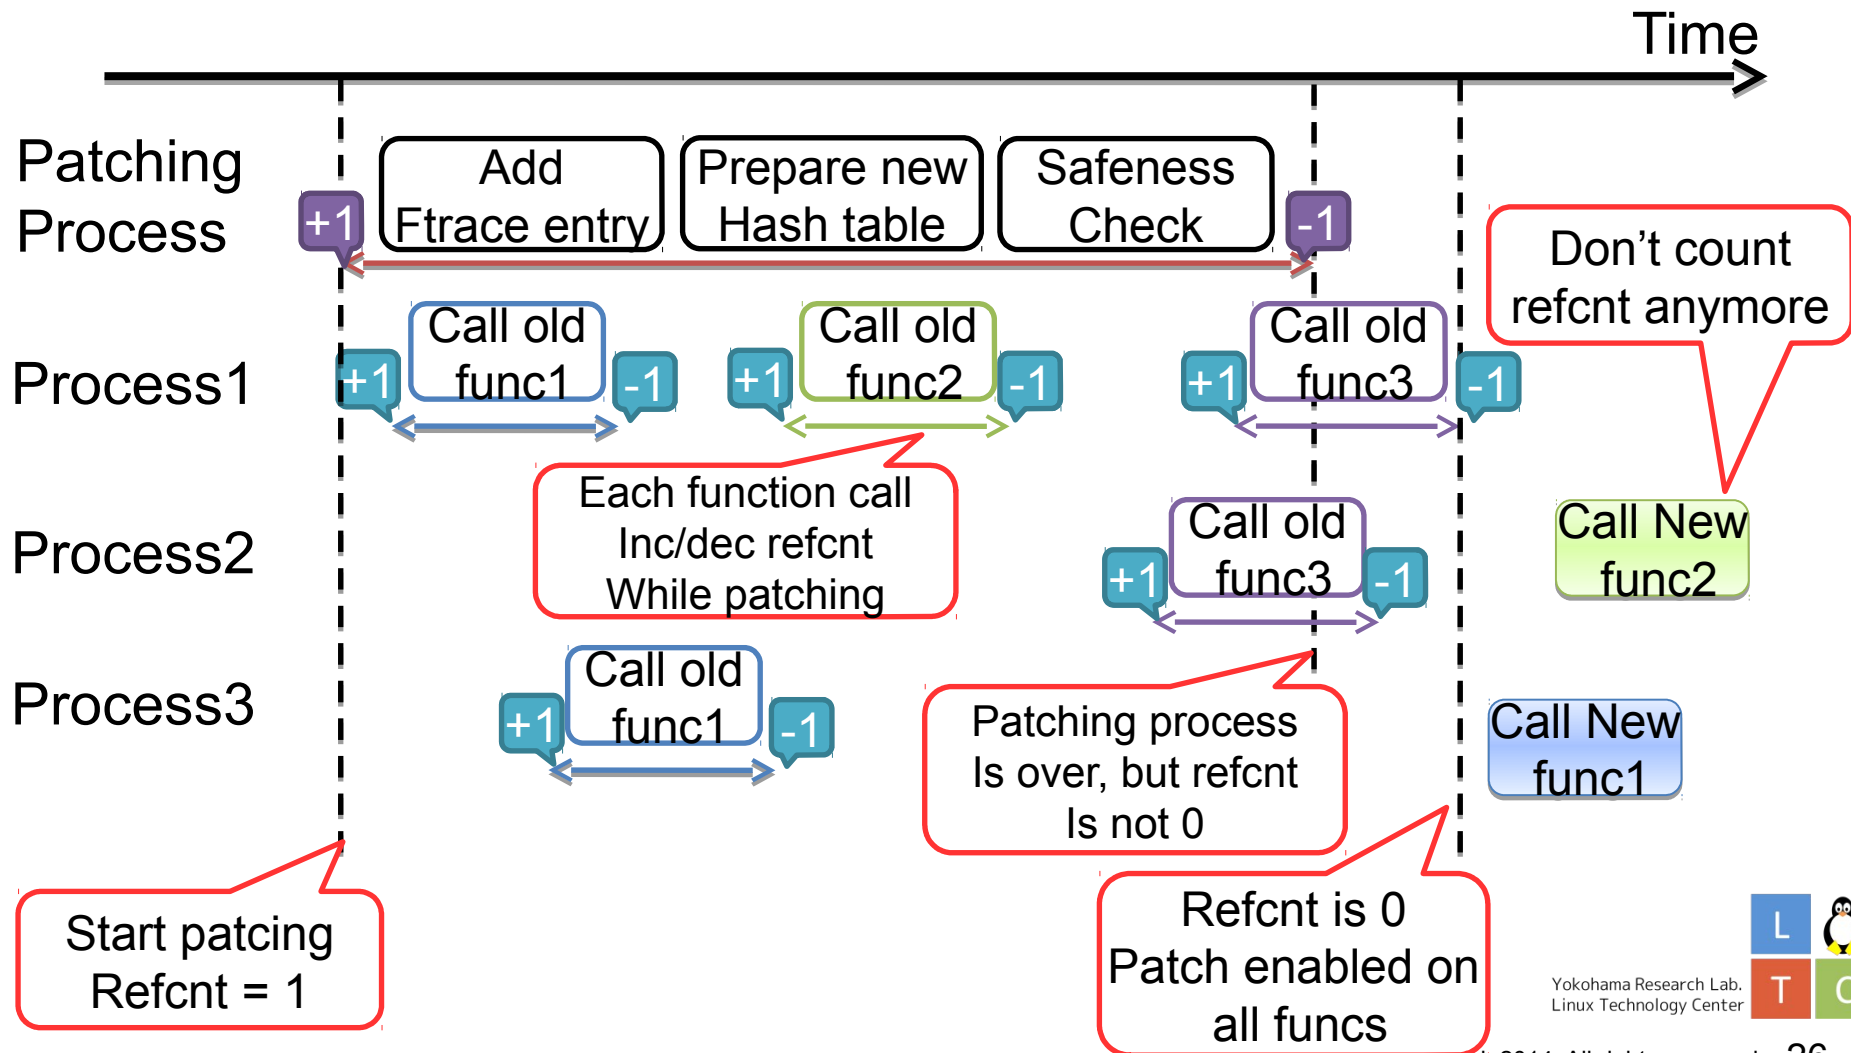
- Patching (switching) controlled by refcount



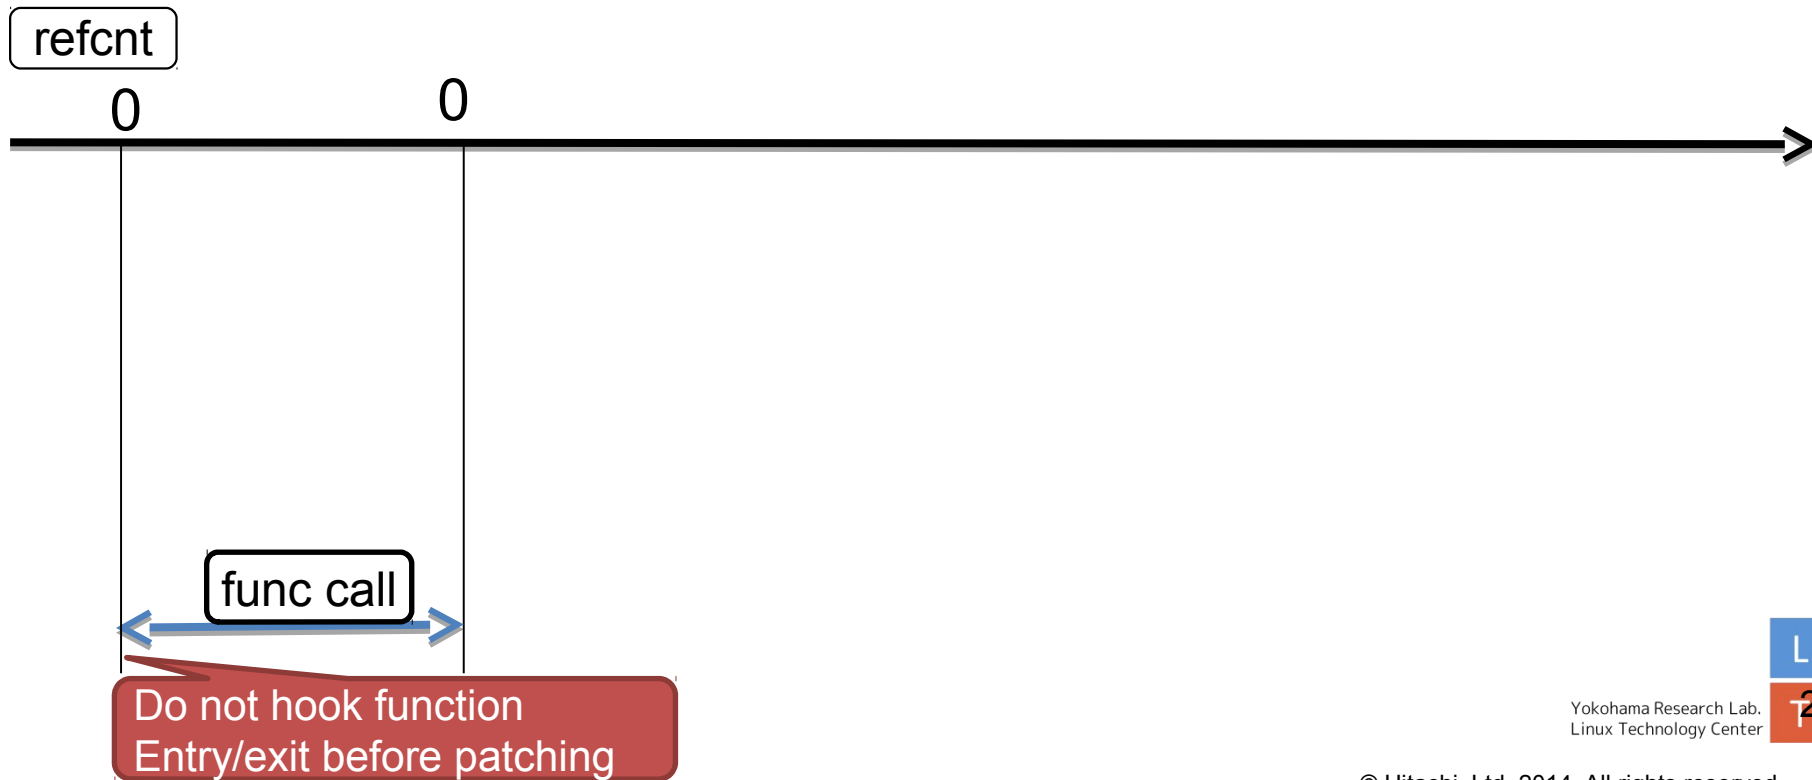
- Patching (switching) controlled by refcount



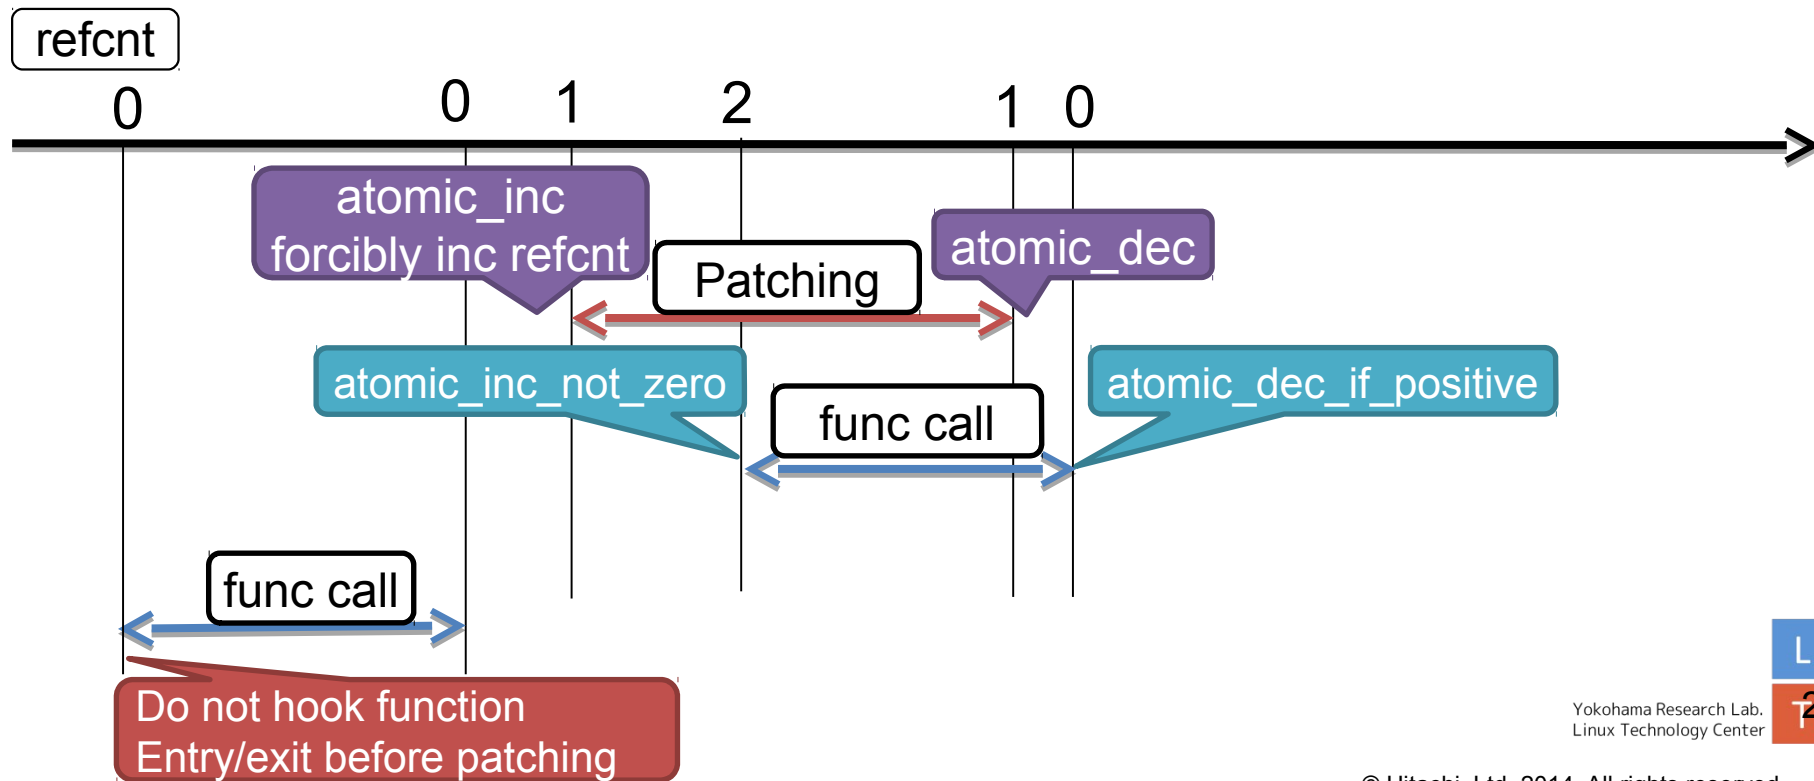
- Patching (switching) controlled by refcount



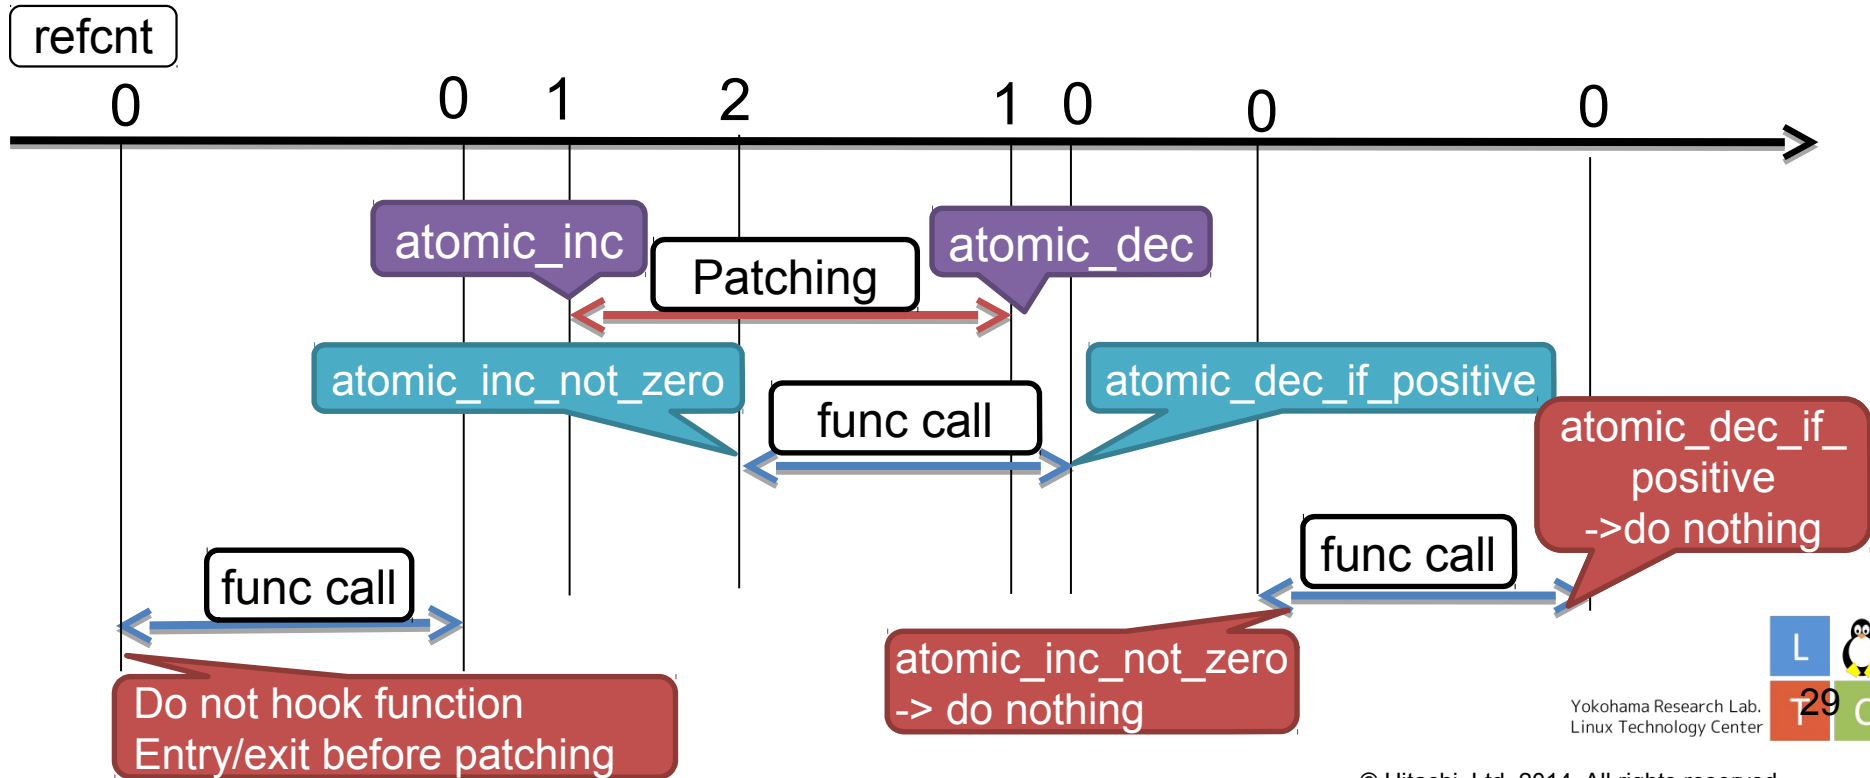
- Control the reference counter
 - Need to stop counting before and after patching
 - Use `atomic_inc_not_zero/dec_if_positive`
 - These are stopped automatically if counter == 0



- Control the reference counter
 - Need to stop counting before and after patching
 - Use `atomic_inc_not_zero/dec_if_positive`
 - These are stopped automatically if counter == 0

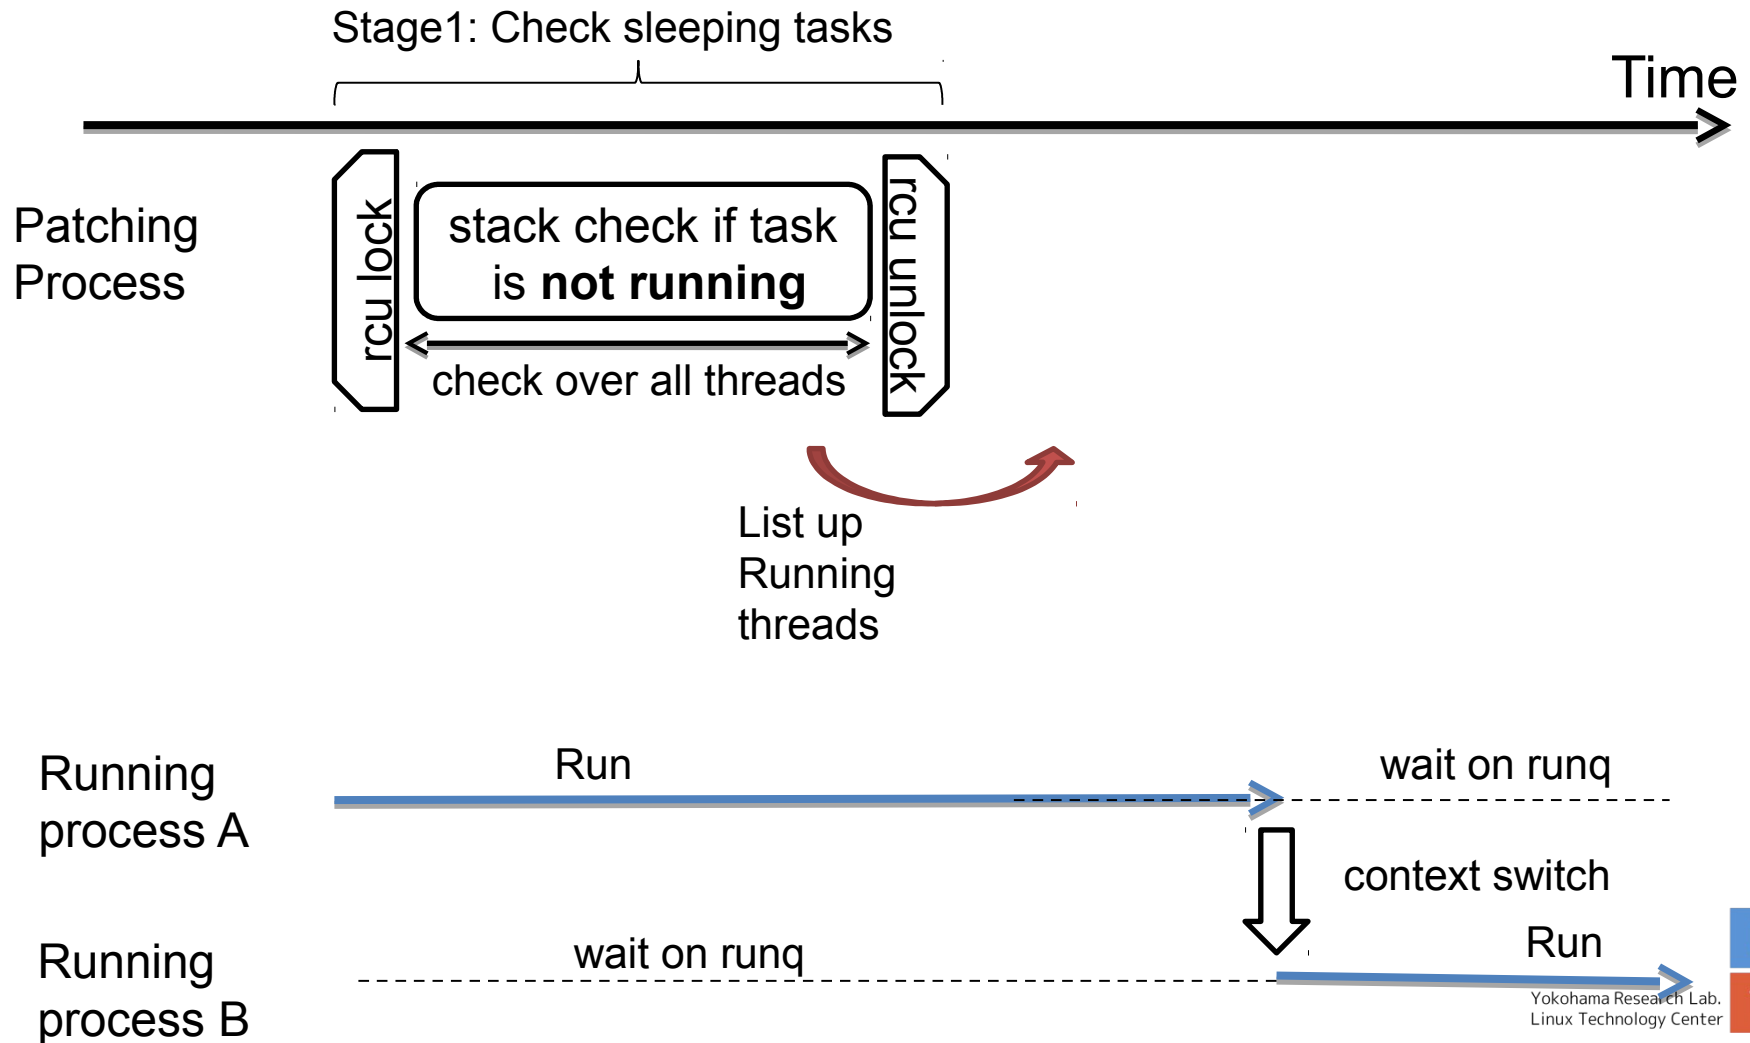


- Control the reference counter
 - Need to stop counting before and after patching
 - Use `atomic_inc_not_zero/dec_if_positive`
 - These are stopped automatically if counter == 0

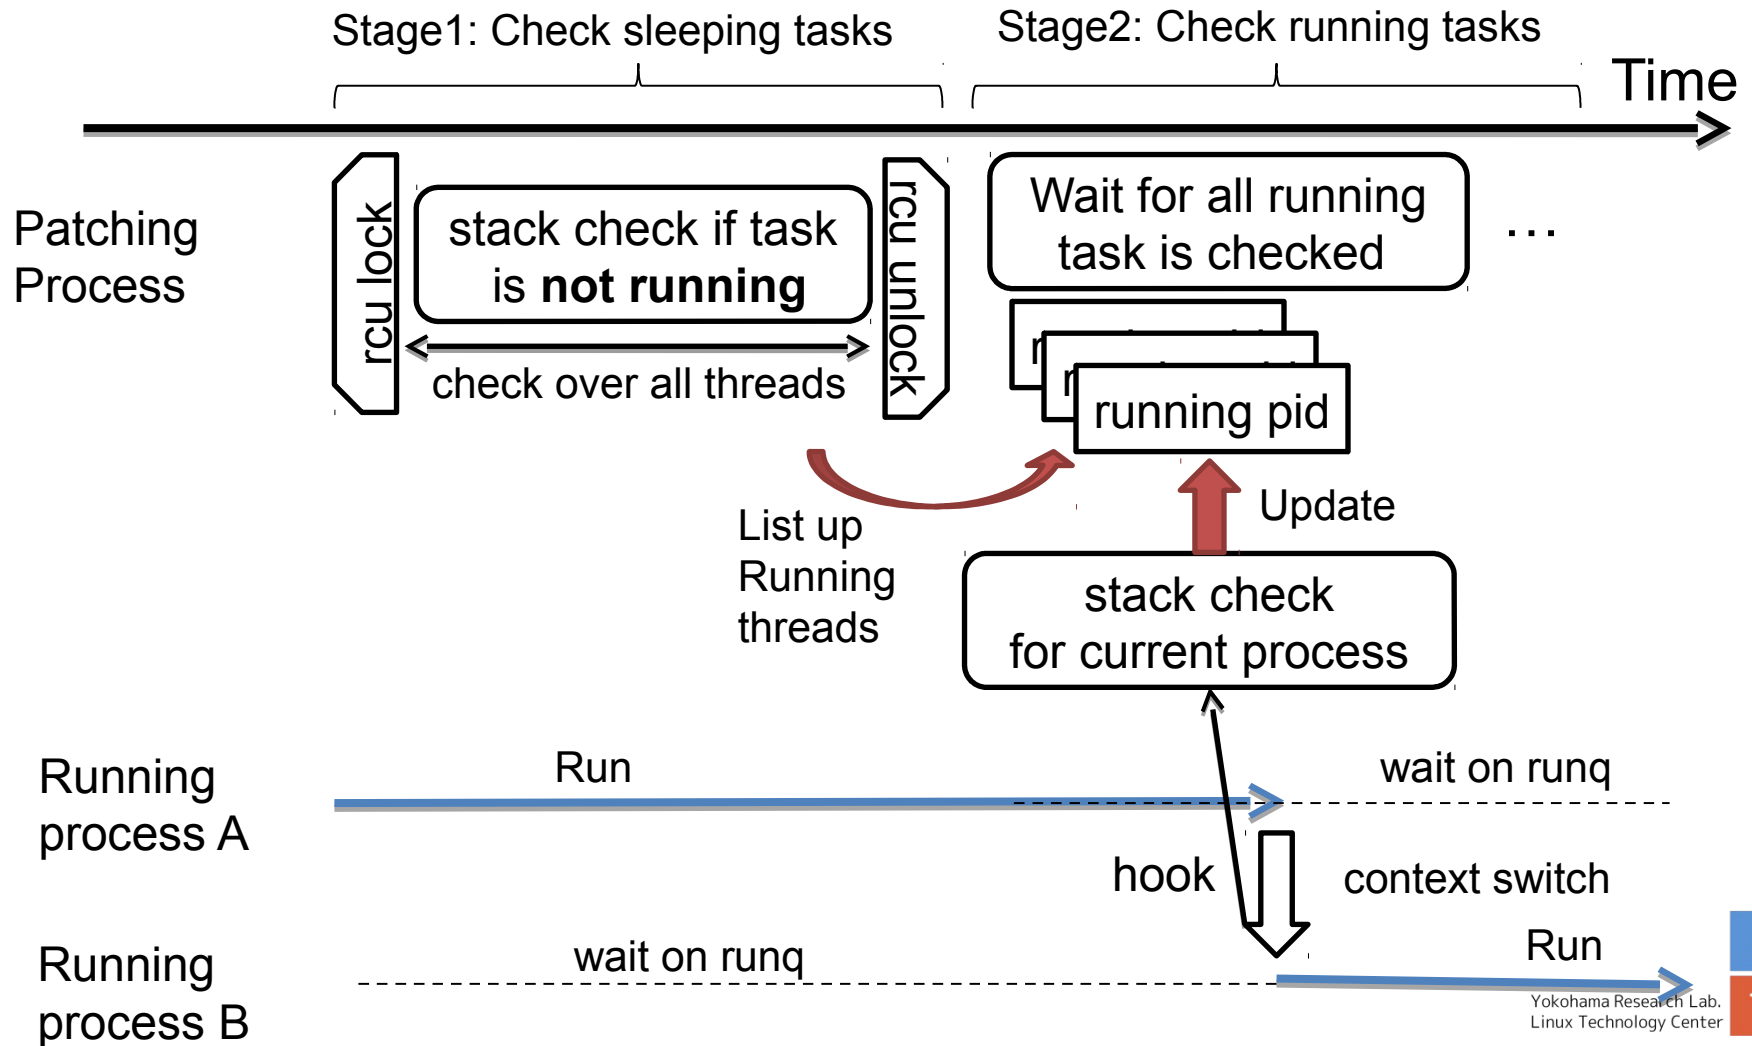


- ## 2. Active safeness check at the context switch
- To find threads sleeping(or going to sleep) on the functions
 - For all running processes, hook the context switch and check stack entries safely.
 - For the sleeping tasks, we can check it safely.

- 2 stages safeness checking



- 2 stages safeness checking



- To hook the function entry/return
 - Use kretprobe to hook it
 - For each function entry/return, inc/dec refcount
- To hook the context switch
 - Use kprobe to hook it
 - Do safeness check (on stack) and update running pid list
- Both are dynamic probe
 - After checking the safeness, all kretprobes/kprobes are removed from the target functions
 - We have minimal overhead

- Demonstrating kpatching with/without stop_machine
 - Using ftrace to trace stop_machine()

(Setup ftrace)

```
# echo stop_machine > /sys/kernel/debug/tracing/set_ftrace_filter  
# echo function_graph > /sys/kernel/debug/tracing/current_tracer
```

(Run the kpatch)

```
# kpatch load kpatch-test-patch.ko
```

(Check the result)

```
# echo 0 > /sys/kernel/debug/tracing/tracing_on  
# cat /sys/kernel/debug/tracing/trace
```

- With stop_machine

```
cat trace
# tracer: function_graph
#
# CPU DURATION          FUNCTION CALLS
# | | |                | | | |
0) ! 6410.455 us | stop_machine();
```

- Without stop_machine

```
cat trace
# tracer: function_graph
#
# CPU DURATION          FUNCTION CALLS
# | | |                | | | |
```

No stop_machine is executed



Background Story

What is kpatch?

Live patching requirements

Major updates and Minor updates

Kpatch internal

Kpatch Overview

Kpatch and Safeness check

Stop_machine

Kpatch without stop_machine

Live Patching Rules

Kpatch Reference Counter

Safeness check without stop_machine

Conclusion and Discussion

- Succeed to get rid of stop_machine() from kpatch
 - This is a proof of concept of stop_machine-free kpatch
 - This means kpatch CAN BE ready for mission critical systems
 - But still under discussion stage
- Upstreaming could be a long way
 - At first, push current stop_machine-based kpatch to upstream
 - Stop_machine-free will be the next step

- Possible to miscount the function reference
 - Kretprobe has no error notification
 - Kretprobe can be failed to handle the function return because of return-address buffer shortage
- Possible to fail patching with big patch
 - We have to monitor all the functions are safe in the patch
 - Big patch has many patched functions
 - Some of them can be always used in the system
 - Incremental patching could be better
- Module unloading using stop_machine
 - This will happen if we replace old patch with new one
 - Incremental patching can avoid this.

- Use a generic return-hook mechanism
 - **Kretprobe** is for PoC, not for general use
 - It can't detect the failure of hooking
 - Should be more safe (e.g. miss-hook handler)
- Context switch hook can be more general
 - Tracepoint/traceevent makes it better.
 - This requires kpatch as embedded feature
- Upstreaming

- Get rid of the stop_machine from kpatch
 - <https://github.com/dynup/kpatch/issues/138>
- My no-stopmachine branch
 - <https://github.com/mhiramathitachi/kpatch/tree/no-stop-machine-v1>
 - This requires IPMODIFY flag patchset for kernel
 - <http://thread.gmane.org/gmane.linux.kernel/1757201>

Questions?

HITACHI

Inspire the Next

Thank you!

Yokohama Research Lab.
Linux Technology Center



- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Other company, product, or service names may be trademarks or service marks of others.