

libcamera: Making Complex Cameras Easy

OSS Japan 2019
Tokyo, Japan

Laurent Pinchart
laurent.pinchart@ideasonboard.com


libcamera

Cameras are complex devices that need heavy hardware image processing operations. Control of the processing is based on advanced algorithms that must run on a programmable processor. This has traditionally been implemented in a dedicated MCU in the camera, but in embedded devices algorithms have been moved to the main CPU to save cost. Blurring the boundary between camera devices and Linux often left the user with no other option than a vendor-specific closed-source solution.

To address this problem the Linux media community has very recently started collaboration with the industry to develop a camera stack that will be open-source-friendly while still protecting vendor core IP. libcamera was born out of that collaboration and will offer modern camera support to Linux-based systems, including traditional Linux distributions, ChromeOS and Android.

Contents


- [libcamera](#)



*In the beginning were
simple pipelines...*

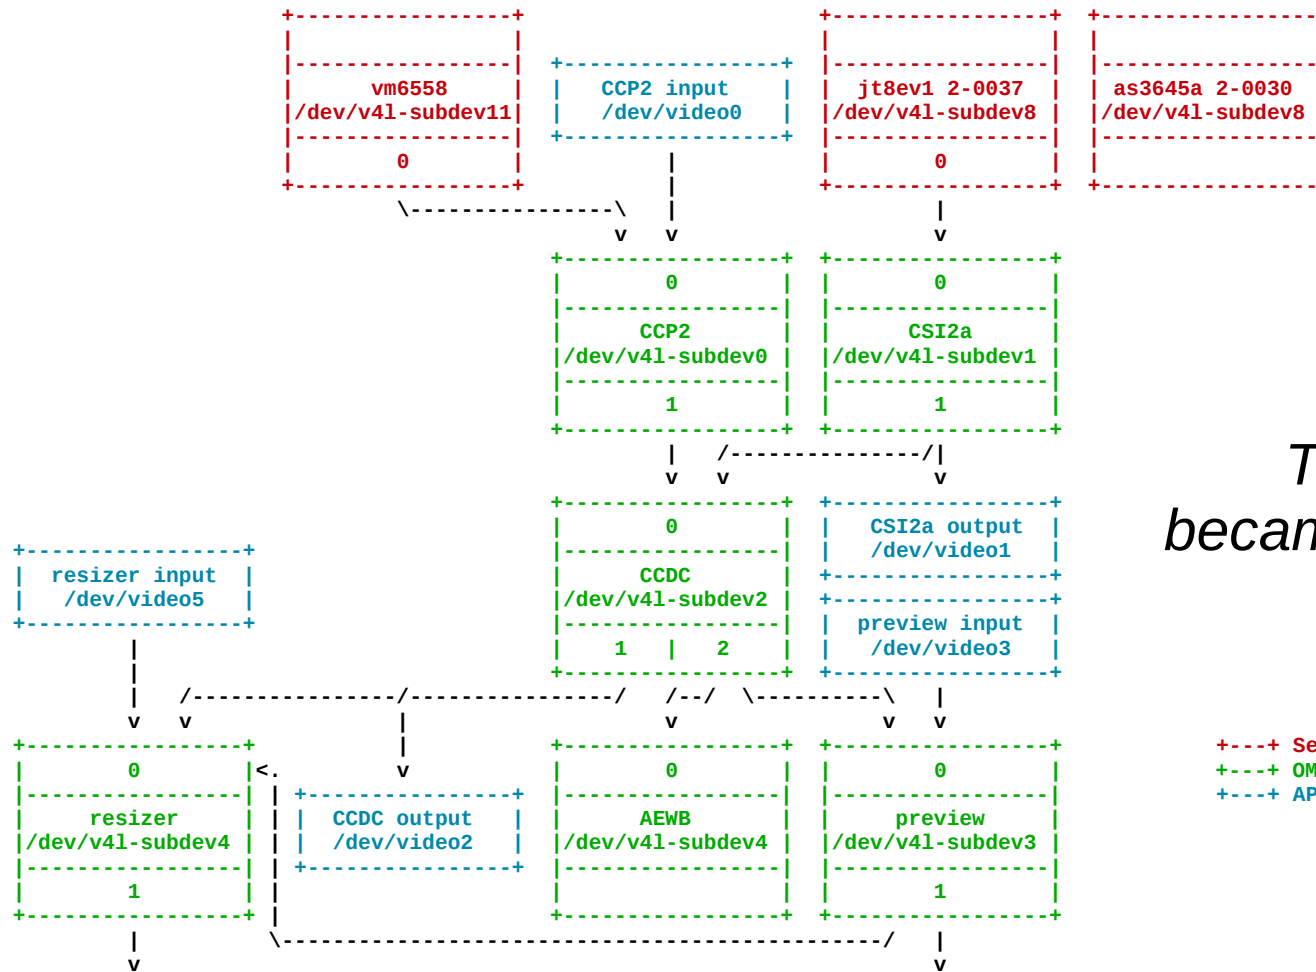
Why?

*... and they were
simple to control,
with a single API.*



Why?






Then the world became complex ...

+---+ Sensors & flash
 +---+ OMAP3 ISP
 +---+ API

Why?

IDEAS
ON BOARD


... and application developers were left suffering.



Why?

**IDEAS
ON BOARD**

*Solutions were
proposed...*



*... but never
implemented.*

Why?

IDEAS
ON BOARD

*Then hope came
back.*



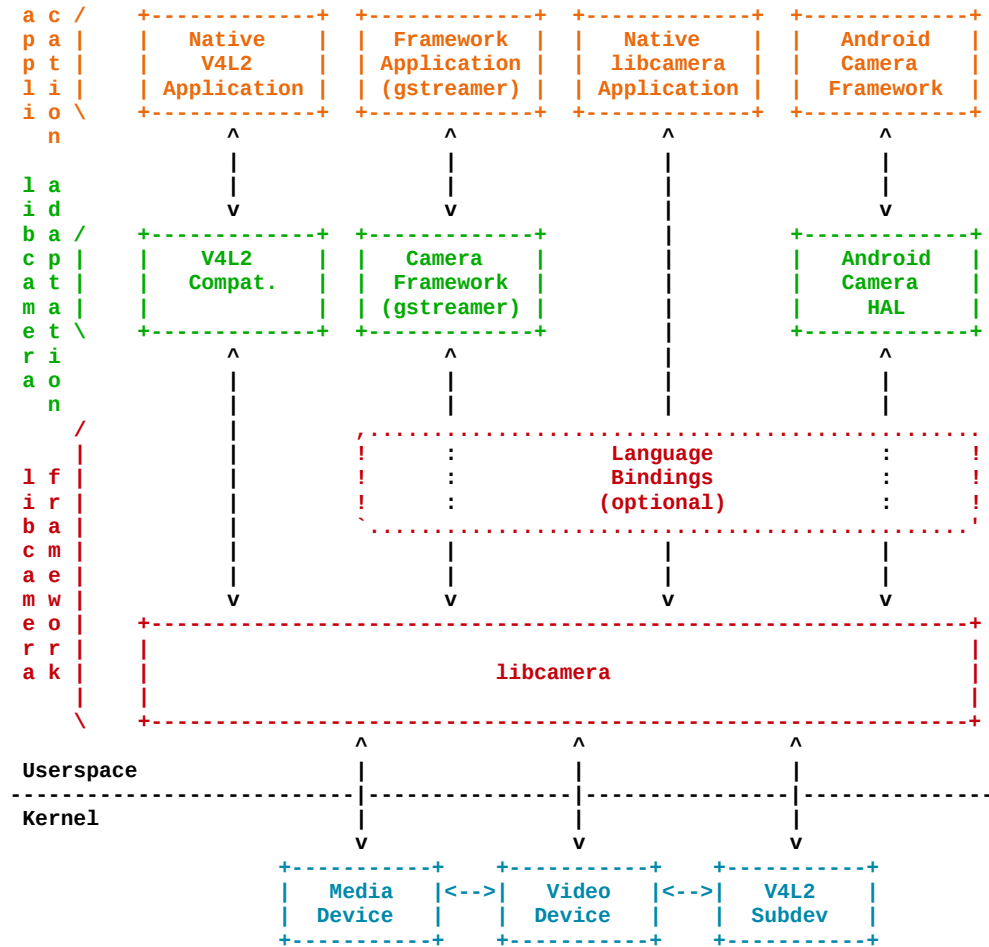
+ - / \ - +

| (o) |

+ - - - - +

libcamera

libcamera provides a complete userspace camera stack.

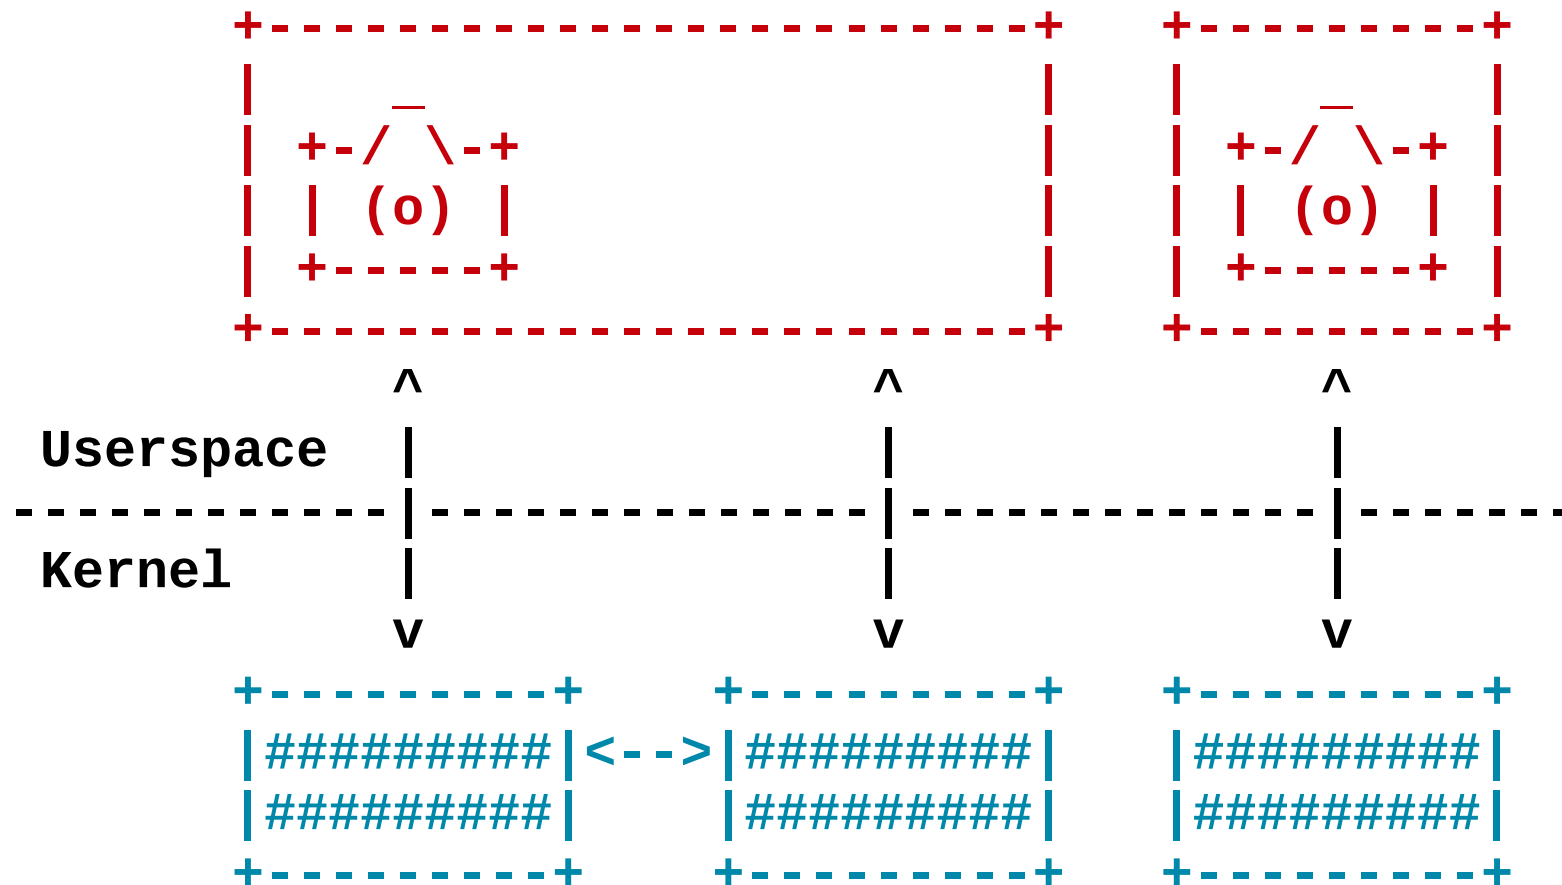


The 'mesa' of the camera world.

Camera Stack




*libcamera
enumerates
cameras...*




Camera Devices & Enumeration

*... and
exposes their
capabilities.*



Capabilities




Profiles

Capabilities & Profiles




*It supports multiple
concurrent streams
for the same
camera...*



Streams

... and per-frame controls.



Per-Frame Controls

*Image Processing
Algorithm are loaded
as external modules.*

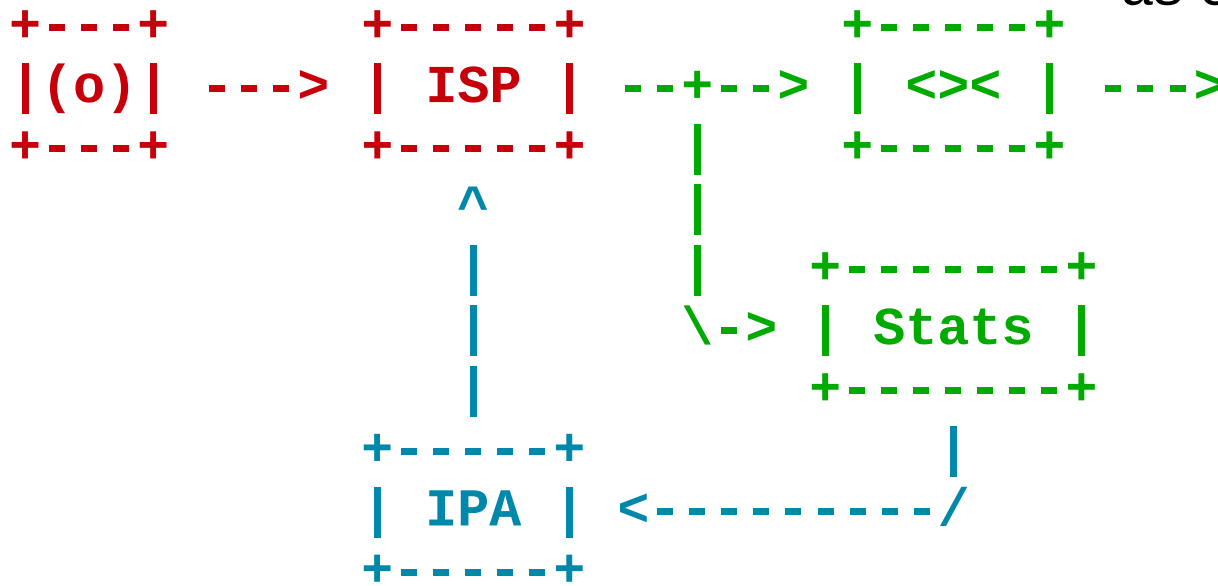


Image Processing Algorithms (3A)

+ - - - - +
| V4L2 App. |
+ - - - - +

+ - - - - +
| V4L2
API |
+ - - - - +

+ - - - - +
| libcamera |
+ - - - - +

*Adaptation layers
offer backward
compatibility with
existing APIs...*

Adaptation

IDEAS
ON BOARD

+-----+
| V4L2 App. |
+-----+

+-----+
| Android |
+-----+

*... and integrate
libcamera with
other operating
systems.*

+-----+
| V4L2
API |
+-----+

+-----+
| _____/ |
/ . . \
!
!

+-----+

+-----+
| libcamera |
+-----+



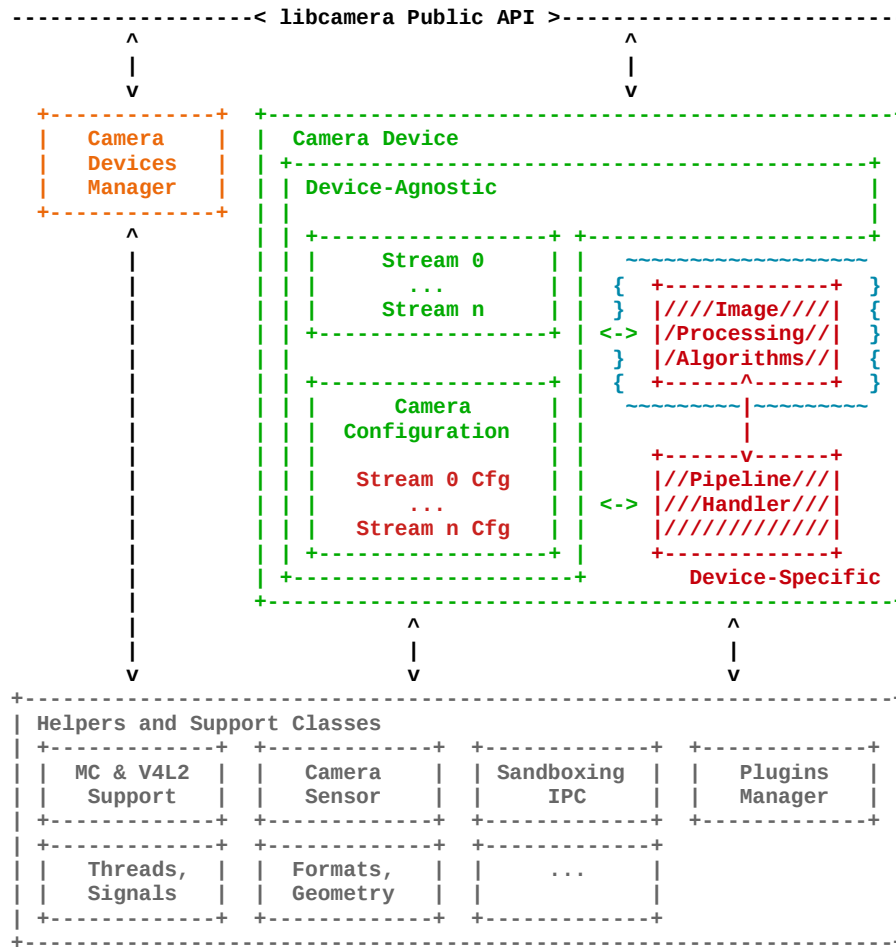
Adaptation

+ - / \ - +

| (o) |

+ - - - - +

libcamera




Central to the stack is the Camera object, interfacing to device-specific pipeline handlers.

/// Device-Specific Components
~~~ Sandboxing

# libcamera architecture






*The Camera Manager  
enumerates media  
devices and instantiates  
corresponding pipeline  
handlers.*




# Camera Devices Manager



*The pipeline handlers create and register one or more cameras.*



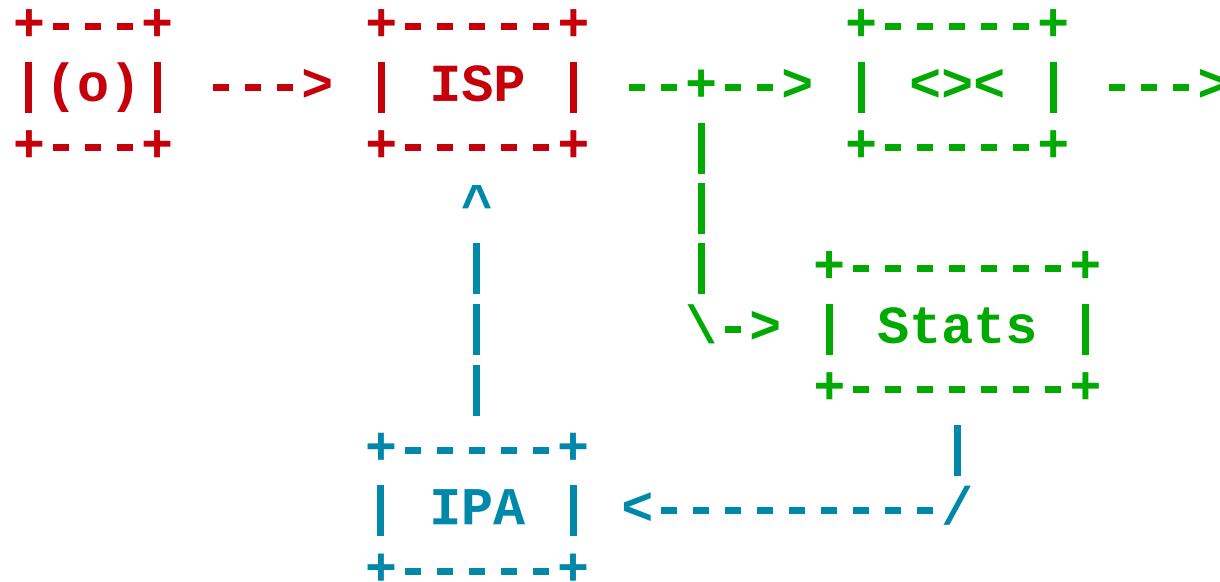
# Camera Devices Manager



*The pipeline handler interfaces with all kernel devices. It abstracts them and exposes video streams to upper layers.*

# Pipeline Handler






*Image Processing Algorithms (IPA) receive statistics from the hardware and compute optimal image parameters.*

# Image Processing Algorithms

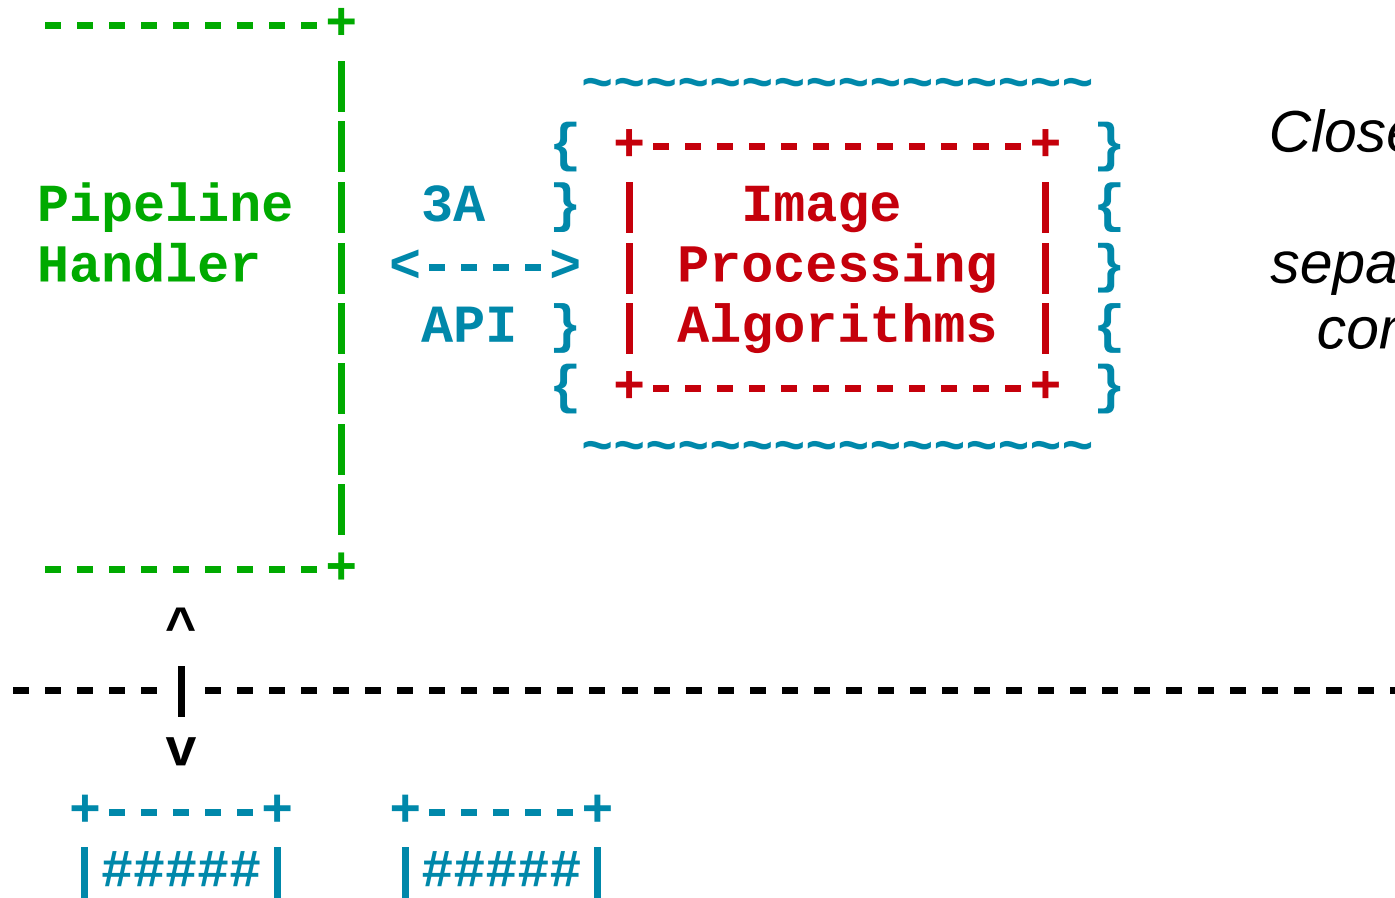




*IPAs are separate modules that don't access kernel devices directly. They only have access to their pipeline handler through the IPA API.*

# Image Processing Algorithms

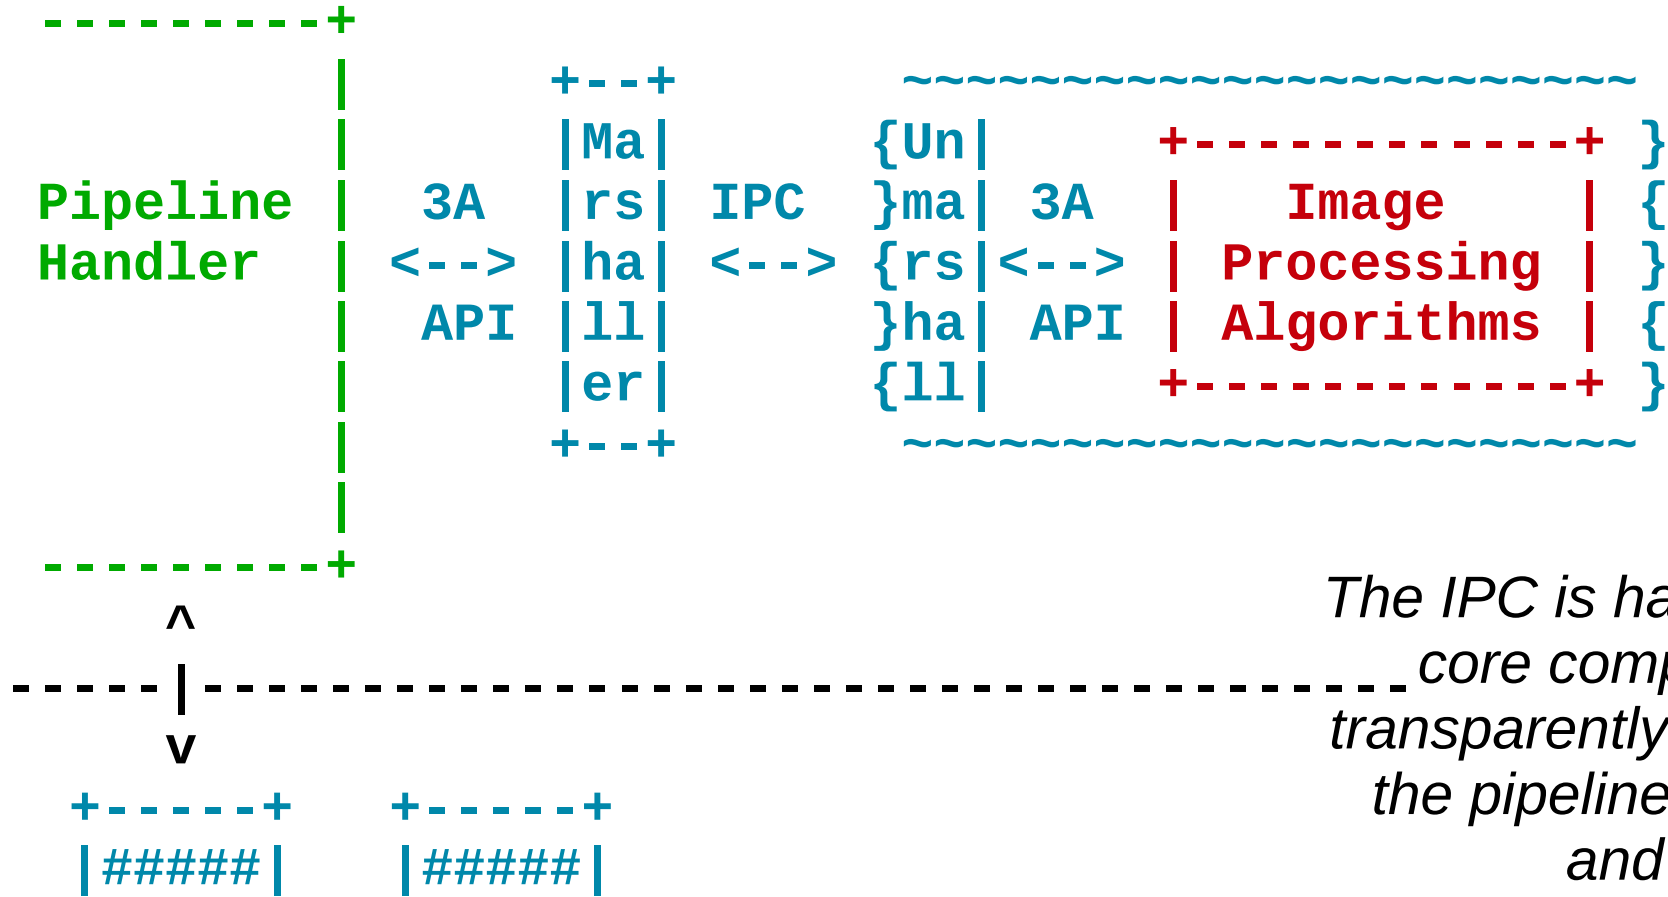




*Closed-source IPAs are sandboxed in a separate process. They communicate with the pipeline handler through IPC.*

# Image Processing Algorithms

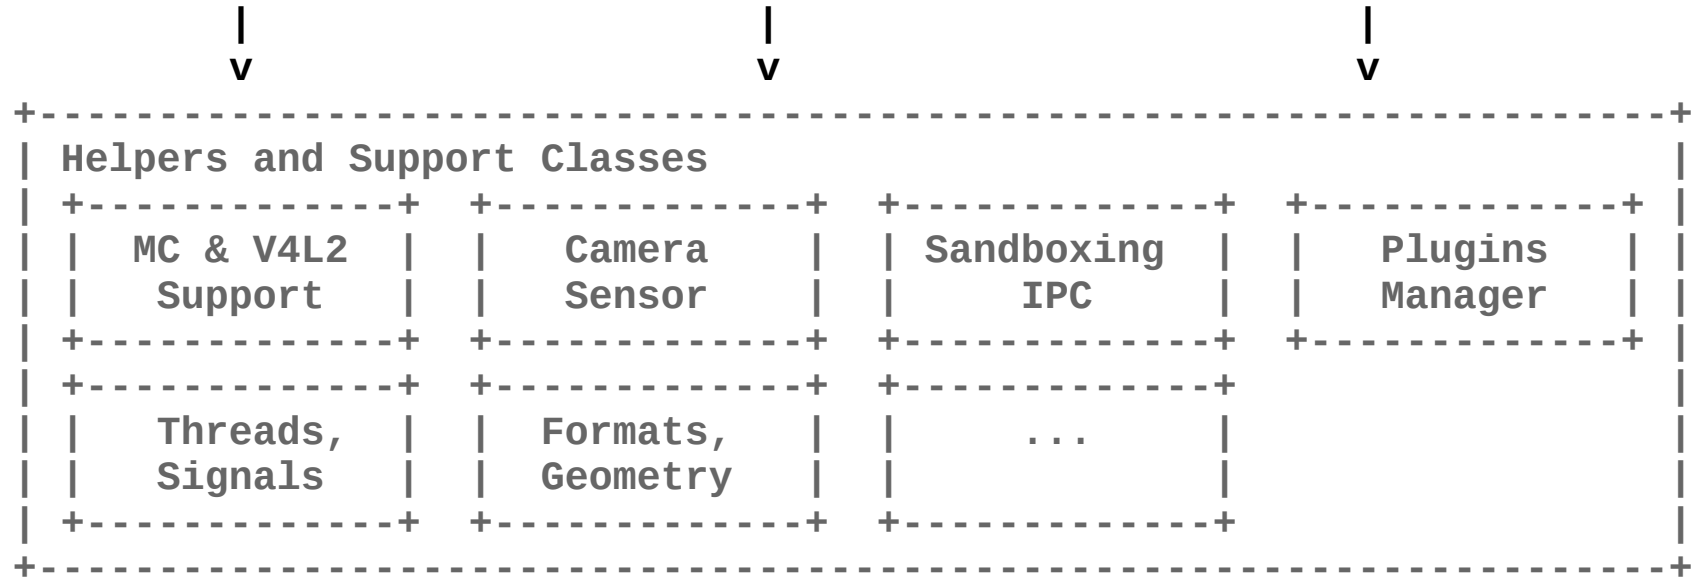




*The IPC is handled in core components, transparently for both the pipeline handler and the IPA.*

# Image Processing Algorithms

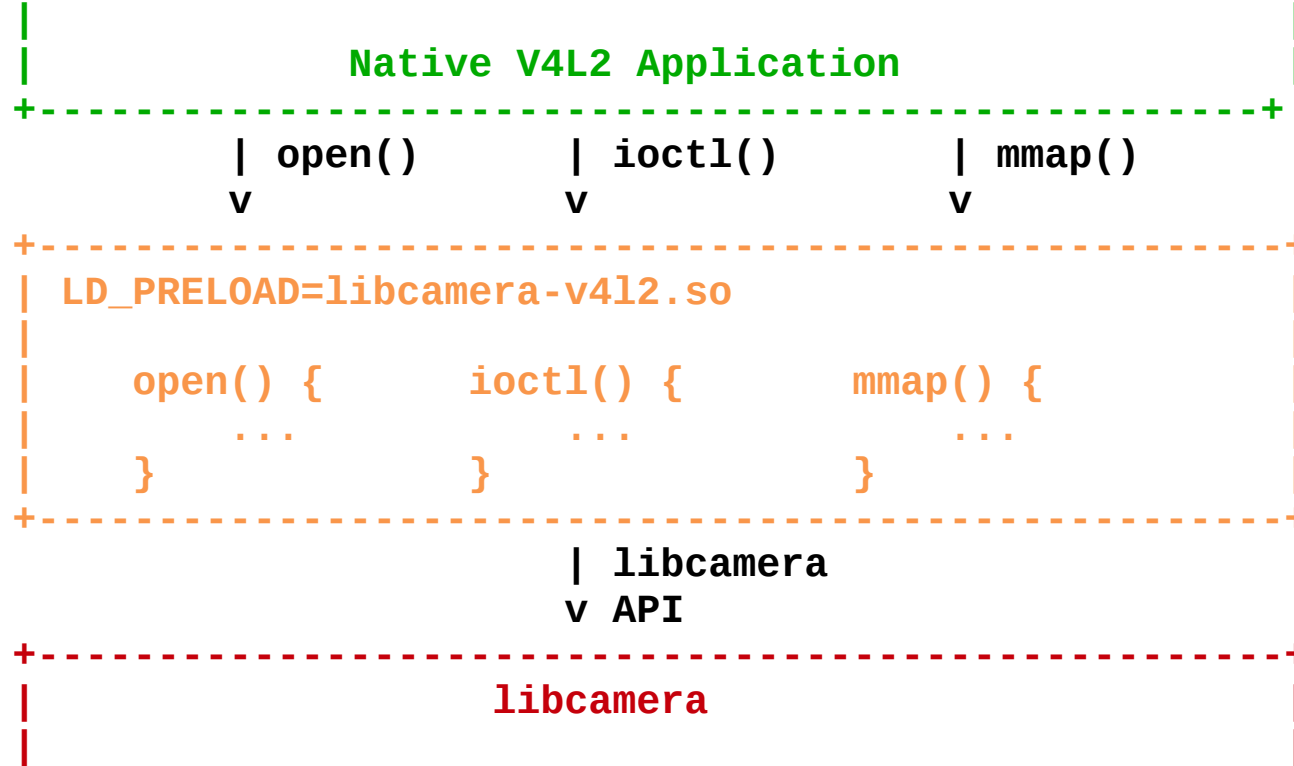




*Many helper classes ease the implementation of pipeline handlers for device vendors.*



# Helpers and Support Classes




*Native V4L2 applications are supported through a transparent compatibility layer.*

# V4L2 Compatibility



# Android Camera Framework



- HW level**
- EXTERNAL
  - LEGACY
  - LIMITED
  - FULL
  - LEVEL\_3

*A single Android camera HAL module implementation for all devices supported by libcamera.*

t  
i  
m  
e

v

# libcamera

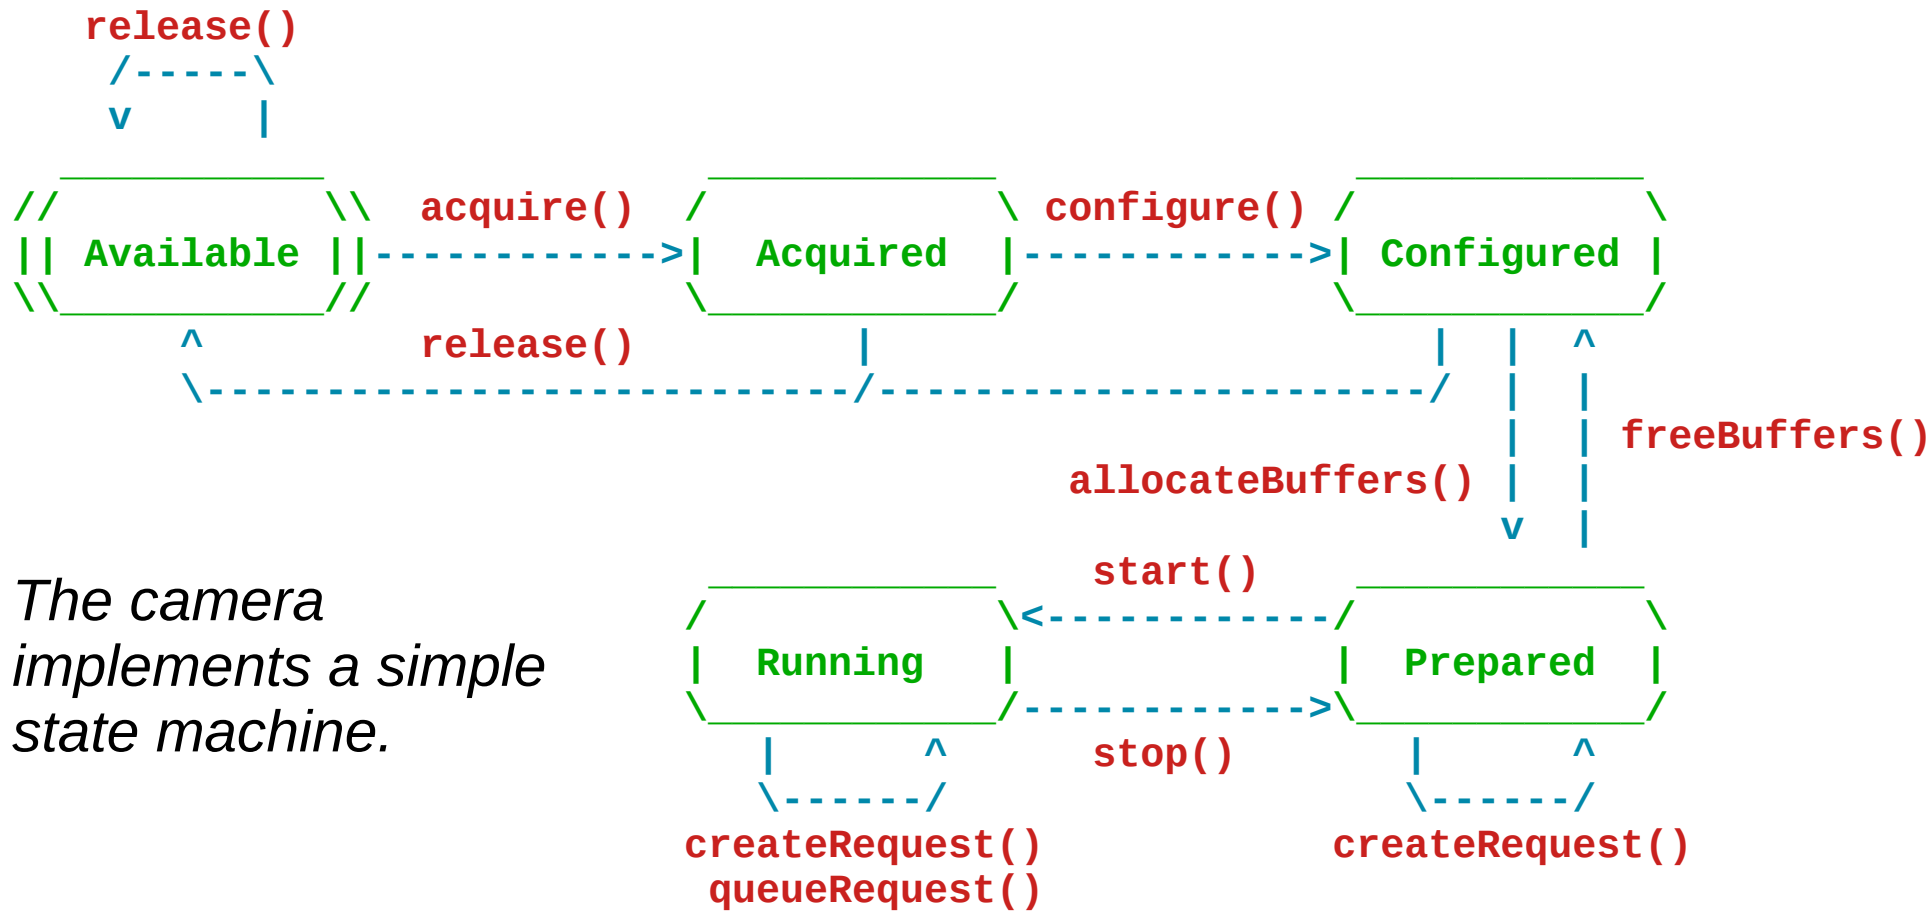
## Android Camera HAL

+ - / \ - +

| (o) |

+ - - - - +

libcamera




*The camera implements a simple state machine.*

# Camera State Machine








The camera generates  
a configuration  
template from roles.

The  
configuration is  
applied to the  
camera.

①  
.  
generateConfiguration({  
Viewfinder,  
StillCapture  
})

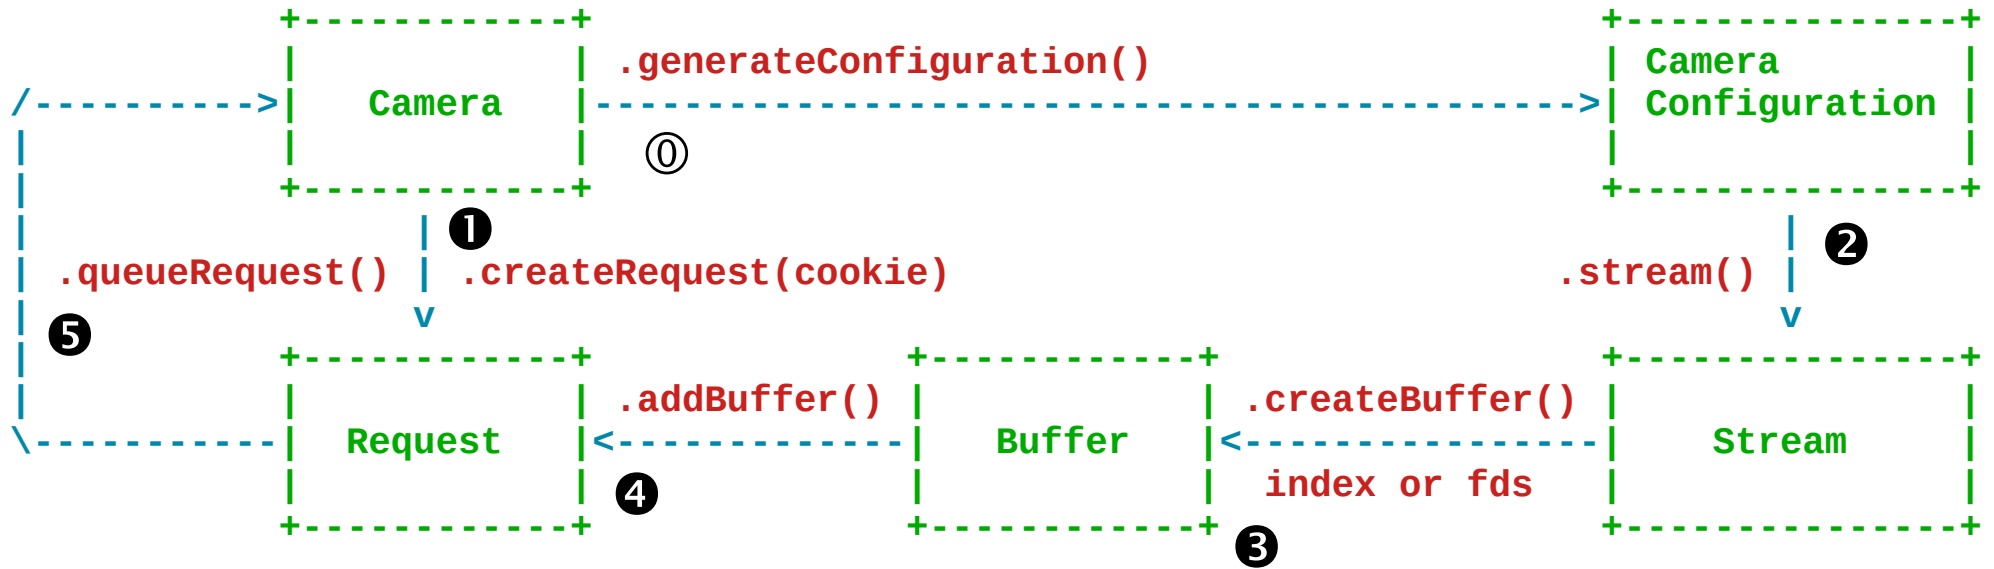


The configuration can  
be modified, and  
shall be validated.

②  
.  
width=720  
height=1280  
validate()  
addConfiguration()

# Camera Configuration

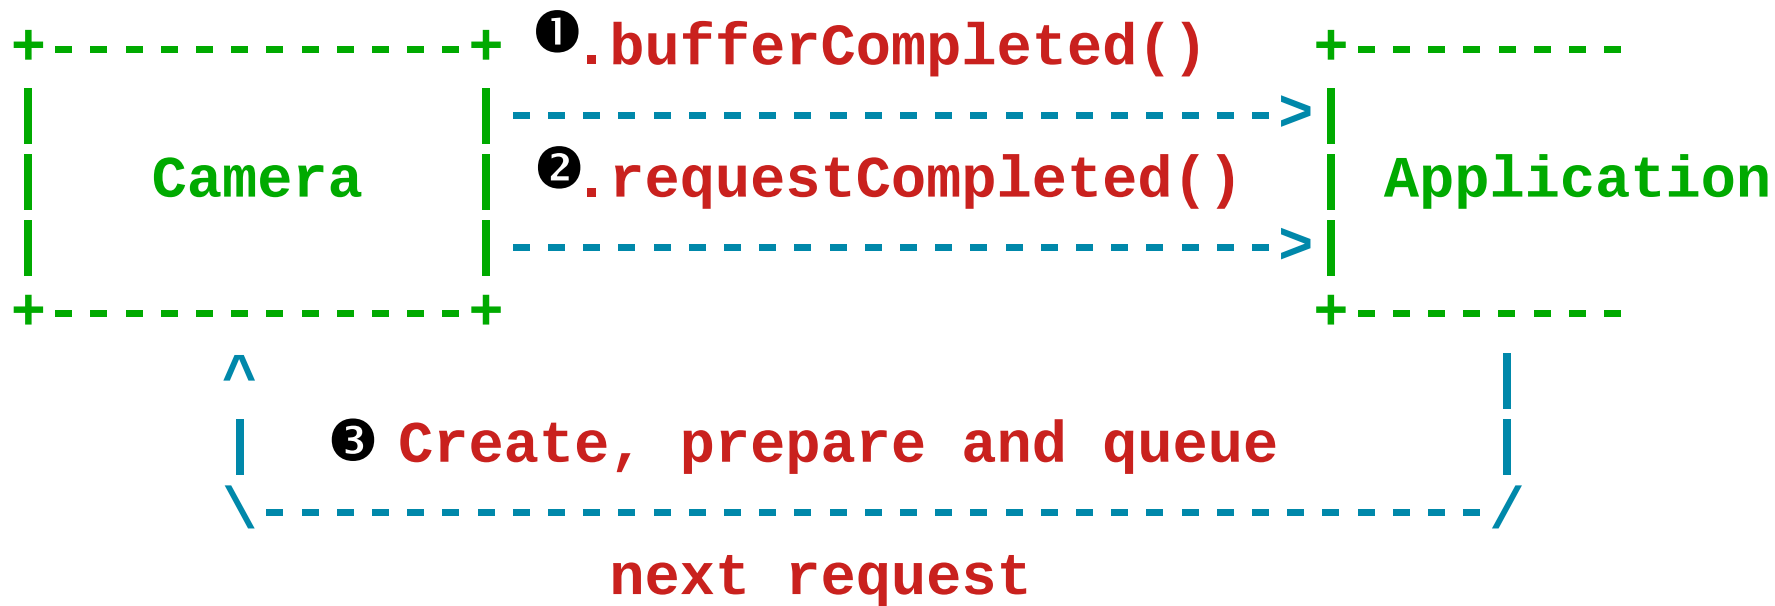




*A request is created on the Camera, populated with a Buffer for each Stream, and queued for capture.*



## Request Queuing



*Buffer and request completion are notified separately.*

*Applications submit new requests to keep the streams going.*



# Request Completion

+ - / \ - +

| (o) |

+ - - - - +

libcamera

# Contributing

libcamera is developed as a free software project and welcomes contributors. Whether you would like to help with coding, documentation, testing, proposing new features, or just discussing the project with the community, you can join our official public communication channels, or simply check out the code.

## Mailing List

We use a public mailing list as our main means of communication. You can find subscription information and the messages archive on the [libcamera-devel](#) list information page.

## IRC Channel

For informal and real time discussions, our IRC channel on Freenode is open to the public. Point your IRC client to #libcamera to say hello, or use the [WebChat](#).

## Source Code

libcamera is in early stages of development, and no releases are available yet. The source code is available from the project's [git tree](#), hosted by [LinuxTV](#).

```
$ git clone git://linuxtv.org/libcamera.git
```

## Documentation

Project documentation is created using [Sphinx](#). Source level documentation uses [Doxygen](#). Please make sure to document all code during development.

Sphinx integration with Doxygen is planned, likely using [Breathe](#) and [Exhale](#).

## Submitting Patches

### Contents

- Contributing
  - Mailing List
  - IRC Channel
  - Source Code
  - Documentation
  - Submitting Patches



# Contribute

|                           |                               |                           |                         |                                     |
|---------------------------|-------------------------------|---------------------------|-------------------------|-------------------------------------|
| <a href="#">Main Page</a> | <a href="#">Related Pages</a> | <a href="#">Classes ▾</a> | <a href="#">Files ▾</a> | <input type="text" value="Search"/> |
|---------------------------|-------------------------------|---------------------------|-------------------------|-------------------------------------|

## Class List

Here are the classes, structs, unions and interfaces with brief descriptions:

[detail level 1 2 3]

|                     |                                     |                                                                                     |
|---------------------|-------------------------------------|-------------------------------------------------------------------------------------|
| ▼ <b>N</b>          | <b>libcamera</b>                    |                                                                                     |
| <a href="#">C</a>   | <a href="#">Buffer</a>              | A buffer handle and dynamic metadata                                                |
| <a href="#">C</a>   | <a href="#">BufferMemory</a>        | A memory buffer to store an image                                                   |
| <a href="#">C</a>   | <a href="#">BufferPool</a>          | A pool of buffers                                                                   |
| <a href="#">C</a>   | <a href="#">Camera</a>              | Camera device                                                                       |
| <a href="#">C</a>   | <a href="#">CameraConfiguration</a> | Hold configuration for streams of the camera                                        |
| <a href="#">C</a>   | <a href="#">CameraData</a>          | Base class for platform-specific data associated with a camera                      |
| <a href="#">C</a>   | <a href="#">CameraManager</a>       | Provide access and manage all cameras in the system                                 |
| <a href="#">C</a>   | <a href="#">CameraSensor</a>        | A camera sensor based on V4L2 subdevices                                            |
| <a href="#">C</a>   | <a href="#">ControlIdentifier</a>   | Describe a ControlId with control specific constant meta-data                       |
| <a href="#">C</a>   | <a href="#">ControlInfo</a>         | Describe the information and capabilities of a Control                              |
| <a href="#">C</a>   | <a href="#">ControlList</a>         | Associate a list of ControlId with their values for a camera                        |
| <a href="#">C</a>   | <a href="#">ControlValue</a>        | Abstract type representing the value of a control                                   |
| <a href="#">C</a>   | <a href="#">DeviceEnumerator</a>    | Enumerate, store and search media devices                                           |
| <a href="#">C</a>   | <a href="#">DeviceMatch</a>         | Description of a media device search pattern                                        |
| <a href="#">C</a>   | <a href="#">EventDispatcher</a>     | Interface to manage the libcamera events and timers                                 |
| <a href="#">C</a>   | <a href="#">EventDispatcherPoll</a> | A poll-based event dispatcher                                                       |
| <a href="#">C</a>   | <a href="#">EventNotifier</a>       | Notify of activity on a file descriptor                                             |
| <a href="#">C</a>   | <a href="#">ImageFormats</a>        | Describe <a href="#">V4L2Device</a> and <a href="#">V4L2SubDevice</a> image formats |
| <a href="#">C</a>   | <a href="#">IPAInterface</a>        | Interface for IPA implementation                                                    |
| <a href="#">C</a>   | <a href="#">IPAManager</a>          | Manager for IPA modules                                                             |
| <a href="#">C</a>   | <a href="#">IPAModule</a>           | Wrapper around IPA module shared object                                             |
| <a href="#">C</a>   | <a href="#">IPAModuleInfo</a>       | Information of an IPA module                                                        |
| <a href="#">C</a>   | <a href="#">IPAProxy</a>            | IPA Proxy                                                                           |
| <a href="#">C</a>   | <a href="#">IPAProxyFactory</a>     | Registration of <a href="#">IPAProxy</a> classes and creation of instances          |
| ▼ <a href="#">C</a> | <a href="#">IPCUnixSocket</a>       | IPC mechanism based on Unix sockets                                                 |
| <a href="#">C</a>   | <a href="#">Payload</a>             | Container for an IPC payload                                                        |
| <a href="#">C</a>   | <a href="#">LogCategory</a>         | A category of log message                                                           |
| <a href="#">C</a>   | <a href="#">Loggable</a>            | Base class to support log message extensions                                        |
| <a href="#">C</a>   | <a href="#">Logger</a>              | <a href="#">Message</a> logger                                                      |

|                   |                                        |                                                                                                               |
|-------------------|----------------------------------------|---------------------------------------------------------------------------------------------------------------|
| <a href="#">C</a> | <a href="#">Logger</a>                 | <a href="#">Message</a> logger                                                                                |
| <a href="#">C</a> | <a href="#">LogMessage</a>             | Internal log message representation                                                                           |
| <a href="#">C</a> | <a href="#">MediaDevice</a>            | The <a href="#">MediaDevice</a> represents a Media Controller device with its full graph of connected objects |
| <a href="#">C</a> | <a href="#">MediaEntity</a>            | The <a href="#">MediaEntity</a> represents an entity in the media graph                                       |
| <a href="#">C</a> | <a href="#">MediaLink</a>              | The <a href="#">MediaLink</a> represents a link between two pads in the media graph                           |
| <a href="#">C</a> | <a href="#">MediaObject</a>            | Base class for all media objects                                                                              |
| <a href="#">C</a> | <a href="#">MediaPad</a>               | The <a href="#">MediaPad</a> represents a pad of an entity in the media graph                                 |
| <a href="#">C</a> | <a href="#">Message</a>                | A message that can be posted to a <a href="#">Thread</a>                                                      |
| <a href="#">C</a> | <a href="#">MessageQueue</a>           | A queue of posted messages                                                                                    |
| <a href="#">C</a> | <a href="#">Object</a>                 | Base object to support automatic signal disconnection                                                         |
| <a href="#">C</a> | <a href="#">PipelineHandler</a>        | Create and manage cameras based on a set of media devices                                                     |
| <a href="#">C</a> | <a href="#">PipelineHandlerFactory</a> | Registration of <a href="#">PipelineHandler</a> classes and creation of instances                             |
| <a href="#">C</a> | <a href="#">Plane</a>                  | A memory region to store a single plane of a frame                                                            |
| <a href="#">C</a> | <a href="#">Process</a>                | <a href="#">Process</a> object                                                                                |
| <a href="#">C</a> | <a href="#">ProcessManager</a>         | Manager of processes                                                                                          |
| <a href="#">C</a> | <a href="#">Rectangle</a>              | Describe a rectangle's position and dimensions                                                                |
| <a href="#">C</a> | <a href="#">Request</a>                | A frame capture request                                                                                       |
| <a href="#">C</a> | <a href="#">Signal</a>                 | Generic signal and slot communication mechanism                                                               |
| <a href="#">C</a> | <a href="#">SignalMessage</a>          | A message carrying a <a href="#">Signal</a> across threads                                                    |
| <a href="#">C</a> | <a href="#">Size</a>                   | Describe a two-dimensional size                                                                               |
| <a href="#">C</a> | <a href="#">SizeRange</a>              | Describe a range of sizes                                                                                     |
| <a href="#">C</a> | <a href="#">Stream</a>                 | Video stream for a camera                                                                                     |
| <a href="#">C</a> | <a href="#">StreamConfiguration</a>    | Configuration parameters for a stream                                                                         |
| <a href="#">C</a> | <a href="#">StreamFormats</a>          | Hold information about supported stream formats                                                               |
| <a href="#">C</a> | <a href="#">Thread</a>                 | A thread of execution                                                                                         |
| <a href="#">C</a> | <a href="#">ThreadData</a>             | Thread-local internal data                                                                                    |
| <a href="#">C</a> | <a href="#">ThreadMain</a>             | <a href="#">Thread</a> wrapper for the main thread                                                            |
| <a href="#">C</a> | <a href="#">Timer</a>                  | Single-shot timer interface                                                                                   |
| <a href="#">C</a> | <a href="#">V4L2Capability</a>         | Struct v4l2_capability object wrapper and helpers                                                             |
| <a href="#">C</a> | <a href="#">V4L2Control</a>            | A V4L2 control value                                                                                          |
| <a href="#">C</a> | <a href="#">V4L2ControlInfo</a>        | Information on a V4L2 control                                                                                 |
| <a href="#">C</a> | <a href="#">V4L2ControlList</a>        | Container of <a href="#">V4L2Control</a> instances                                                            |
| <a href="#">C</a> | <a href="#">V4L2Device</a>             | Base class for <a href="#">V4L2VideoDevice</a> and <a href="#">V4L2Subdevice</a>                              |
| <a href="#">C</a> | <a href="#">V4L2DeviceFormat</a>       | The V4L2 video device image format and sizes                                                                  |
| <a href="#">C</a> | <a href="#">V4L2Subdevice</a>          | A V4L2 subdevice as exposed by the Linux kernel                                                               |
| <a href="#">C</a> | <a href="#">V4L2SubdeviceFormat</a>    | The V4L2 sub-device image format and sizes                                                                    |
| <a href="#">C</a> | <a href="#">V4L2VideoDevice</a>        | <a href="#">V4L2VideoDevice</a> object and API                                                                |



libcamera-devel@lists.libcamera.org  
irc://chat.freenode.net/#libcamera

laurent.pinchart@ideasonboard.com



## Contact

? !



ご清聴  
ありがとう  
ございました！

