

# Linux Plumbers Conference 2022

>> Dublin, Ireland / September 12-14, 2022



## fw\_devlink & parallelization

Saravana Kannan

# Linux Plumbers Conference 2022

>> Dublin, Ireland / September 12-14, 2022



Upstream status update



# fw\_devlink for parallelization

fw\_devlink started off as a way to avoid unnecessary probe attempts and add `sync_state()`

Trying to expand it to support parallelized boot and suspend/resume.

This means fw\_devlink needs to be as complete as possible about capturing dependencies.



# Landed since LPC 2021

- Add flag for parent devices that need child to probe on addition.
- Improved logging and /debug/devices\_deferred.
- Remove some unnecessary device link creation.
- Fixed some issues with simple-pm-bus, simple-bus & similar devices.
- Added wildcard "\*" support to driver\_async\_probe= (Please give it a shot)
- Extend driver\_deferred\_probe\_timeout= on new driver registration.
- Fix driver\_deferred\_probe\_timeout conflict with IP auto config
- driver\_deferred\_probe\_timeout=10 with CONFIG\_MODULES
- fw\_devlink.strict=1 by default
- Added module.async\_probe to set default async probe for modules.
- Removed deferred addition of AMBA devices.

# Linux Plumbers Conference 2022

>> Dublin, Ireland / September 12-14, 2022




Work in progress/TODOs



# Cyclic dependencies

- A, B, B', B'' and C should not have probe dependencies enforced between them.
- X should still be deferred until A probes.
- Y should still be deferred until C probes.
- Maybe create suspend/resume ordering device links as devices in the cycle probe. Eg: If A probes first, then mark B as consumer of A?



Pointy end of arrow is at supplier/parent  
Red arrows indicate parent-child relationship  
These cycles are from real world scenarios



# Improvements in progress

- No longer requires “compatible” property for identifying what DT node will become a device. This is done lazily as devices are created.
- Support corner cases where DT child is supplier of DT parent, but the parent/child relationship is not maintained in struct device model.
- Trying to remove `driver_deferred_probe_check_state()` so “Not waiting for optional suppliers” doesn’t need to be handled in each framework.



# TODOs for base features

- Add `sync_state()` ops to struct class.
- Add MANAGED device link support for “class” devices.
- Add support for linking consumers to per-resource devices for better `sync_state()` granularity.
- Add `sync_state()` support in regulator and power domain frameworks.
- Address comments and upstream `sync_state()` for clock framework.
- Add support for `phy-handle` - forced binding of PHYs to generic driver complicates things.



# Linux Plumbers Conference 2022

>> Dublin, Ireland / September 12-14, 2022




## Parallelization




# Parallelize boot


Synchronous probing (~900ms)



Asynchronous probing (~450ms) - Please give `driver_async_probe=*` a shot and report issues. Want to enable it by default at some point if `fw_devlink=on`.



Parallel module loading (~250ms) - Supported by Android 13's `init` (`androidboot.load_modules_parallel=1`)





# Parallelize suspend/resume

- Drivers need to opt-in right now.
- Should I just add a new command line param to enable it by default for all devices?
- Keep track of device suspend/resume time and don't auto opt-in for devices which never take longer than a configurable threshold (say 250 us)?
- Add flag for drivers to opt-out?

# Linux Plumbers Conference 2022

>> Dublin, Ireland / September 12-14, 2022



Thank you!