UNORTHORDOX DESIGN PATTERNS IN RABBITMQ

AYANDA DUBE

ERLANG SOLUTIONS

ACKNOWLEDGEMENTS

- **Joe Armstrong** - the legacy
- Erlang Solutions conferences (Gurpreet & crew)

OVERVIEW: RABBITMQ

- Erlang & Elixir AMQP implementation
- Client libraries (Java, .Net, Objective-C, JMS, PHP,)
- Been around for a while - over 10years
- More than 35,000 known deployments

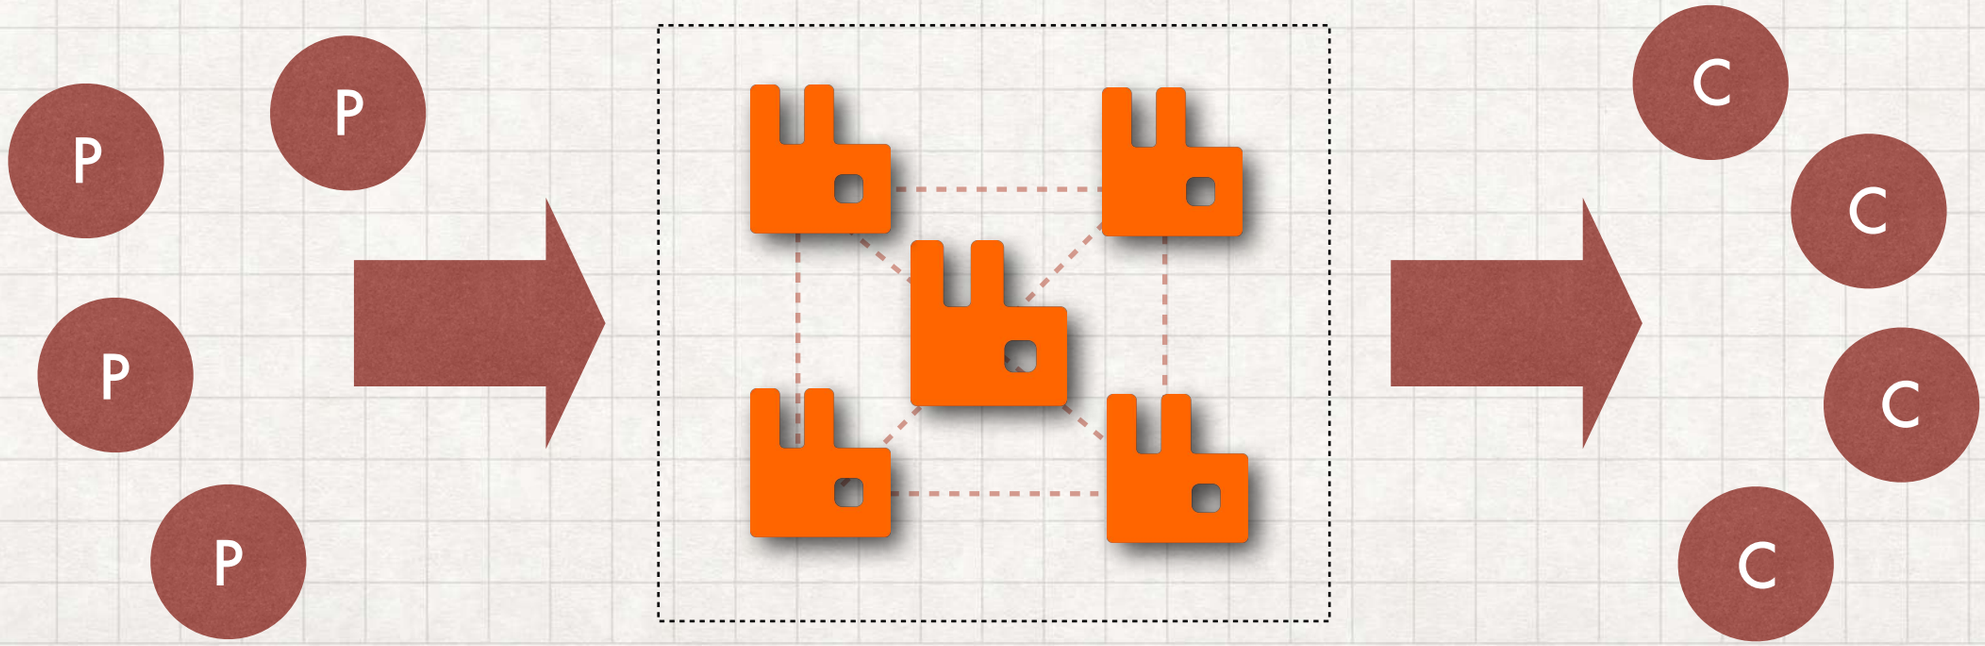

OVERVIEW: INTERNALS

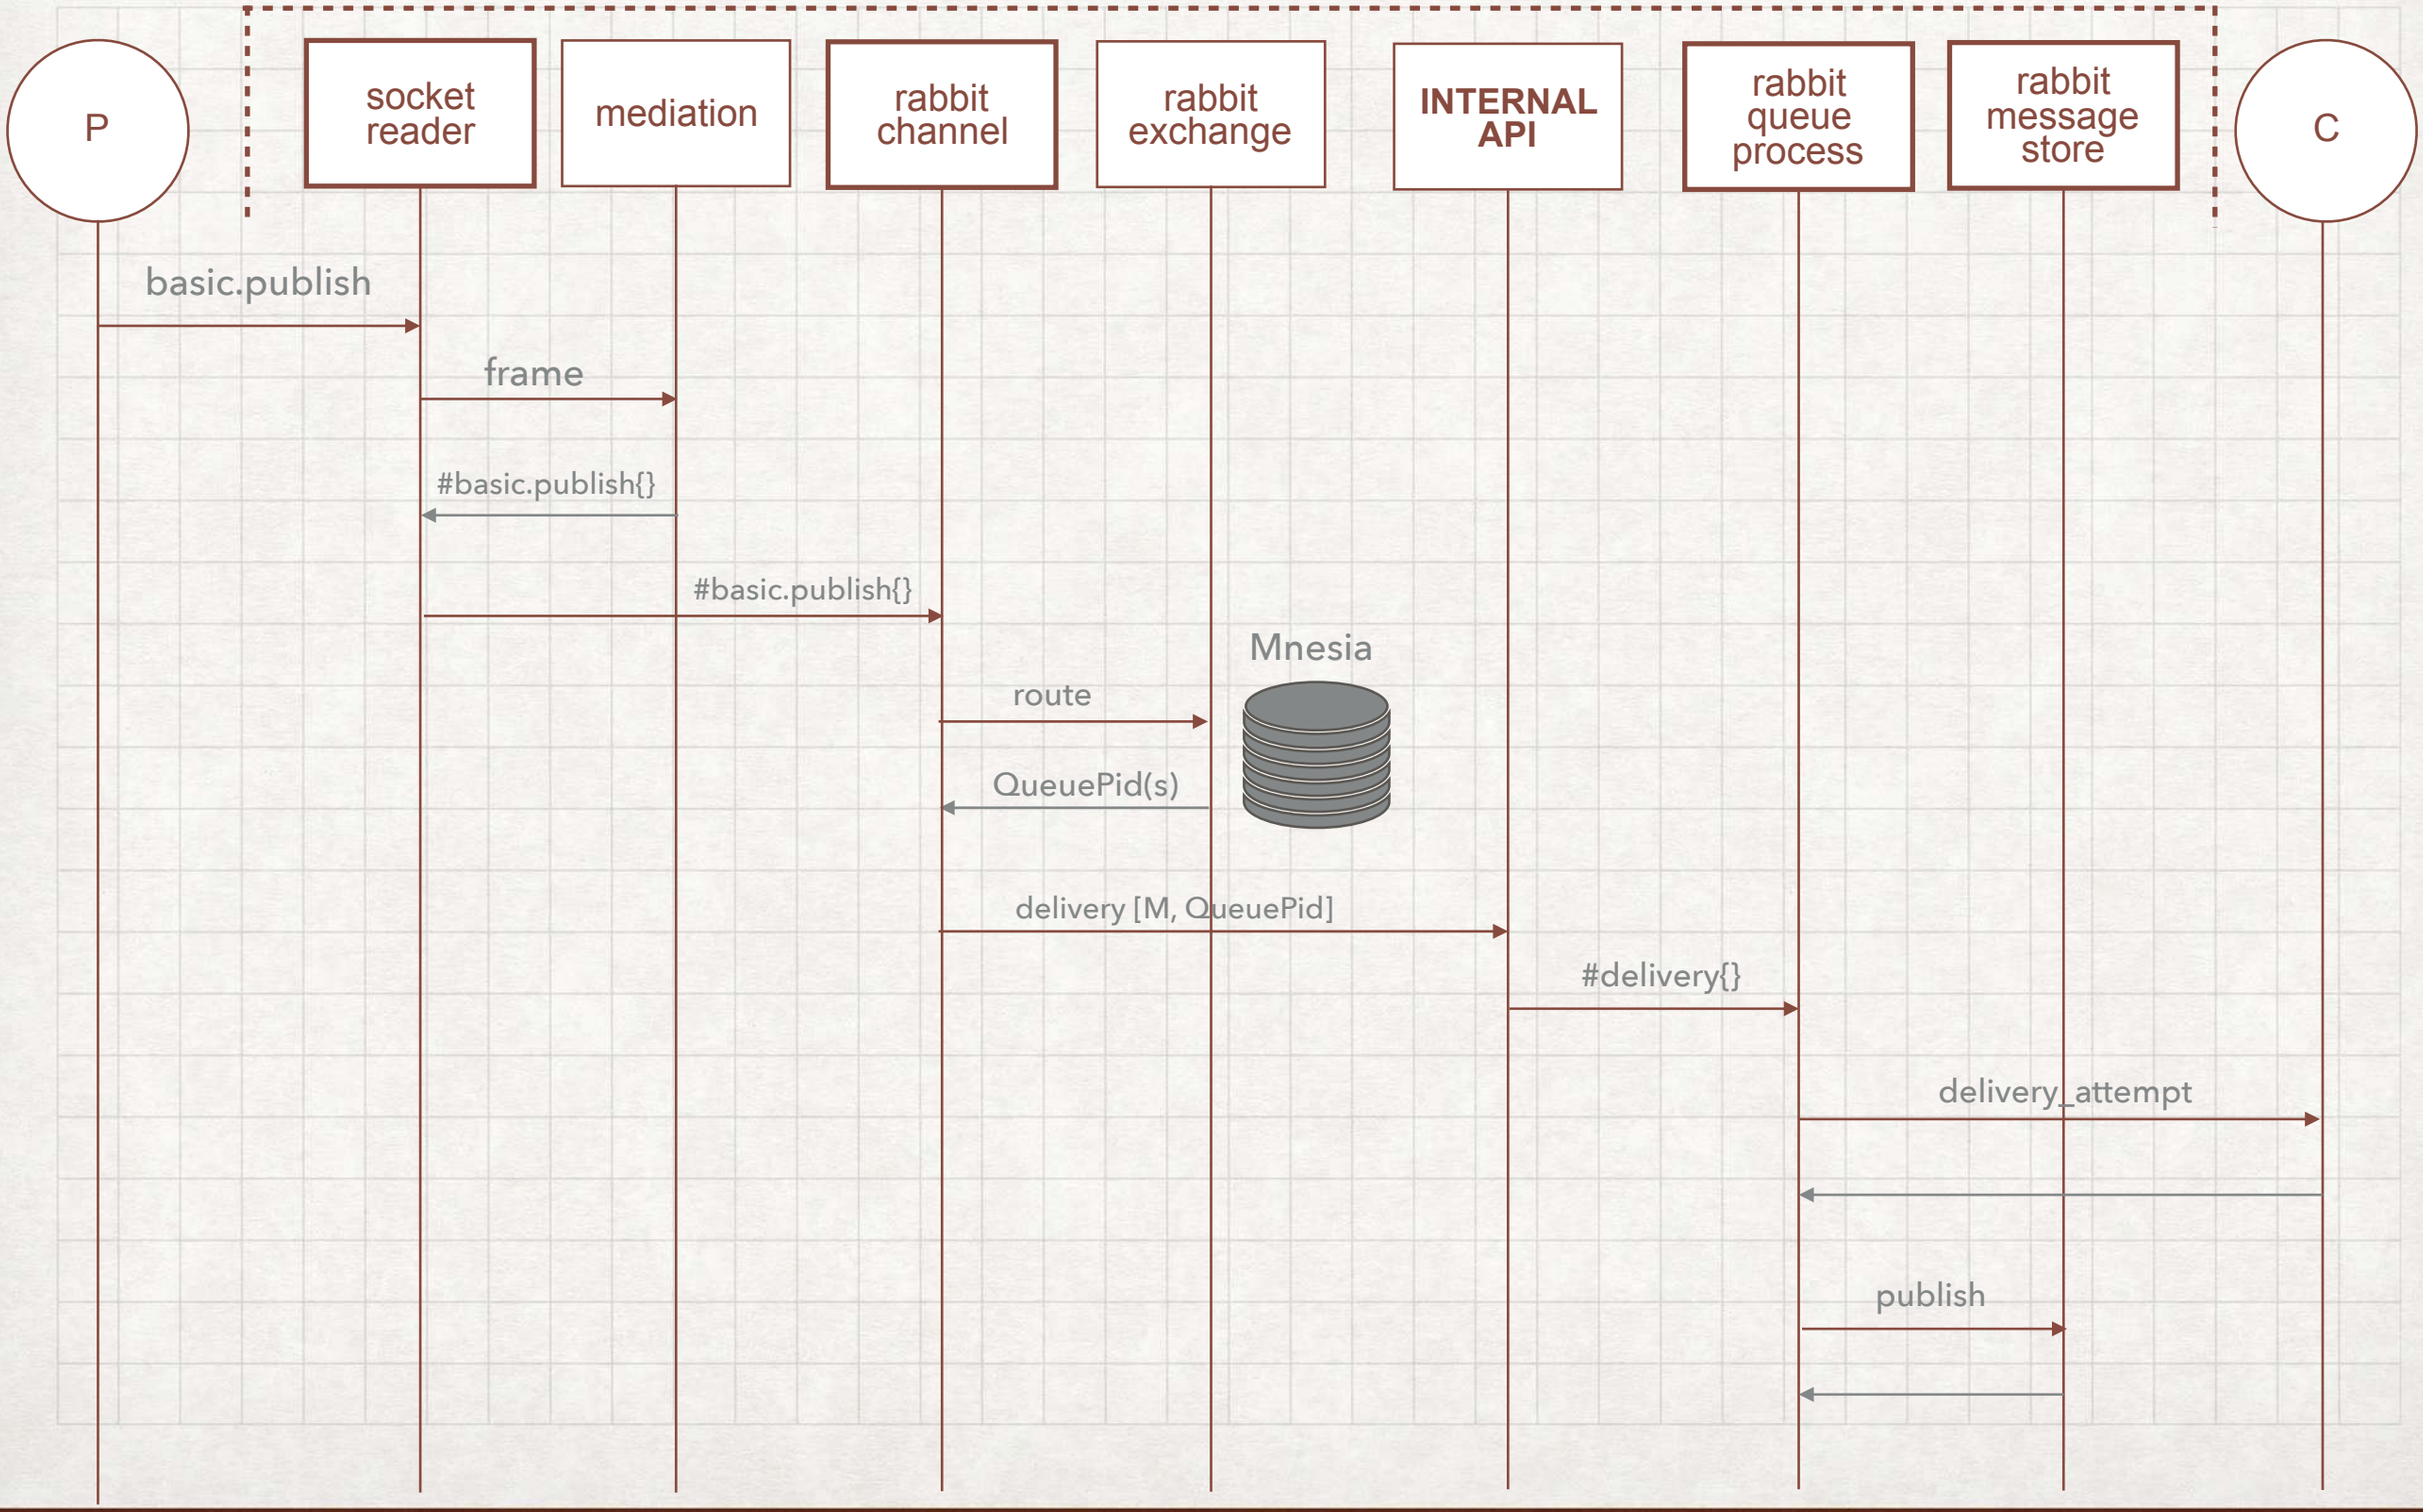

OVERVIEW: ROADMAP

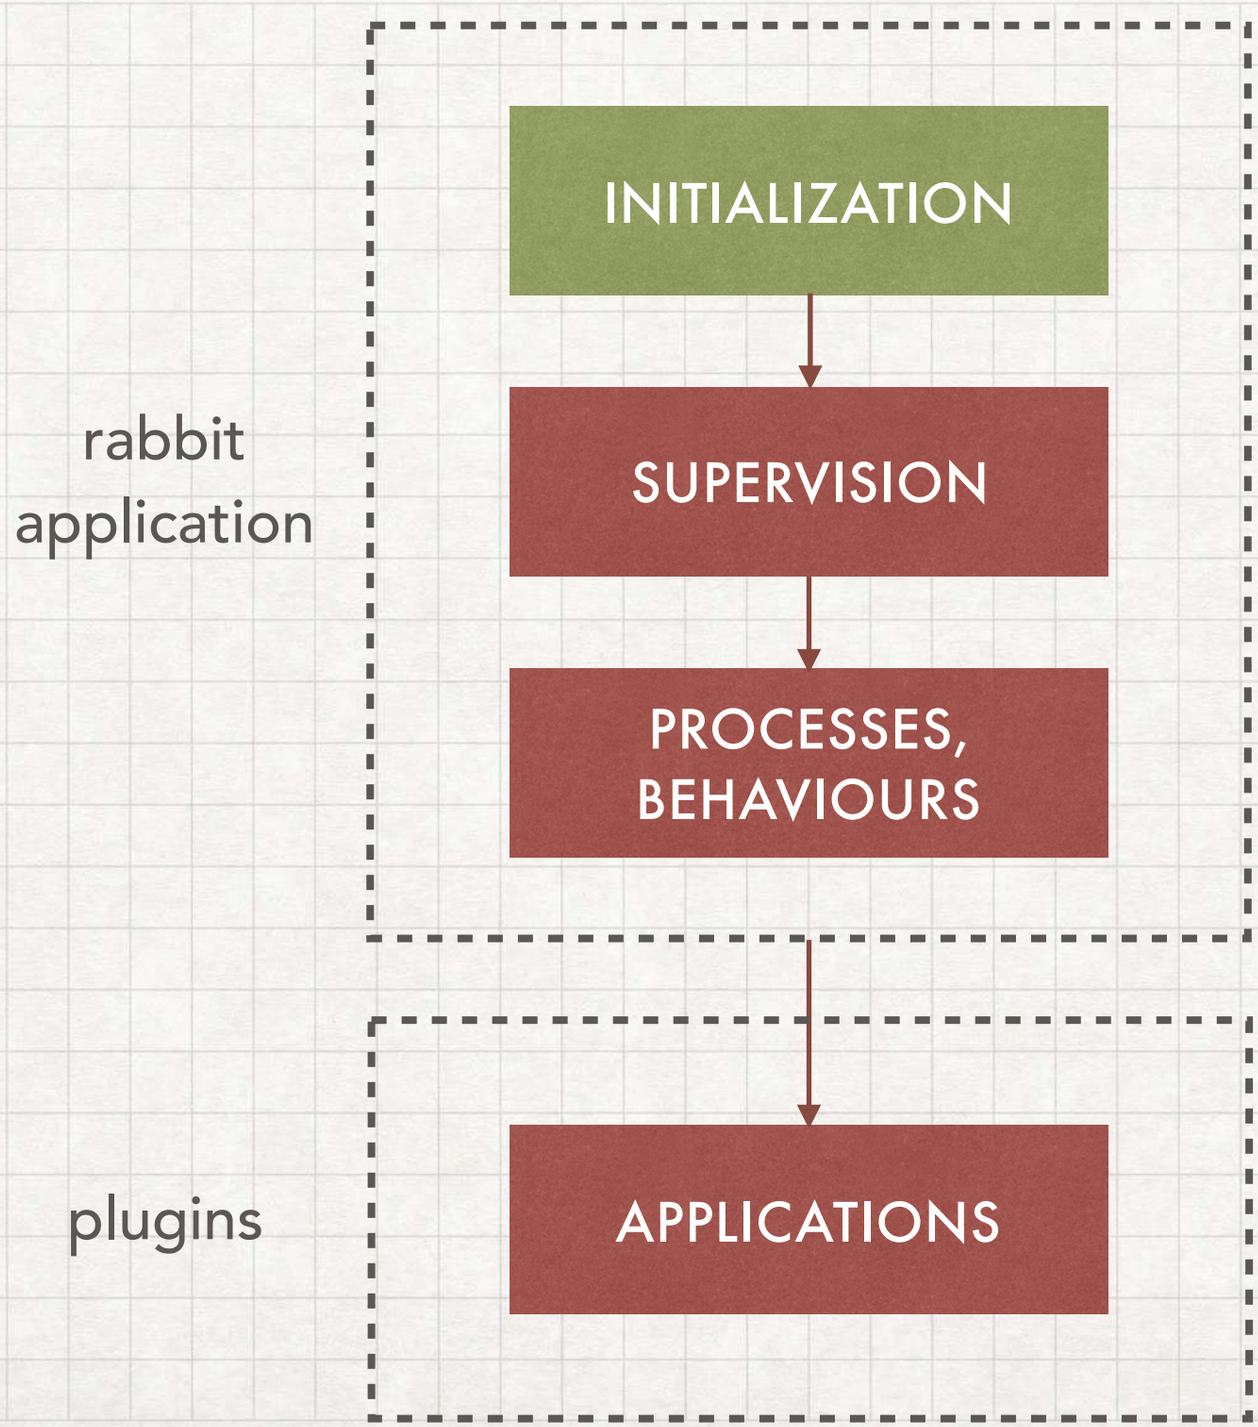

INITIALIZATION: BOOTSTEPS

- Initialisation: **boot steps**
- Loosely coupled application procedures/steps
- Ordered execution of initialisation steps
- Defined and set as **module attributes**
- Pre-conditions and post-conditions
- Cleanup capabilities on shutdown
- Alternative to **OTP** application start-phases

INITIALIZATION: BOOTSTEPS

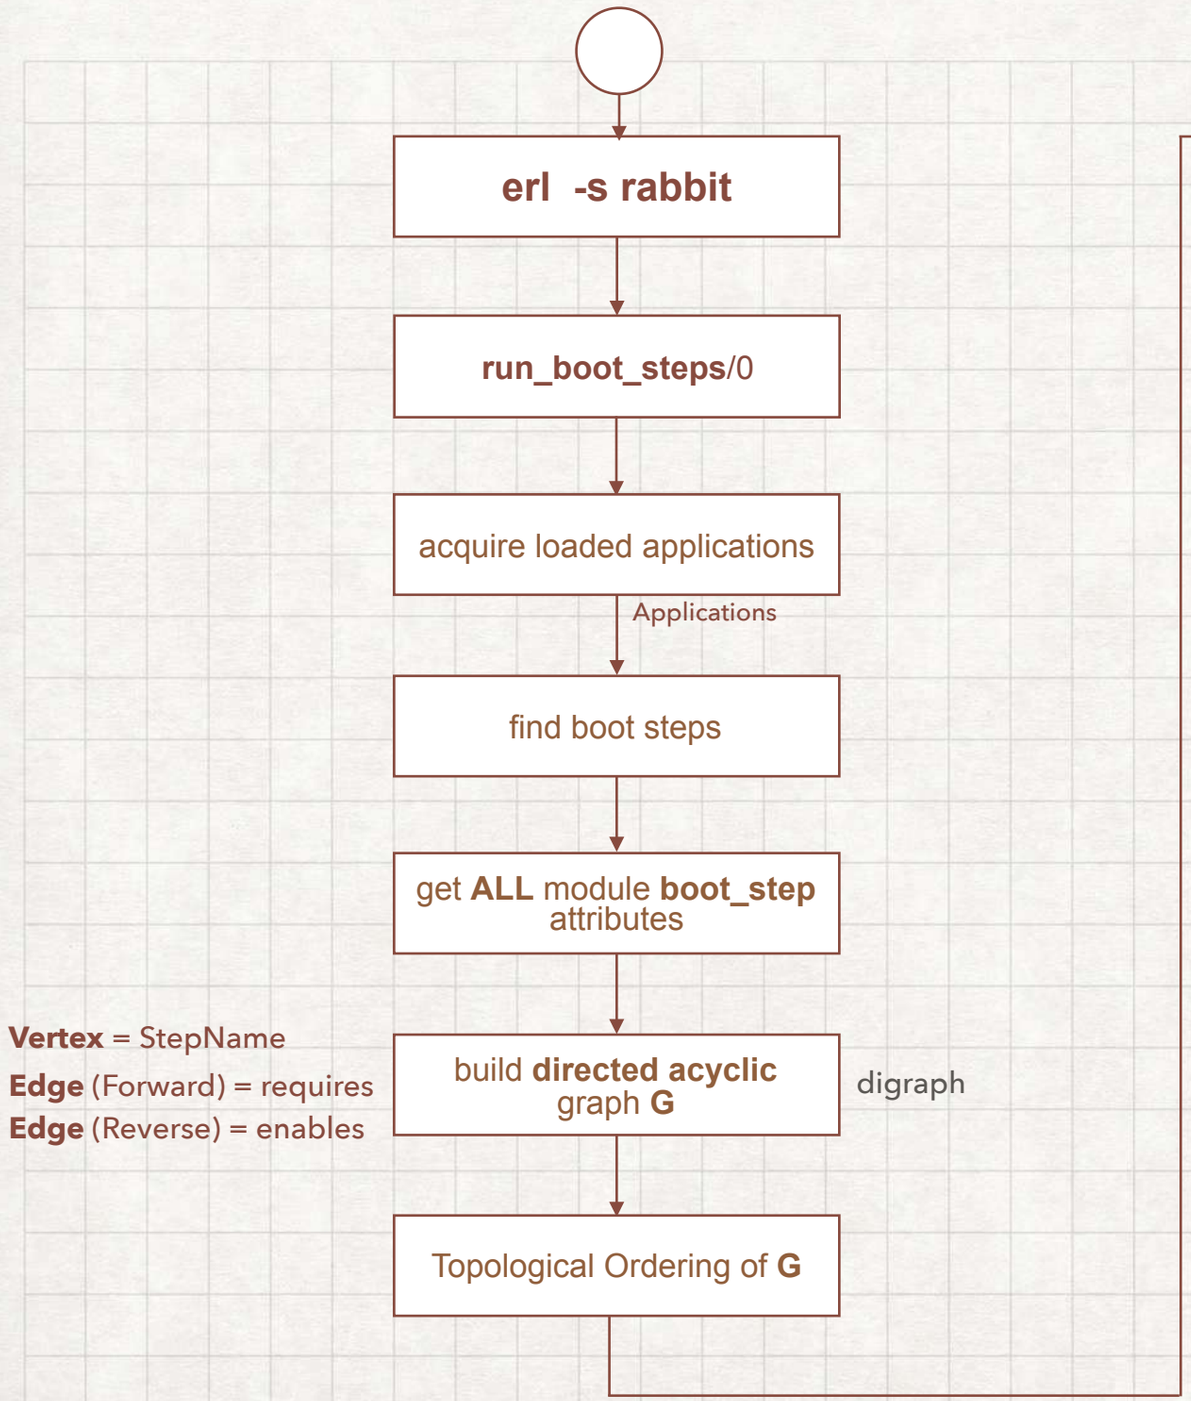

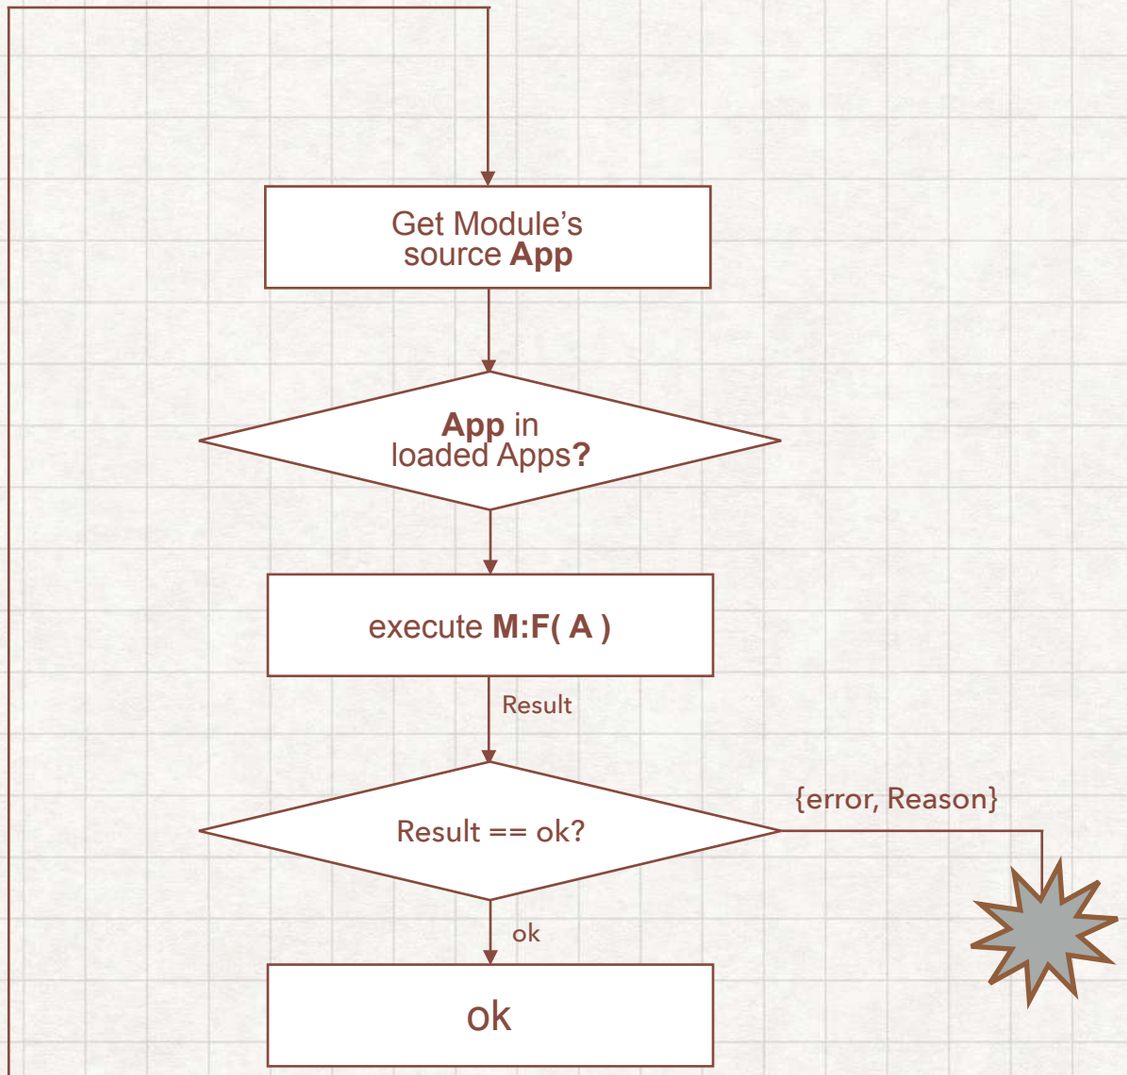

INITIALIZATION: BOOTSTEPS

```
17  -module(rabbit_exchange_type_fanout).
18  -include("rabbit.hrl").
19
20  -behaviour(rabbit_exchange_type).
21
22  -export([description/0, serialise_events/0, route/2]).
23  -export([validate/1, validate_binding/2,
24           create/2, delete/3, policy_changed/2, add_binding/3,
25           remove_bindings/3, assert_args_equivalence/2]).
26  -export([info/1, info/2]).
27
28  -rabbit_boot_step({?MODULE,
29                     [{description, "exchange type fanout"},
30                      {mfa,         {rabbit_registry, register,
31                                     [exchange, <<"fanout">>, ?MODULE]}},
32                      {requires,    rabbit_registry},
33                      {enables,     kernel_ready}]}).
```

INITIALIZATION: BOOTSTEPS

```
-rabbit_boot_step({rabbit_core_metrics,
                   [{description, "core metrics storage"},
                    {mfa,         {rabbit_sup, start_child,
                                   [rabbit_metrics]}},
                    {requires,    pre_boot},
                    {enables,     external_infrastructure}]}).

-rabbit_boot_step({rabbit_event,
                   [{description, "statistics event manager"},
                    {mfa,         {rabbit_sup, start_restartable_child,
                                   [rabbit_event]}},
                    {requires,    external_infrastructure},
                    {enables,     kernel_ready}]}).

-rabbit_boot_step({kernel_ready,
                   [{description, "kernel ready"},
                    {requires,    external_infrastructure}]}).

-rabbit_boot_step({rabbit_memory_monitor,
                   [{description, "memory monitor"},
                    {mfa,         {rabbit_sup, start_restartable_child,
                                   [rabbit_memory_monitor]}},
                    {requires,    rabbit_alarm},
                    {enables,     core_initialized}]}).
```

INITIALIZATION: BOOTSTEPS

- Topologically sorted: digraph_utils:topsort(G)

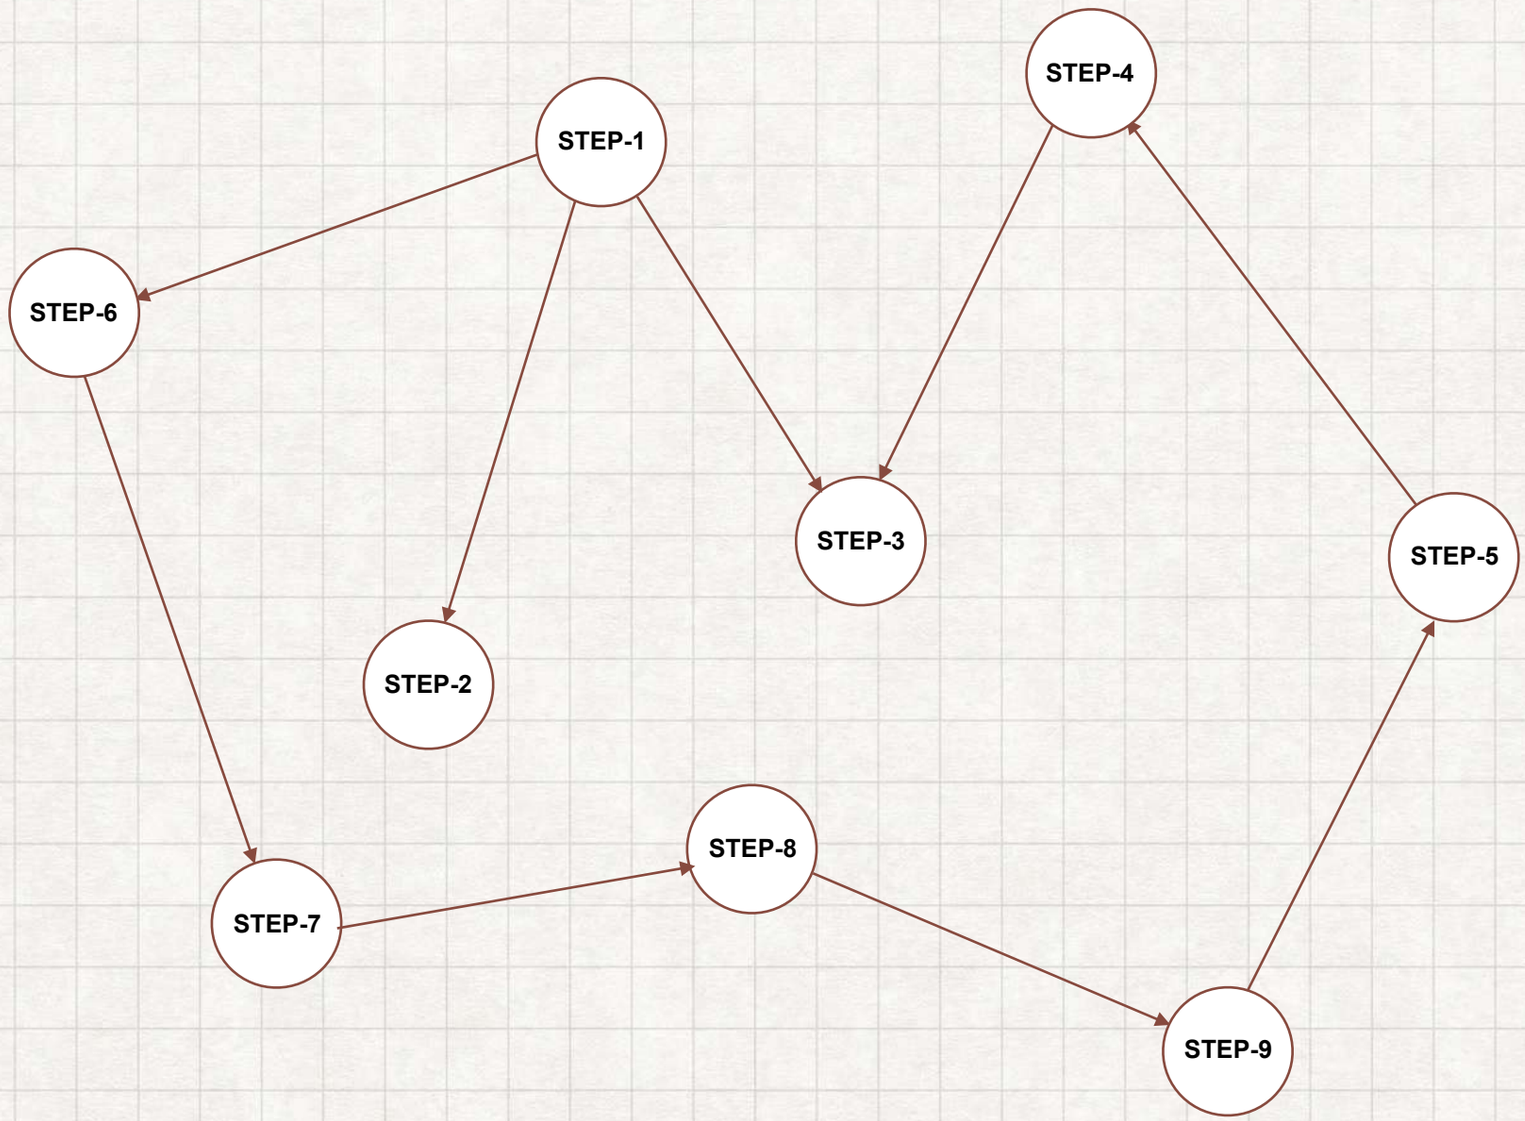

[STEP-1, STEP-6, STEP-7, STEP-8, STEP-9, STEP-5, STEP-4, STEP-2, STEP-3]

SUPERVISION

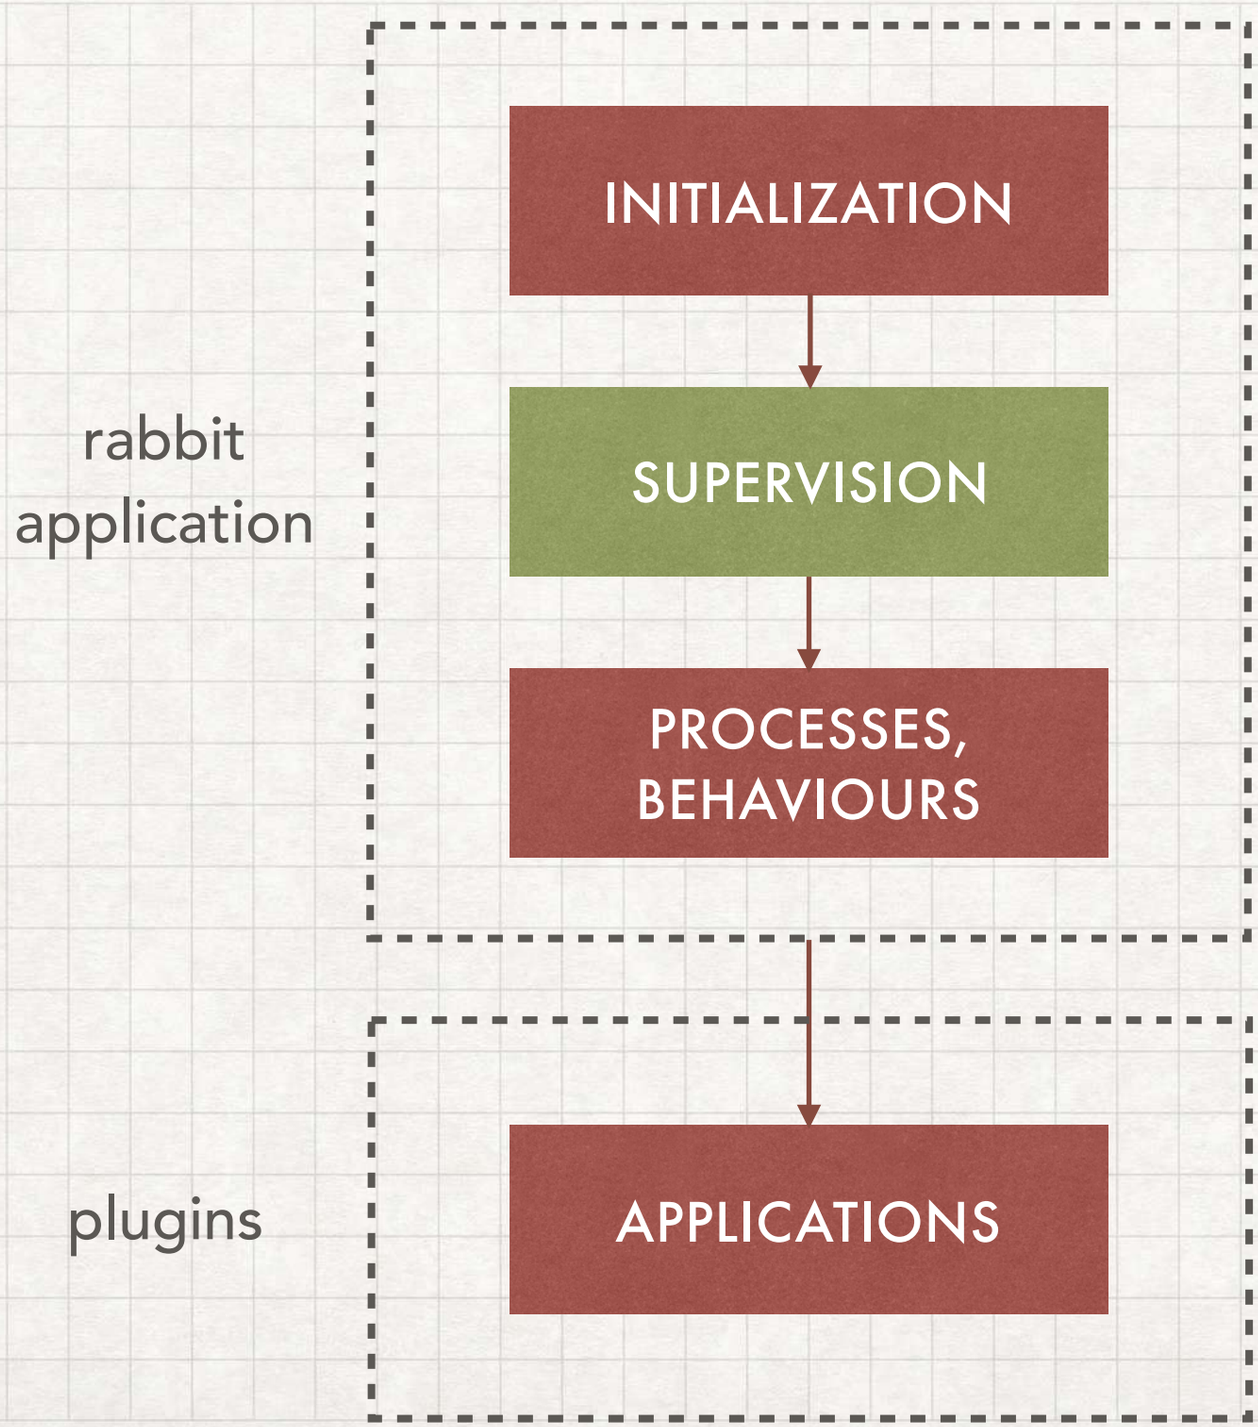

SUPERVISORS: SUPERVISOR2

- Extension of OTP supervisor
- **Intrinsic** restart type (restarts on abnormal exists)
  - If child exists normally, sup also exits normally
- Delayed restart types, e.g. {intrinsic, Delay}
  - Sup continues after Delay to restart child if MaxRestarts and MaxTime were exceeded
- Find child utilities, ...

SUPERVISORS: MIRRORED SUPERVISOR

- Multiple supervisors within a single process group
- Child specifications retained in Mnesia
- Processes than need to exist once in a cluster
- Low state footprint
- Process recovery on separate node in case of node failure
- {global, Name} registration **not** supported

SUPERVISORS: MIRRORED SUPERVISOR

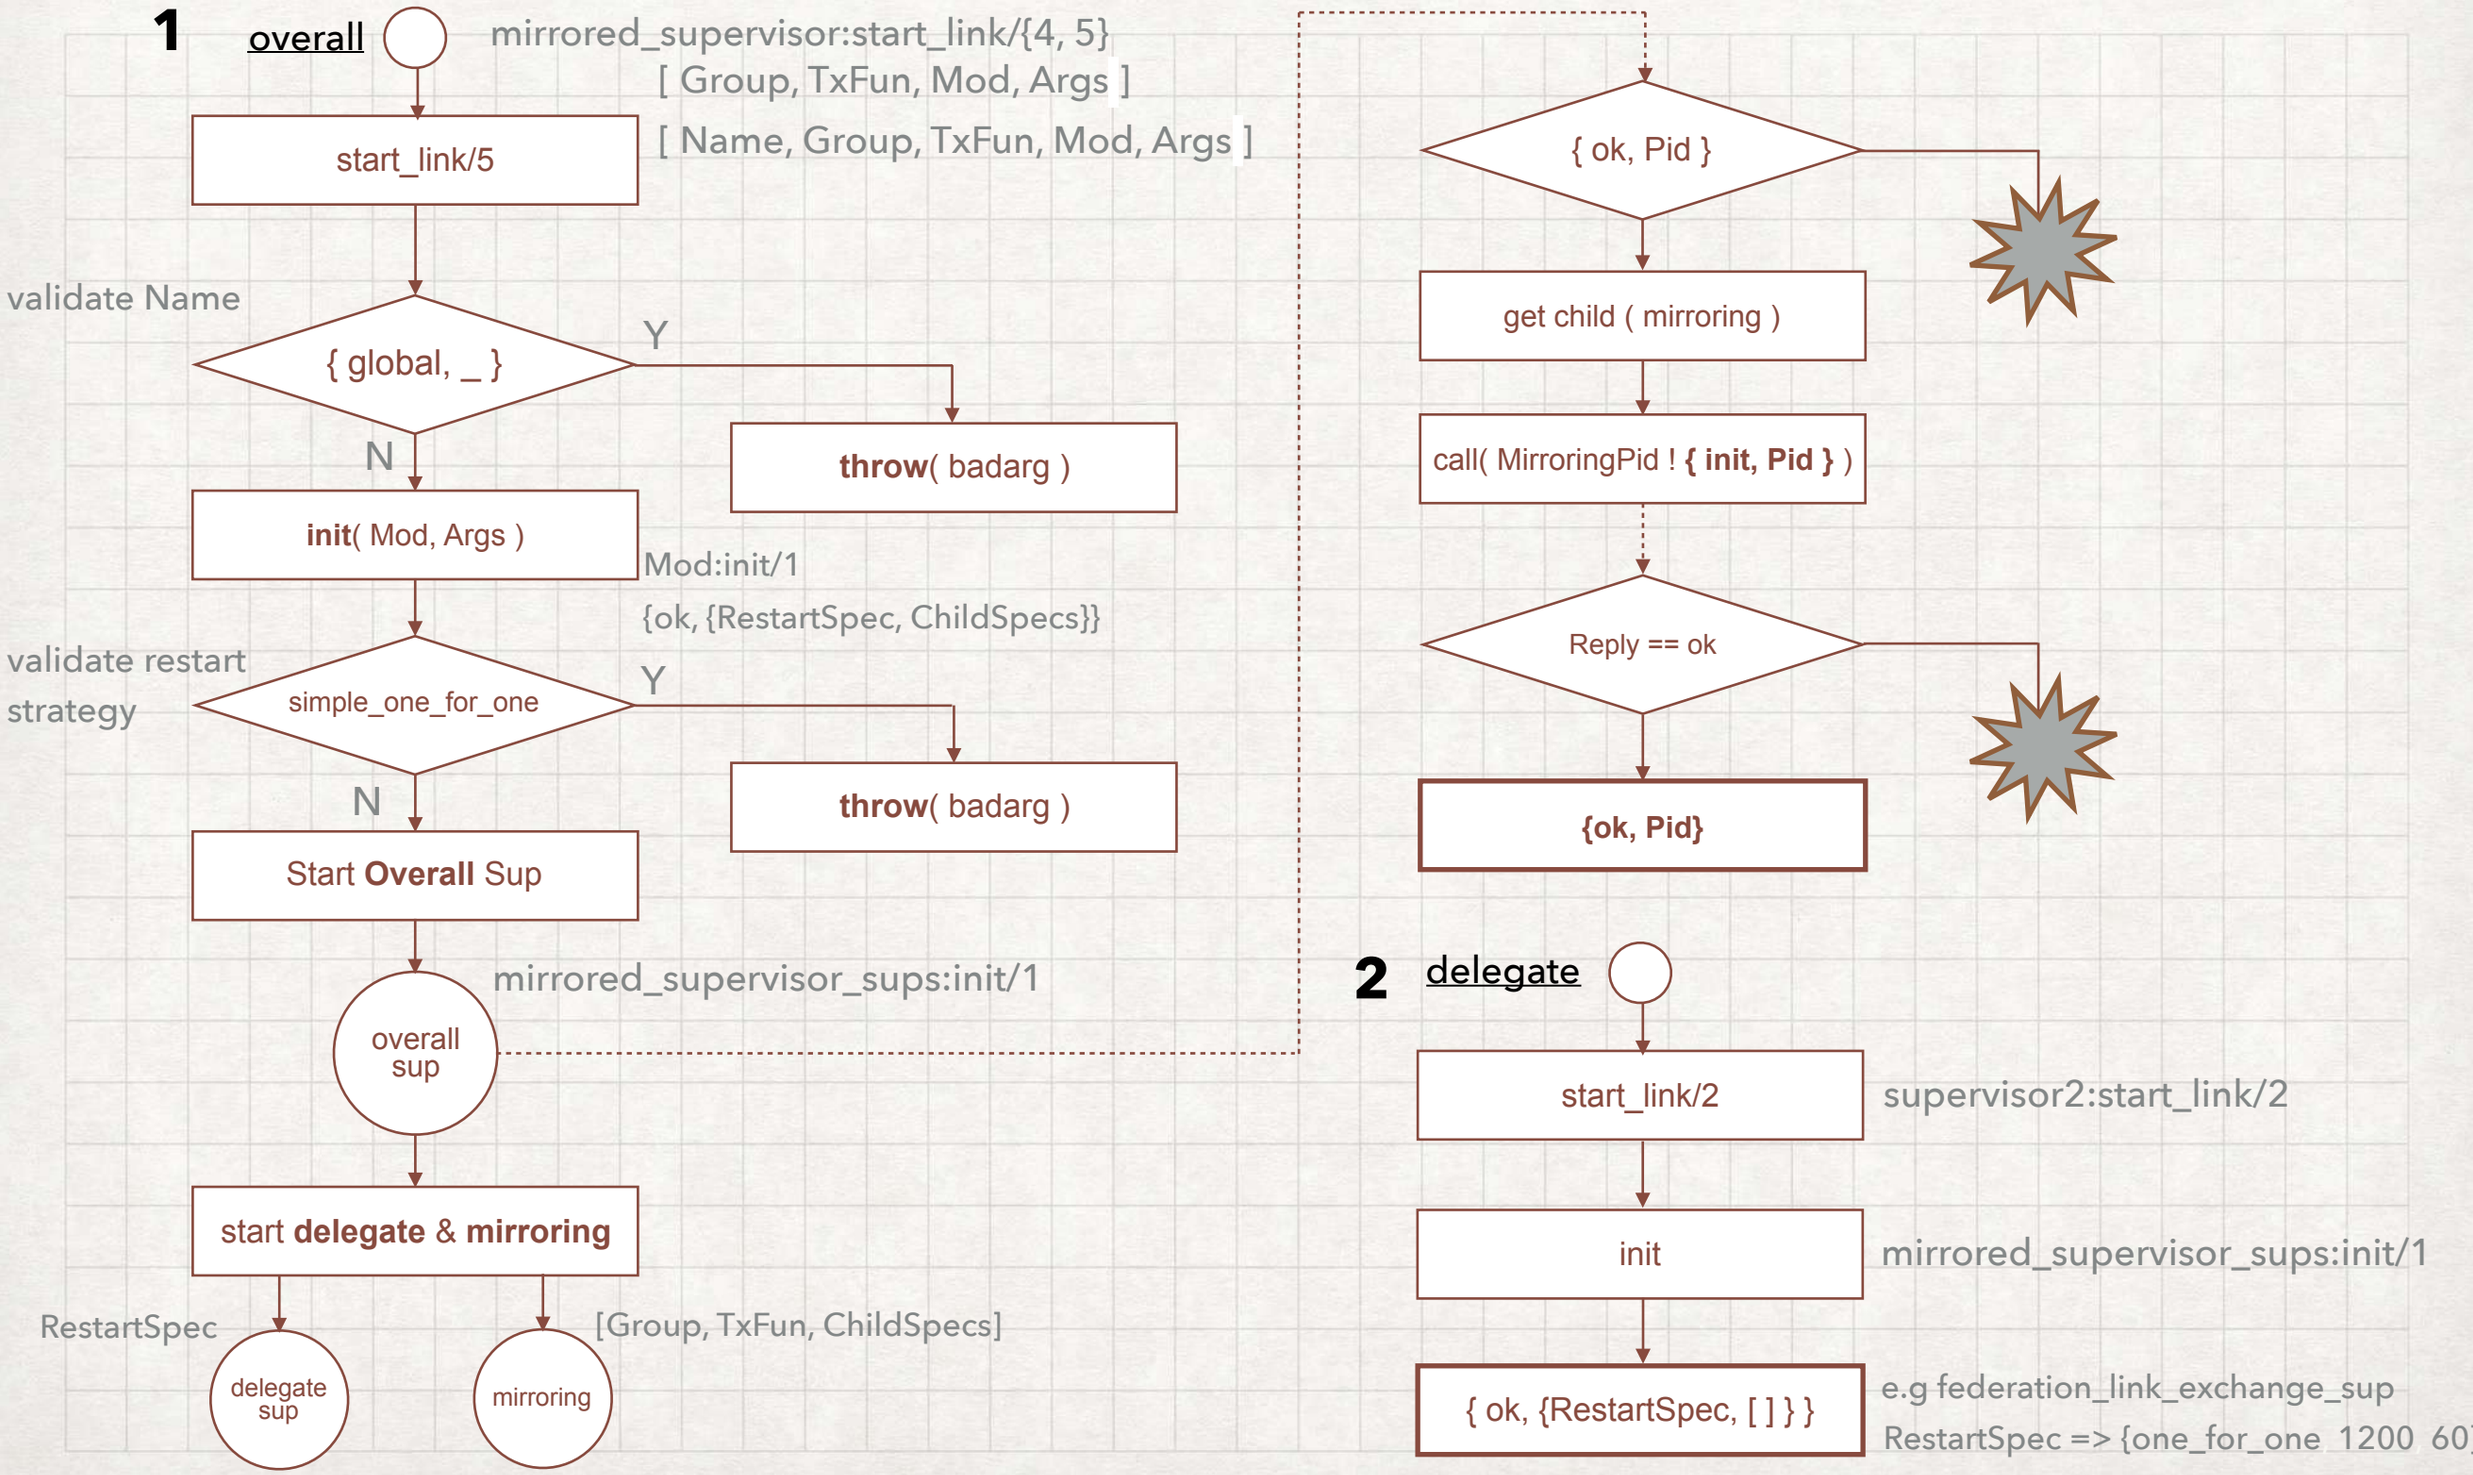

SUPERVISORS: MIRRORED SUPERVISOR

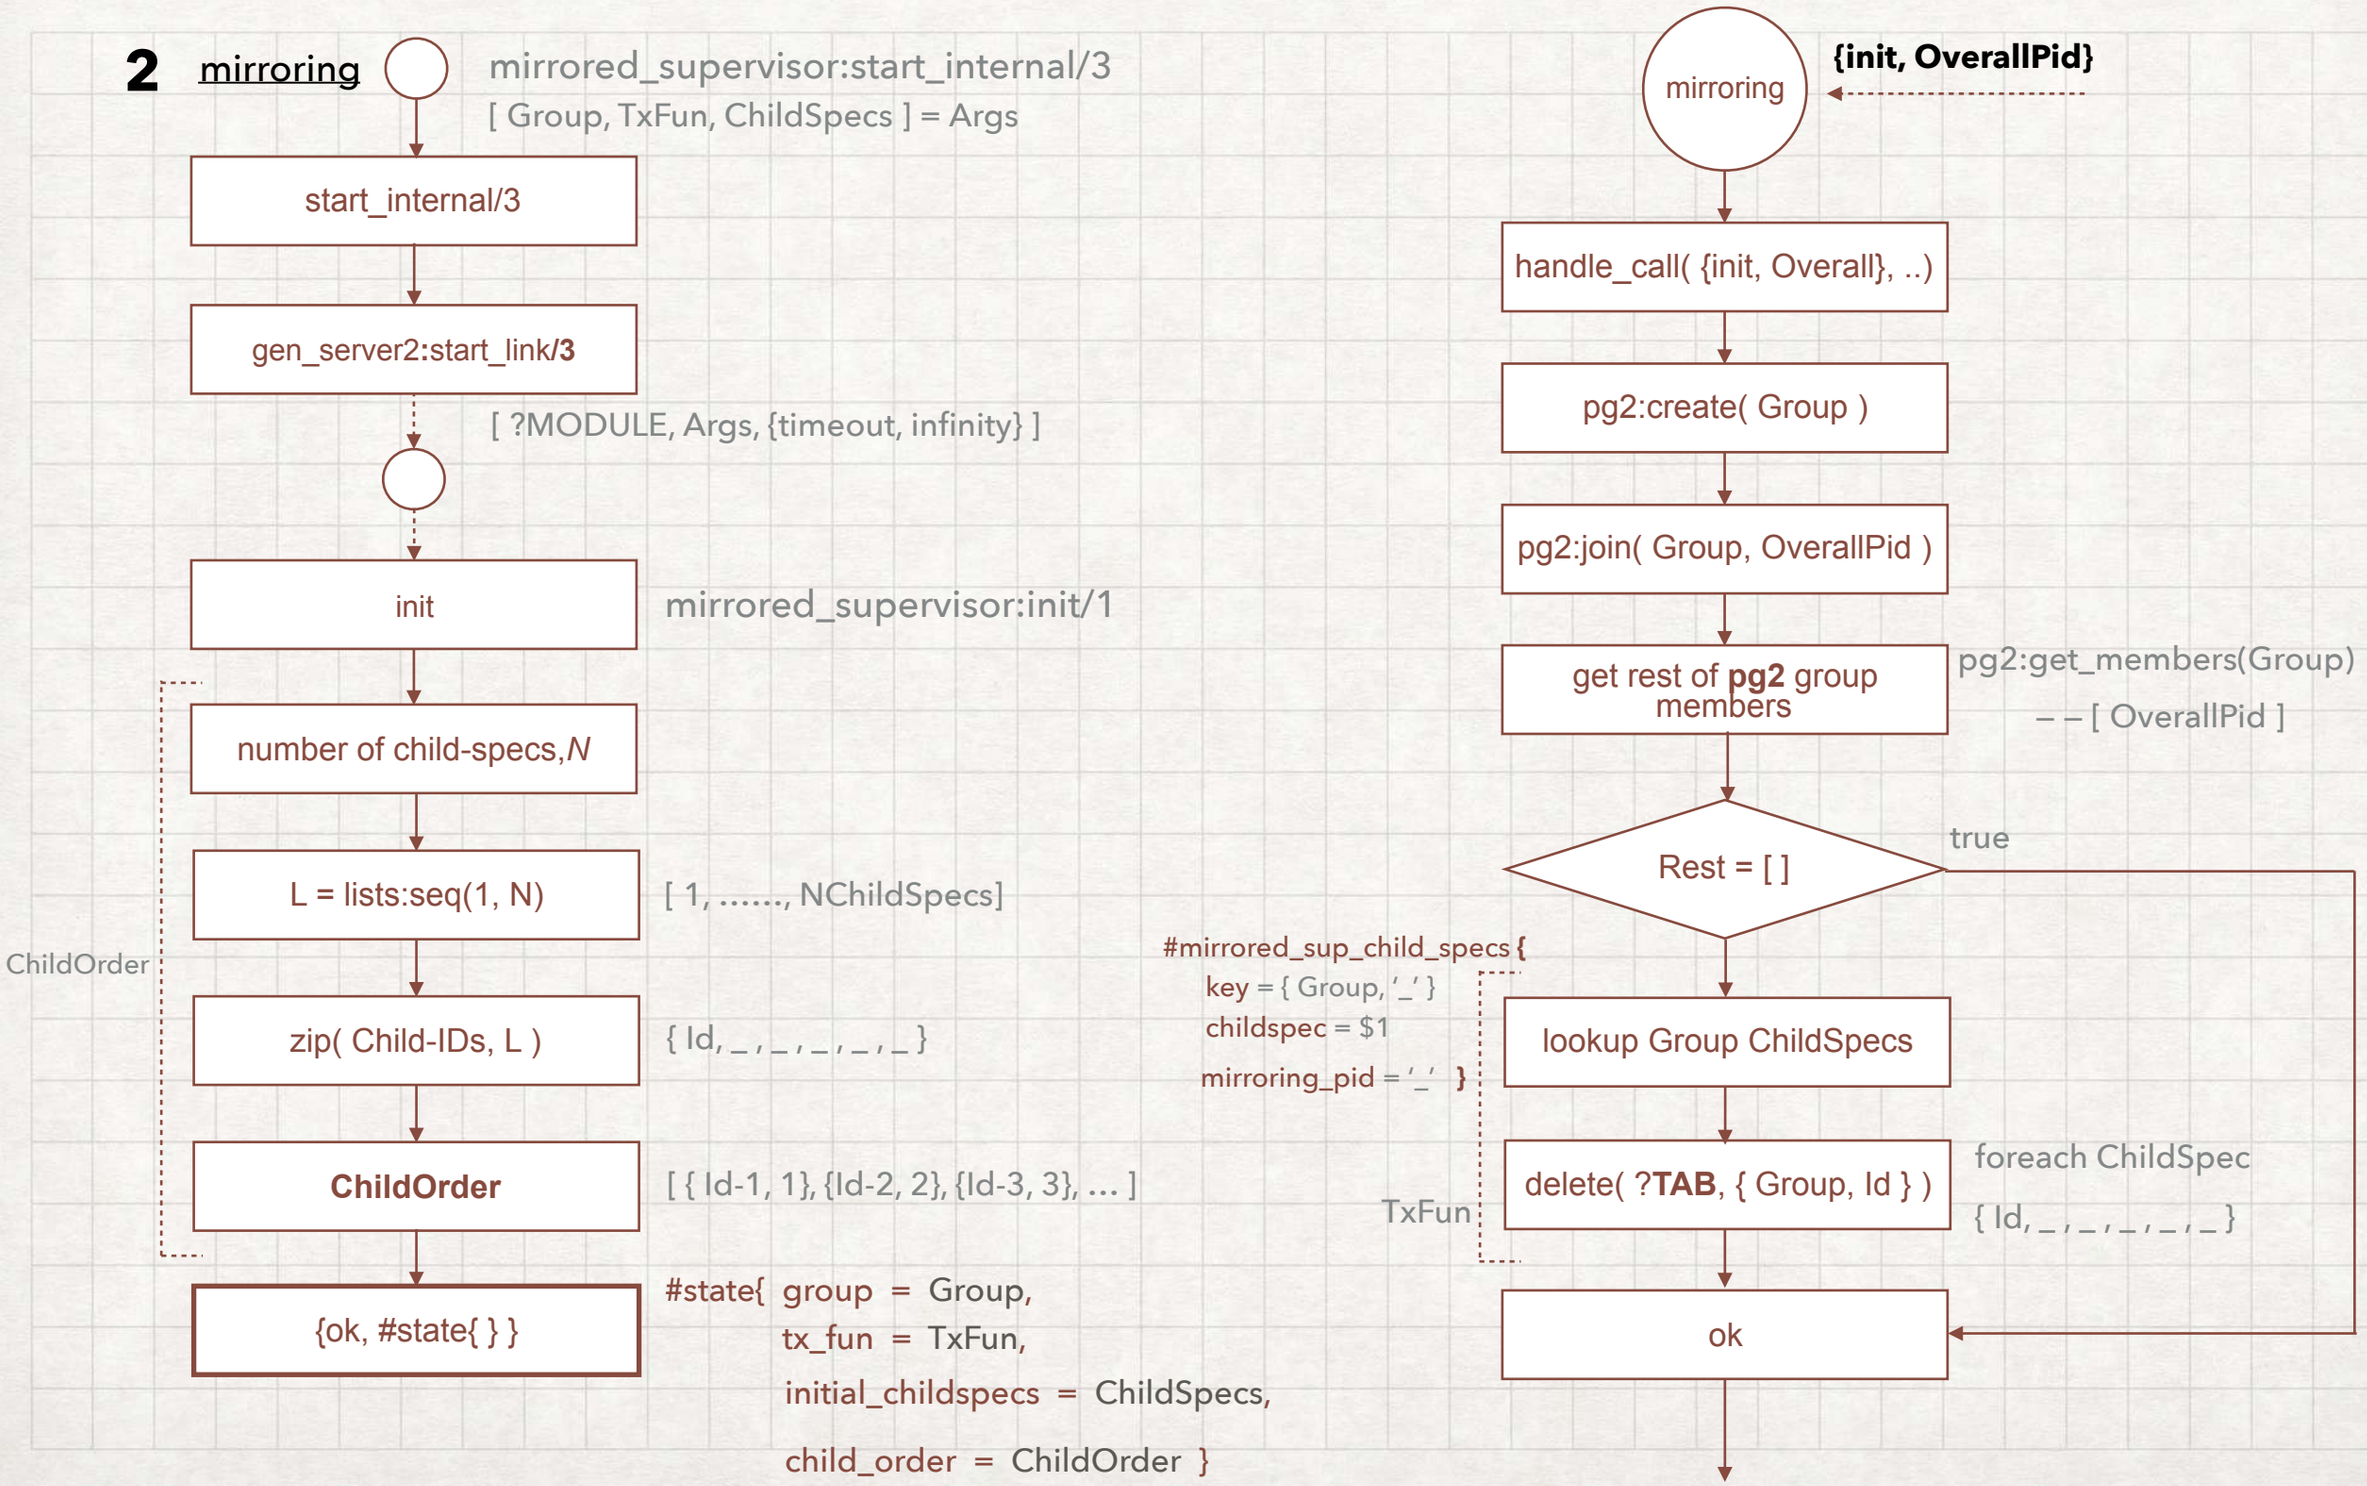

SUPERVISORS: MIRRORED SUPERVISOR

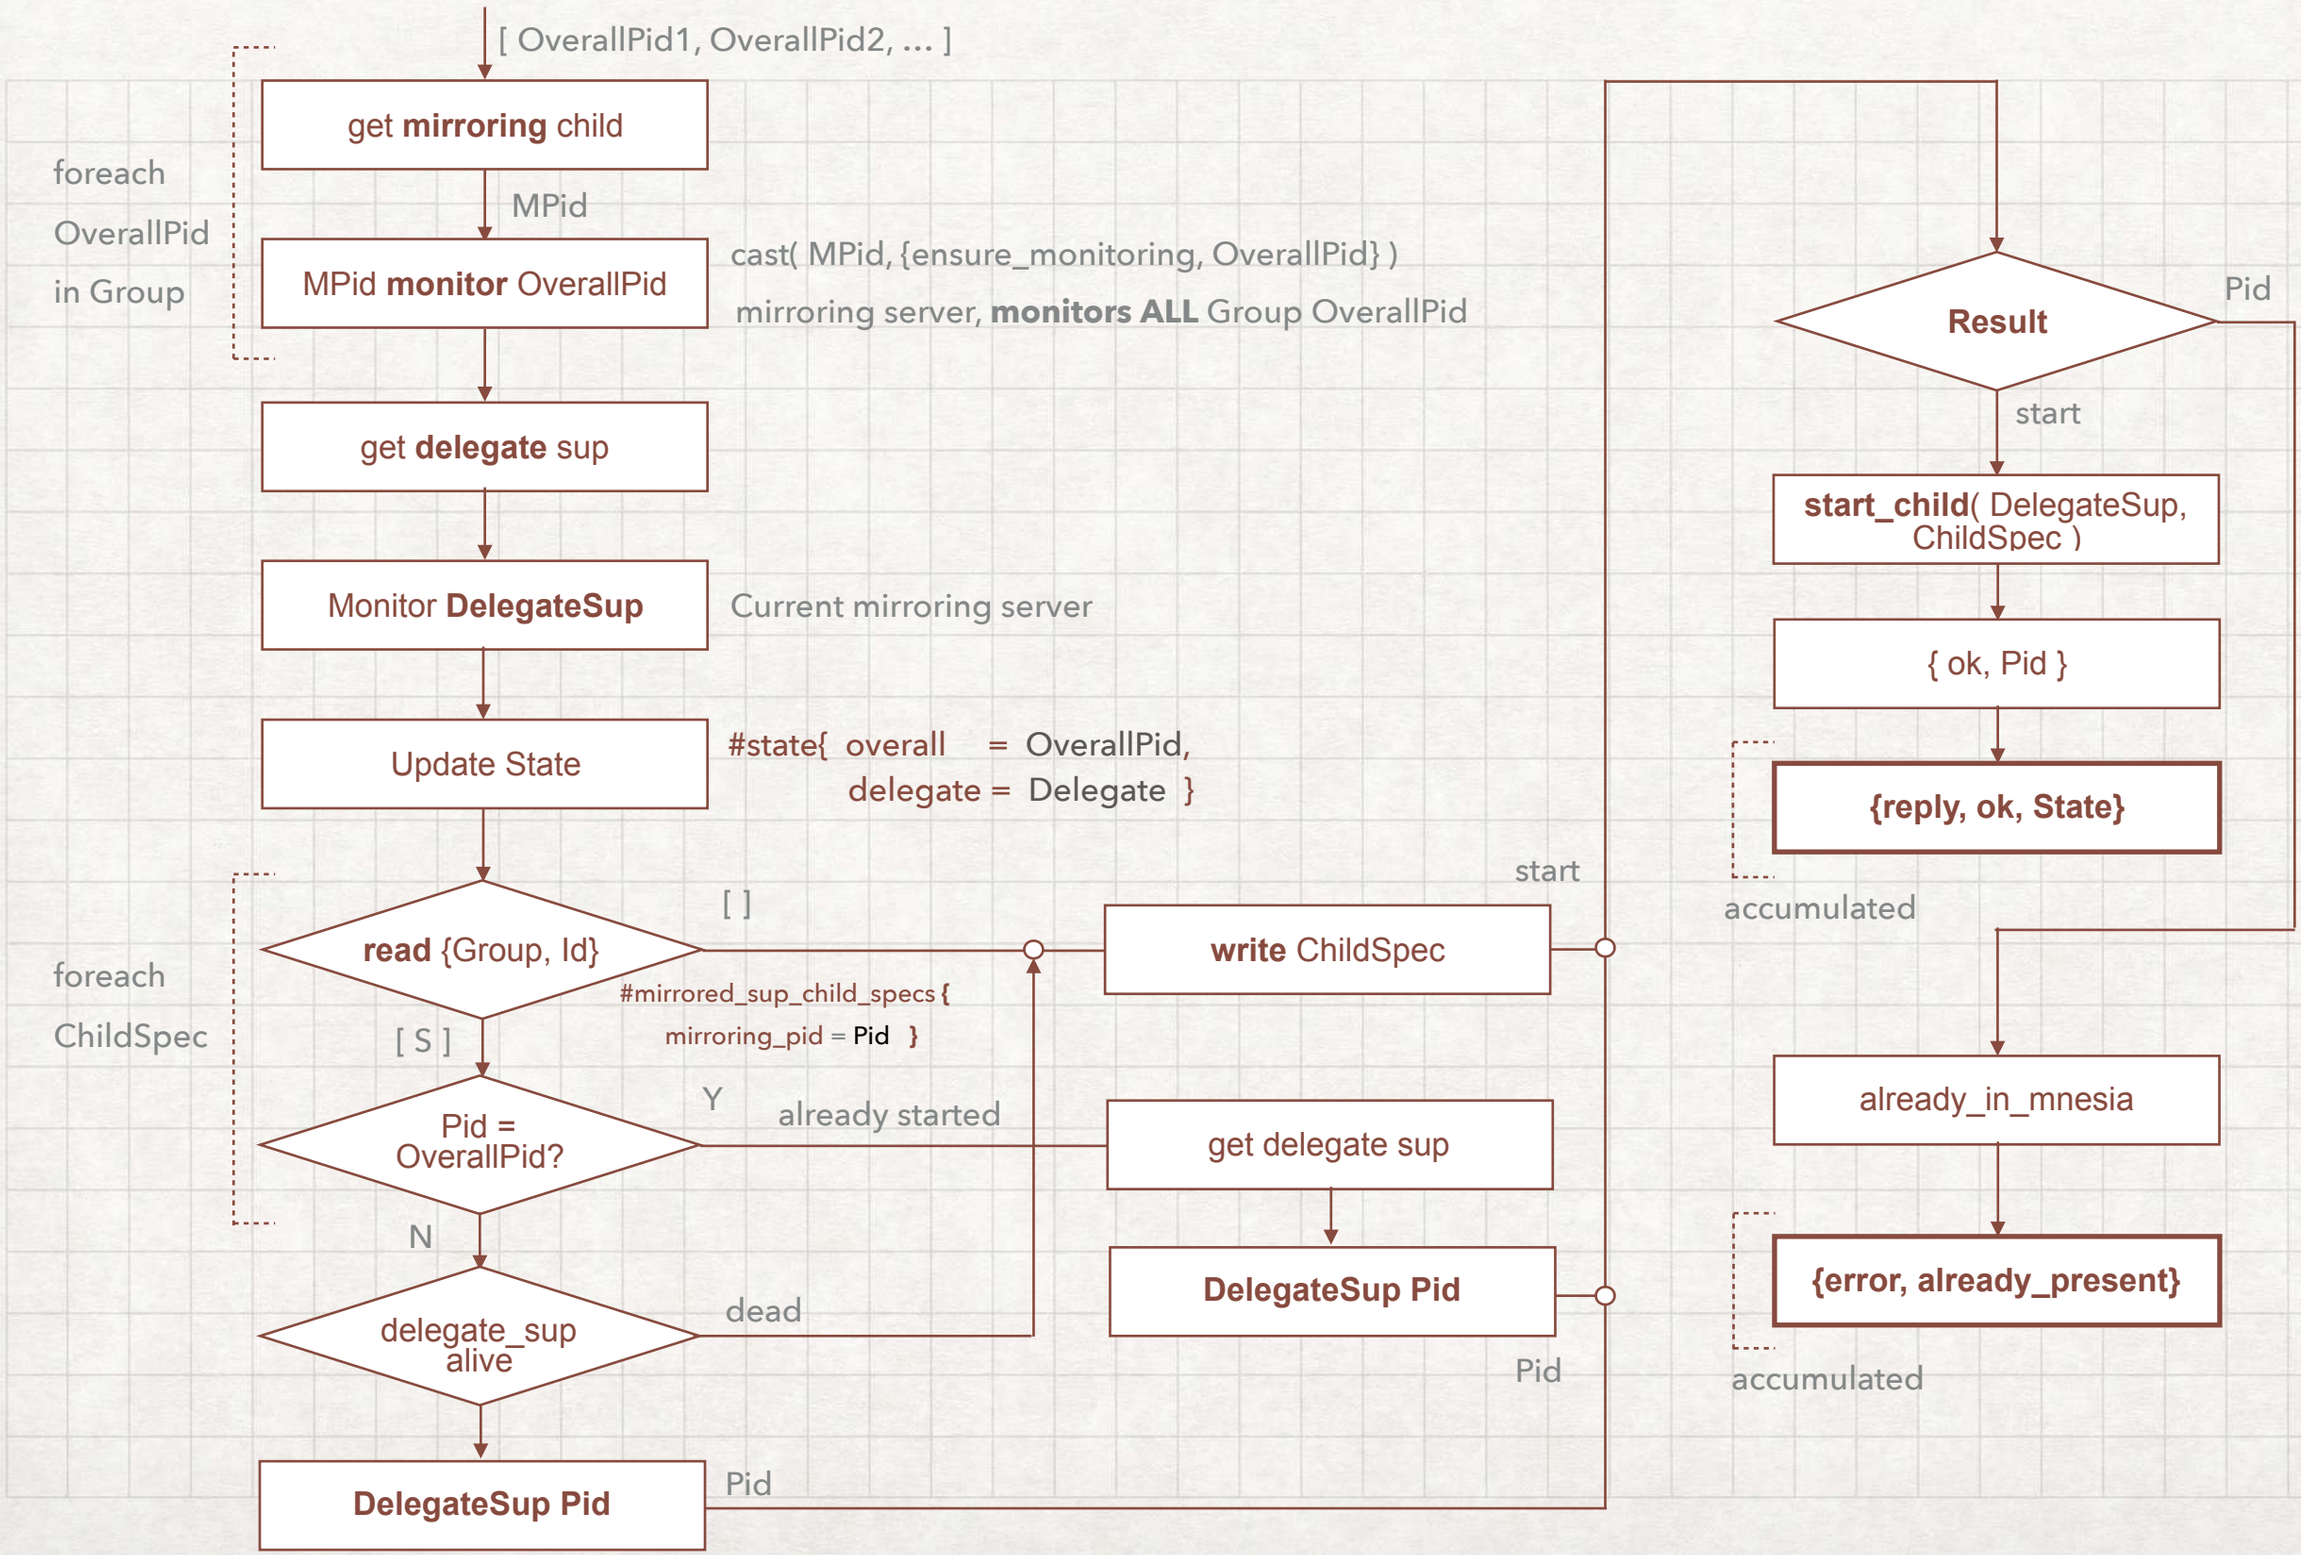

SUPERVISORS: MIRRORED SUPERVISOR

mnesia

```
#mirrored_sup_child_specs {
    key = { Group, Id }
    childspec = ChildSpec,
    mirroring_pid = Overall   }
```

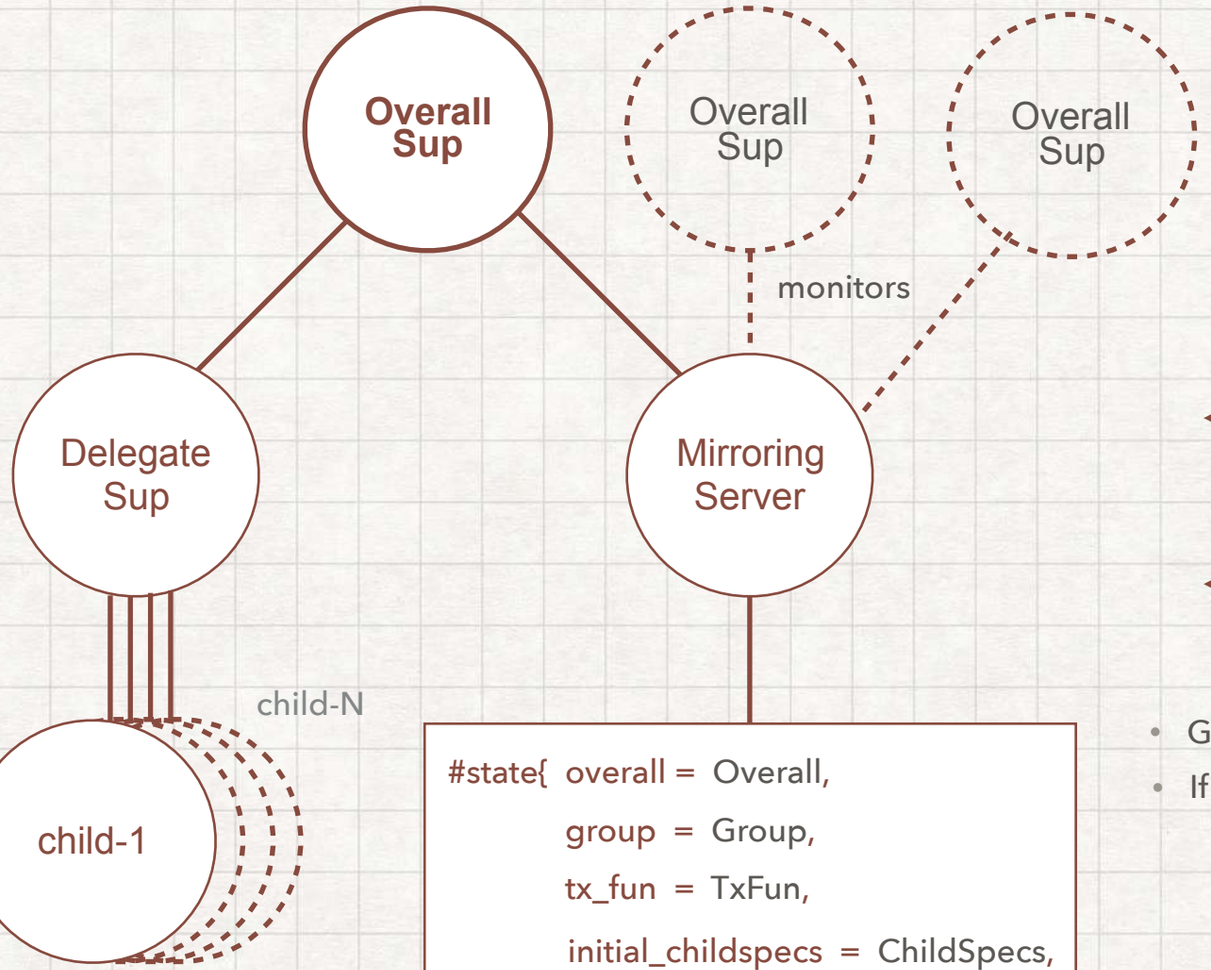


```
#state{  overall =  Overall,
         group  =  Group,
         tx_fun  =  TxFun,
         initial_childspecs  =  ChildSpecs,
         child_order  =  ChildOrder  }
```

start_child/2

delete_child/2

restart_child/2

terminate_child/2

which_children/1

count_children/1

SUPERVISOR API CALLS

{'DOWN', Ref, process, Pid, Info}

- Get **NewOverallPID** from PG2
- If current **Overall** of mirror is head of list
  - Update all **OldOverall** ChildSpecs in MNESIA, with **NewOverallPID**
  - Restore child-ordering
  - Start child specs

SUPERVISORS: MIRRORED SUPERVISOR

- Federation link top-level supervisors

```
17  -module(rabbit_federation_exchange_link_sup_sup).
18
19  -behaviour(mirrored_supervisor).
20
21  -include_lib("rabbit_common/include/rabbit.hrl").
22  -define(SUPERVISOR, ?MODULE).
23
24  %% Supervises the upstream links for all exchanges (but not queues). We need
25  %% different handling here since exchanges want a mirrored sup.
26
27  -export([start_link/0, start_child/1, adjust/1, stop_child/1]).
28  -export([init/1]).
29
30  %%----------------------------------------------------------------------------
31
32  start_link() ->
33      mirrored_supervisor:start_link({local, ?SUPERVISOR}, ?SUPERVISOR,
34                                     fun rabbit_misc:execute_mnesia_transaction/1,
35                                     ?MODULE, []).
```

PROCESSES & BEHAVIOURS

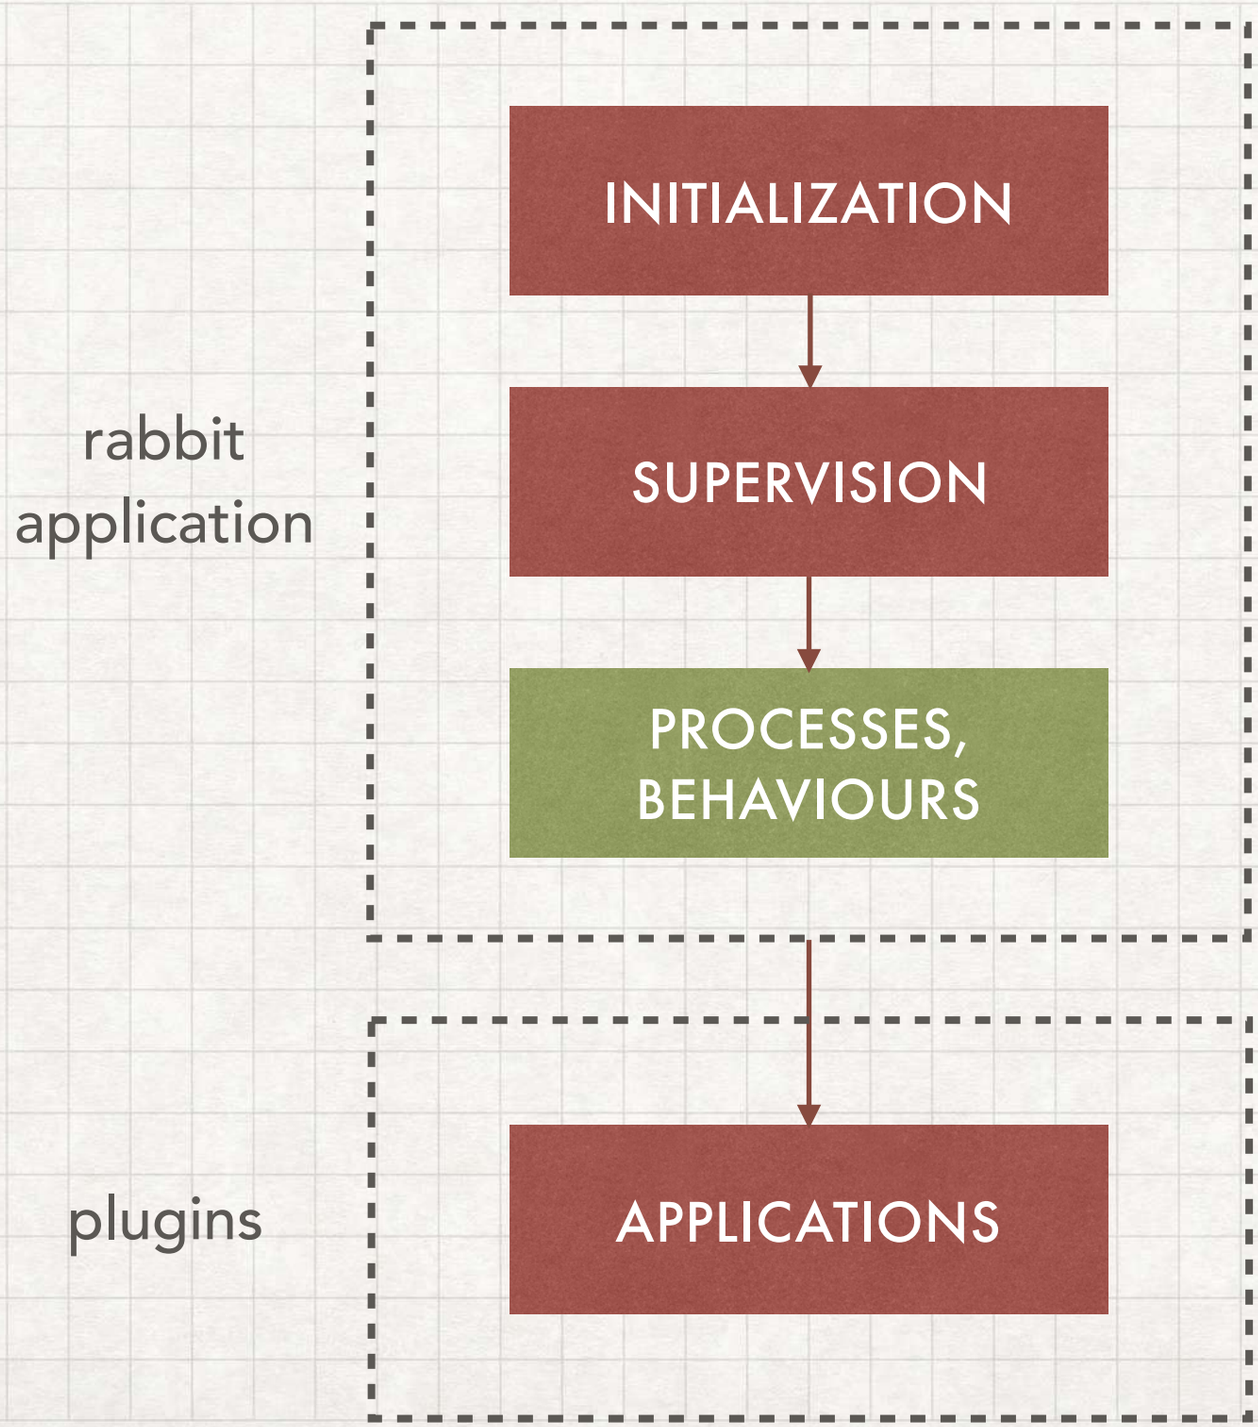

PROCESSES & BEHAVIOURS: GEN-SERVER-2

- Optimised selective receives - internal buffer extending (& draining) the message queue
- Additional callbacks - prioritised **call**, **casts** & **info**
- Pre- and Post- hibernation callbacks
- Backoff capabilities for delayed hibernation and variable timeouts
- Dynamic switching of callbacks (**become**)
- Debugging and formatting capabilities

PROCESSES & BEHAVIOURS: CREDIT FLOW

- Flow control on peer Erlang processes
- Lightweight - based on process dictionary
- Single control Erlang message (on demand)
- Simple, effective principle of operation
  - Sender **granted credit** by receiver, to send more
  - Sender **blocks** if it runs out of credit
  - Transceivers cannot grant more credit if blocked
- {InitialCredit, MoreCreditAfter}

PROCESSES & BEHAVIOURS: CREDIT FLOW

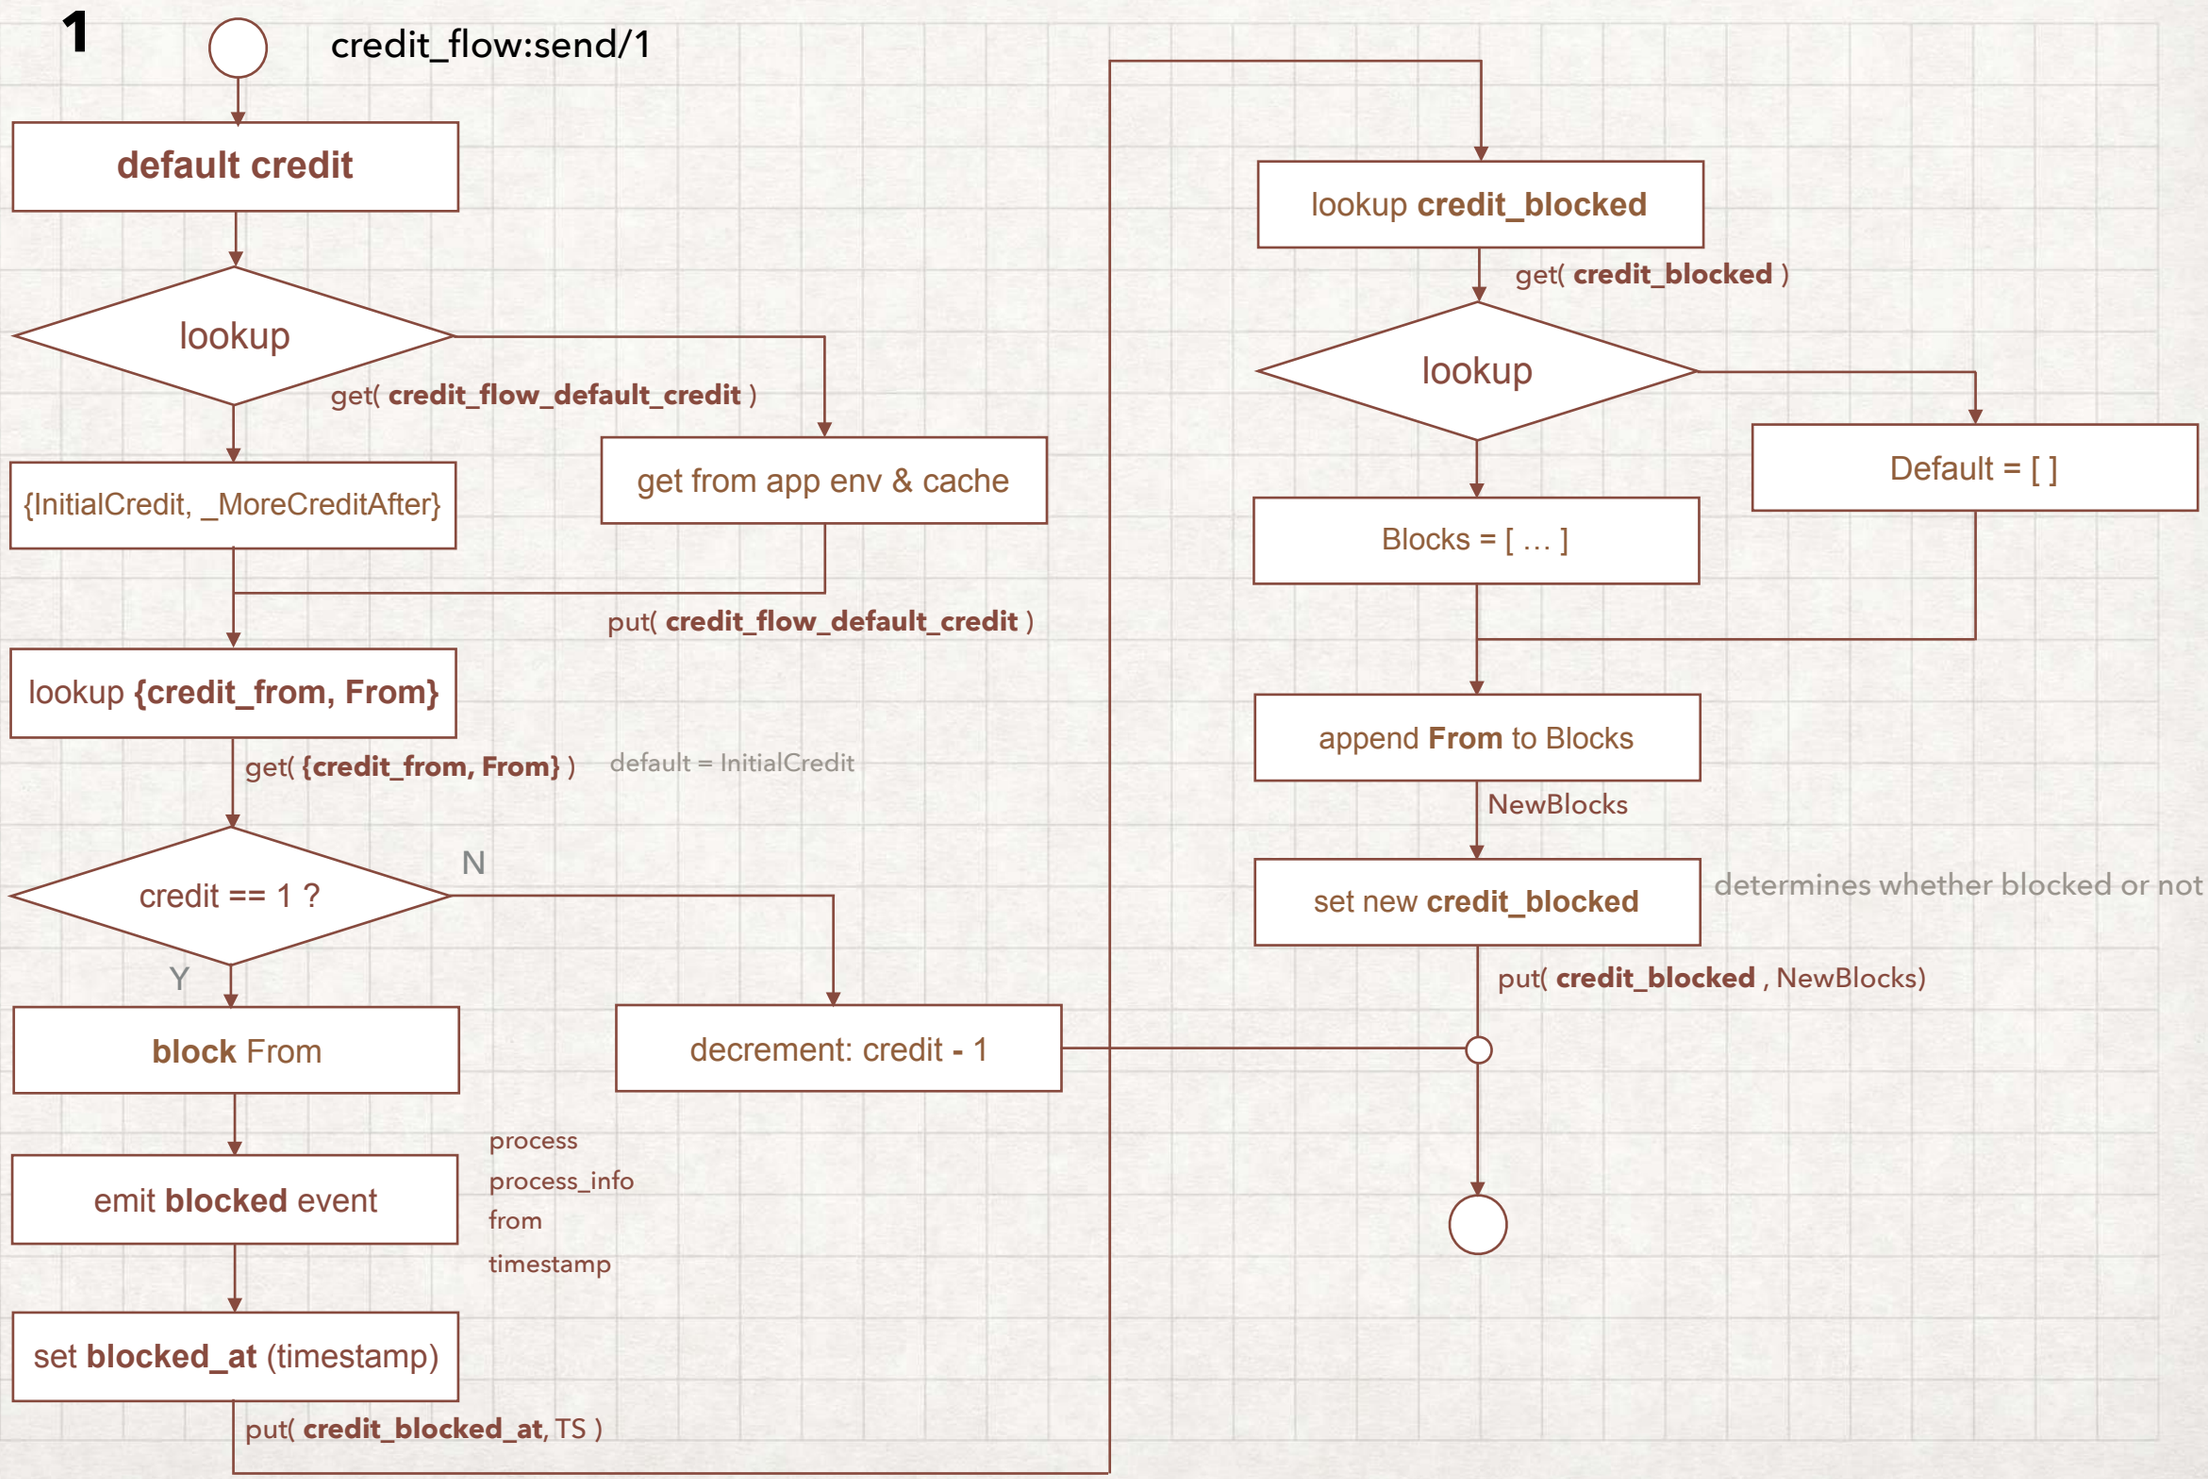

PROCESSES & BEHAVIOURS: CREDIT FLOW

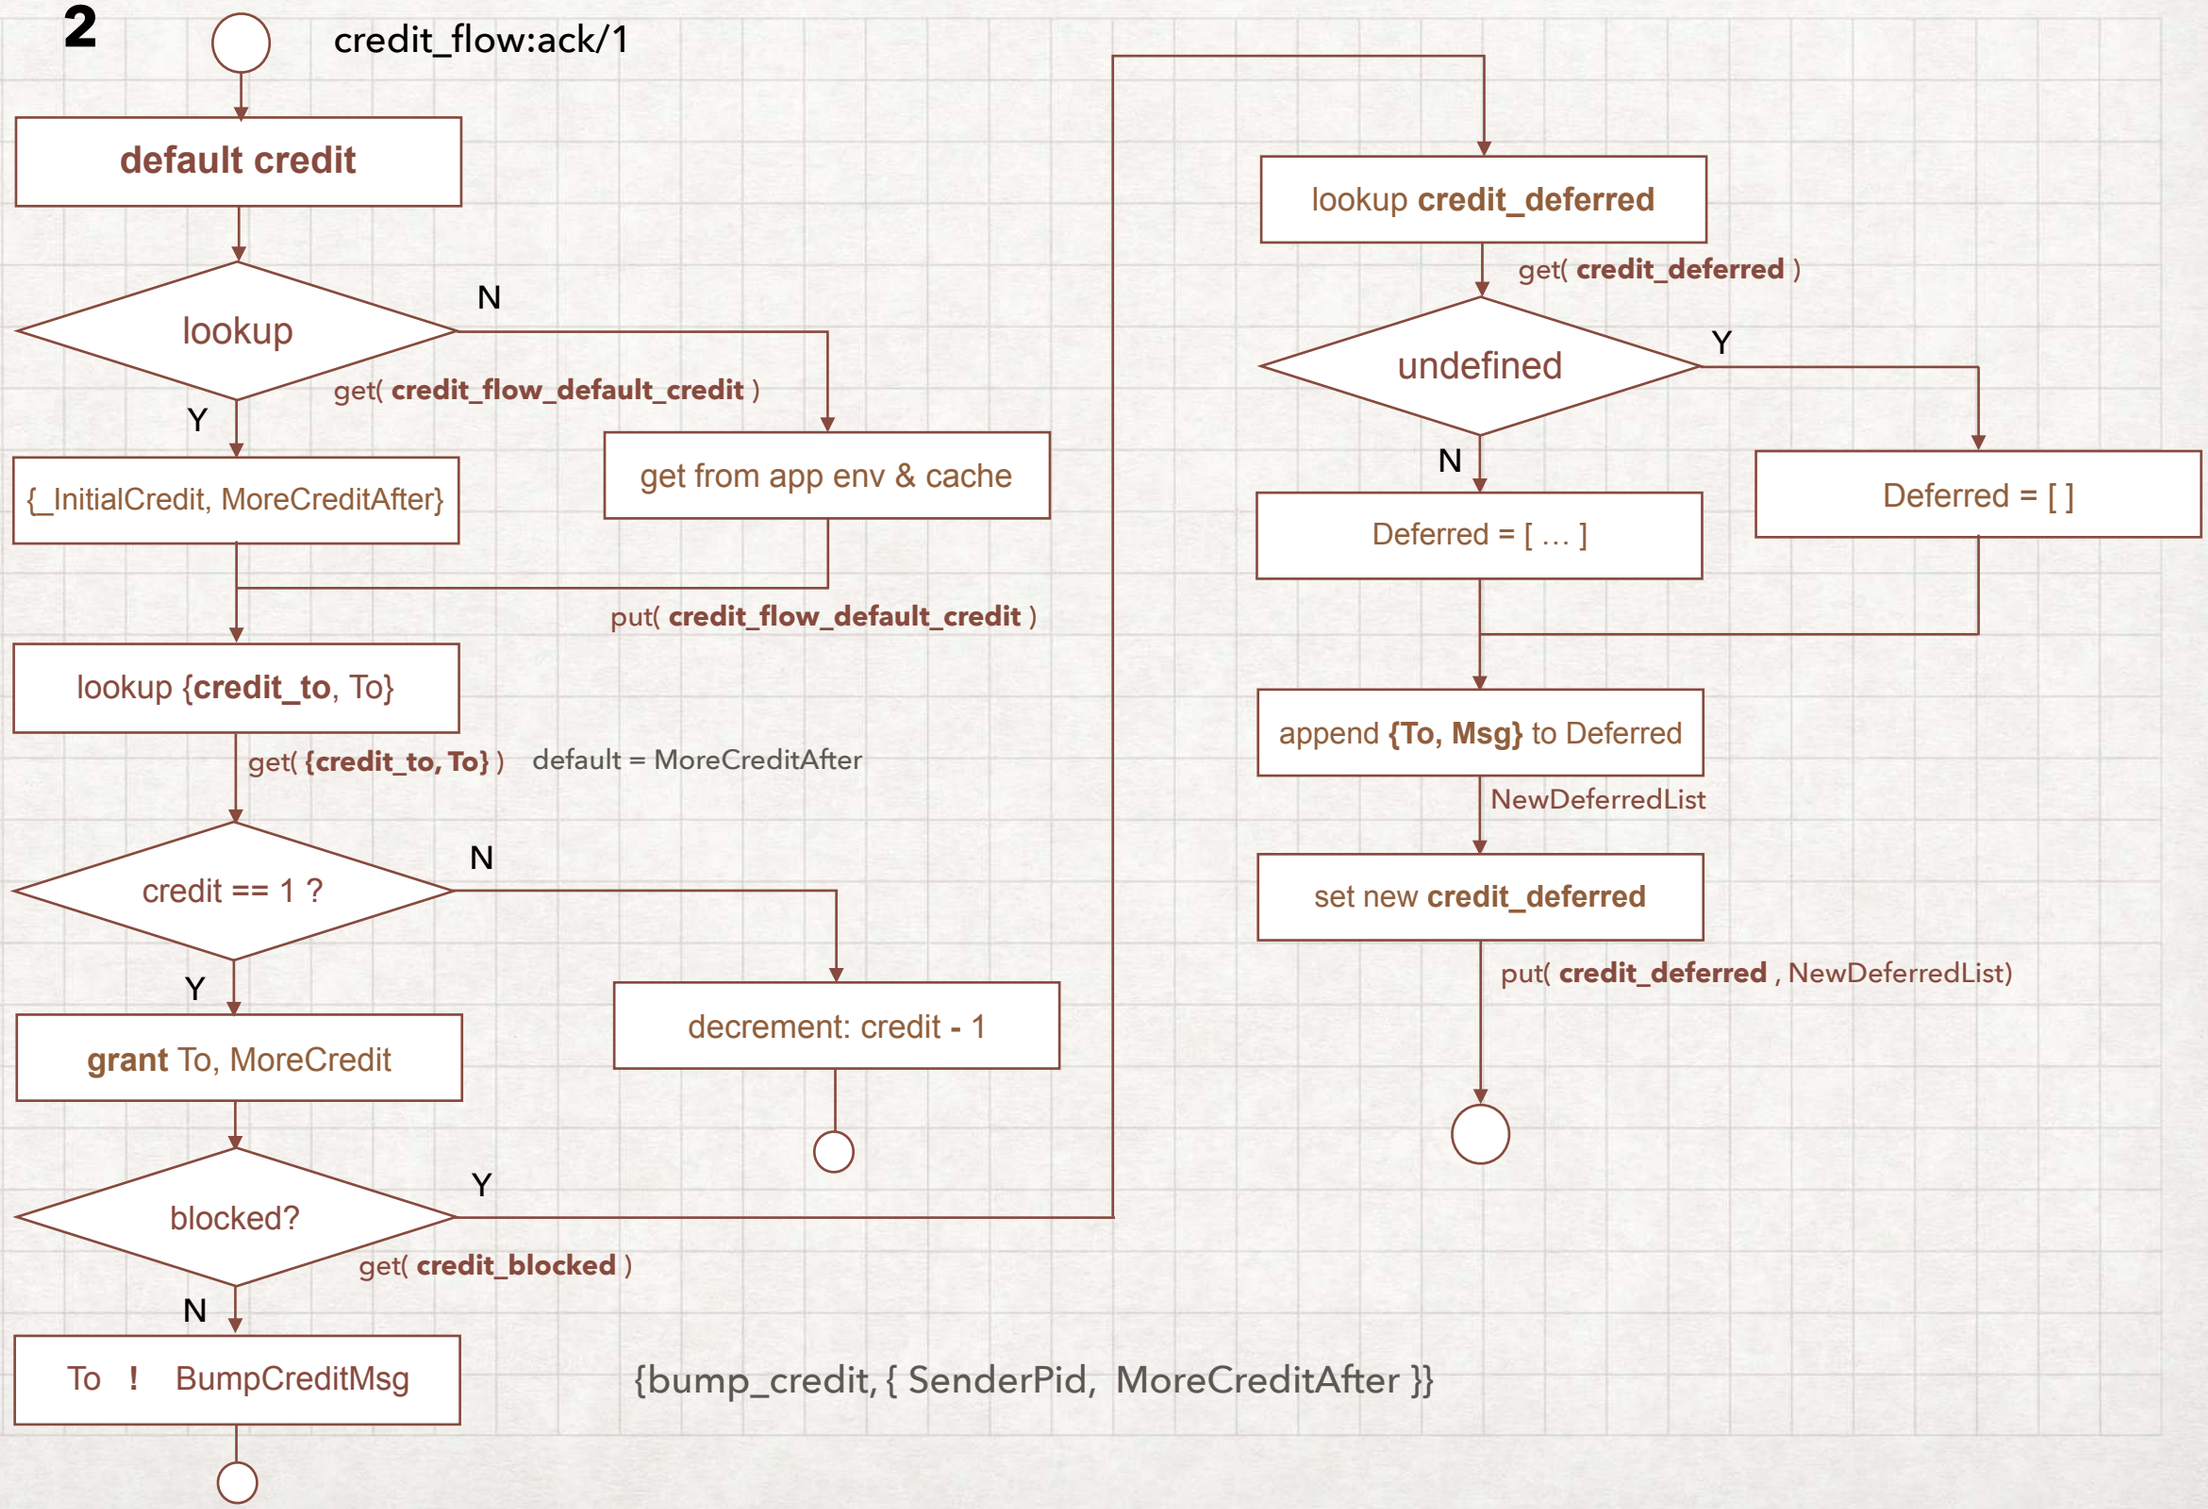

{bump_credit, { SenderPid, MoreCreditAfter }}

PROCESSES & BEHAVIOURS: CREDIT FLOW

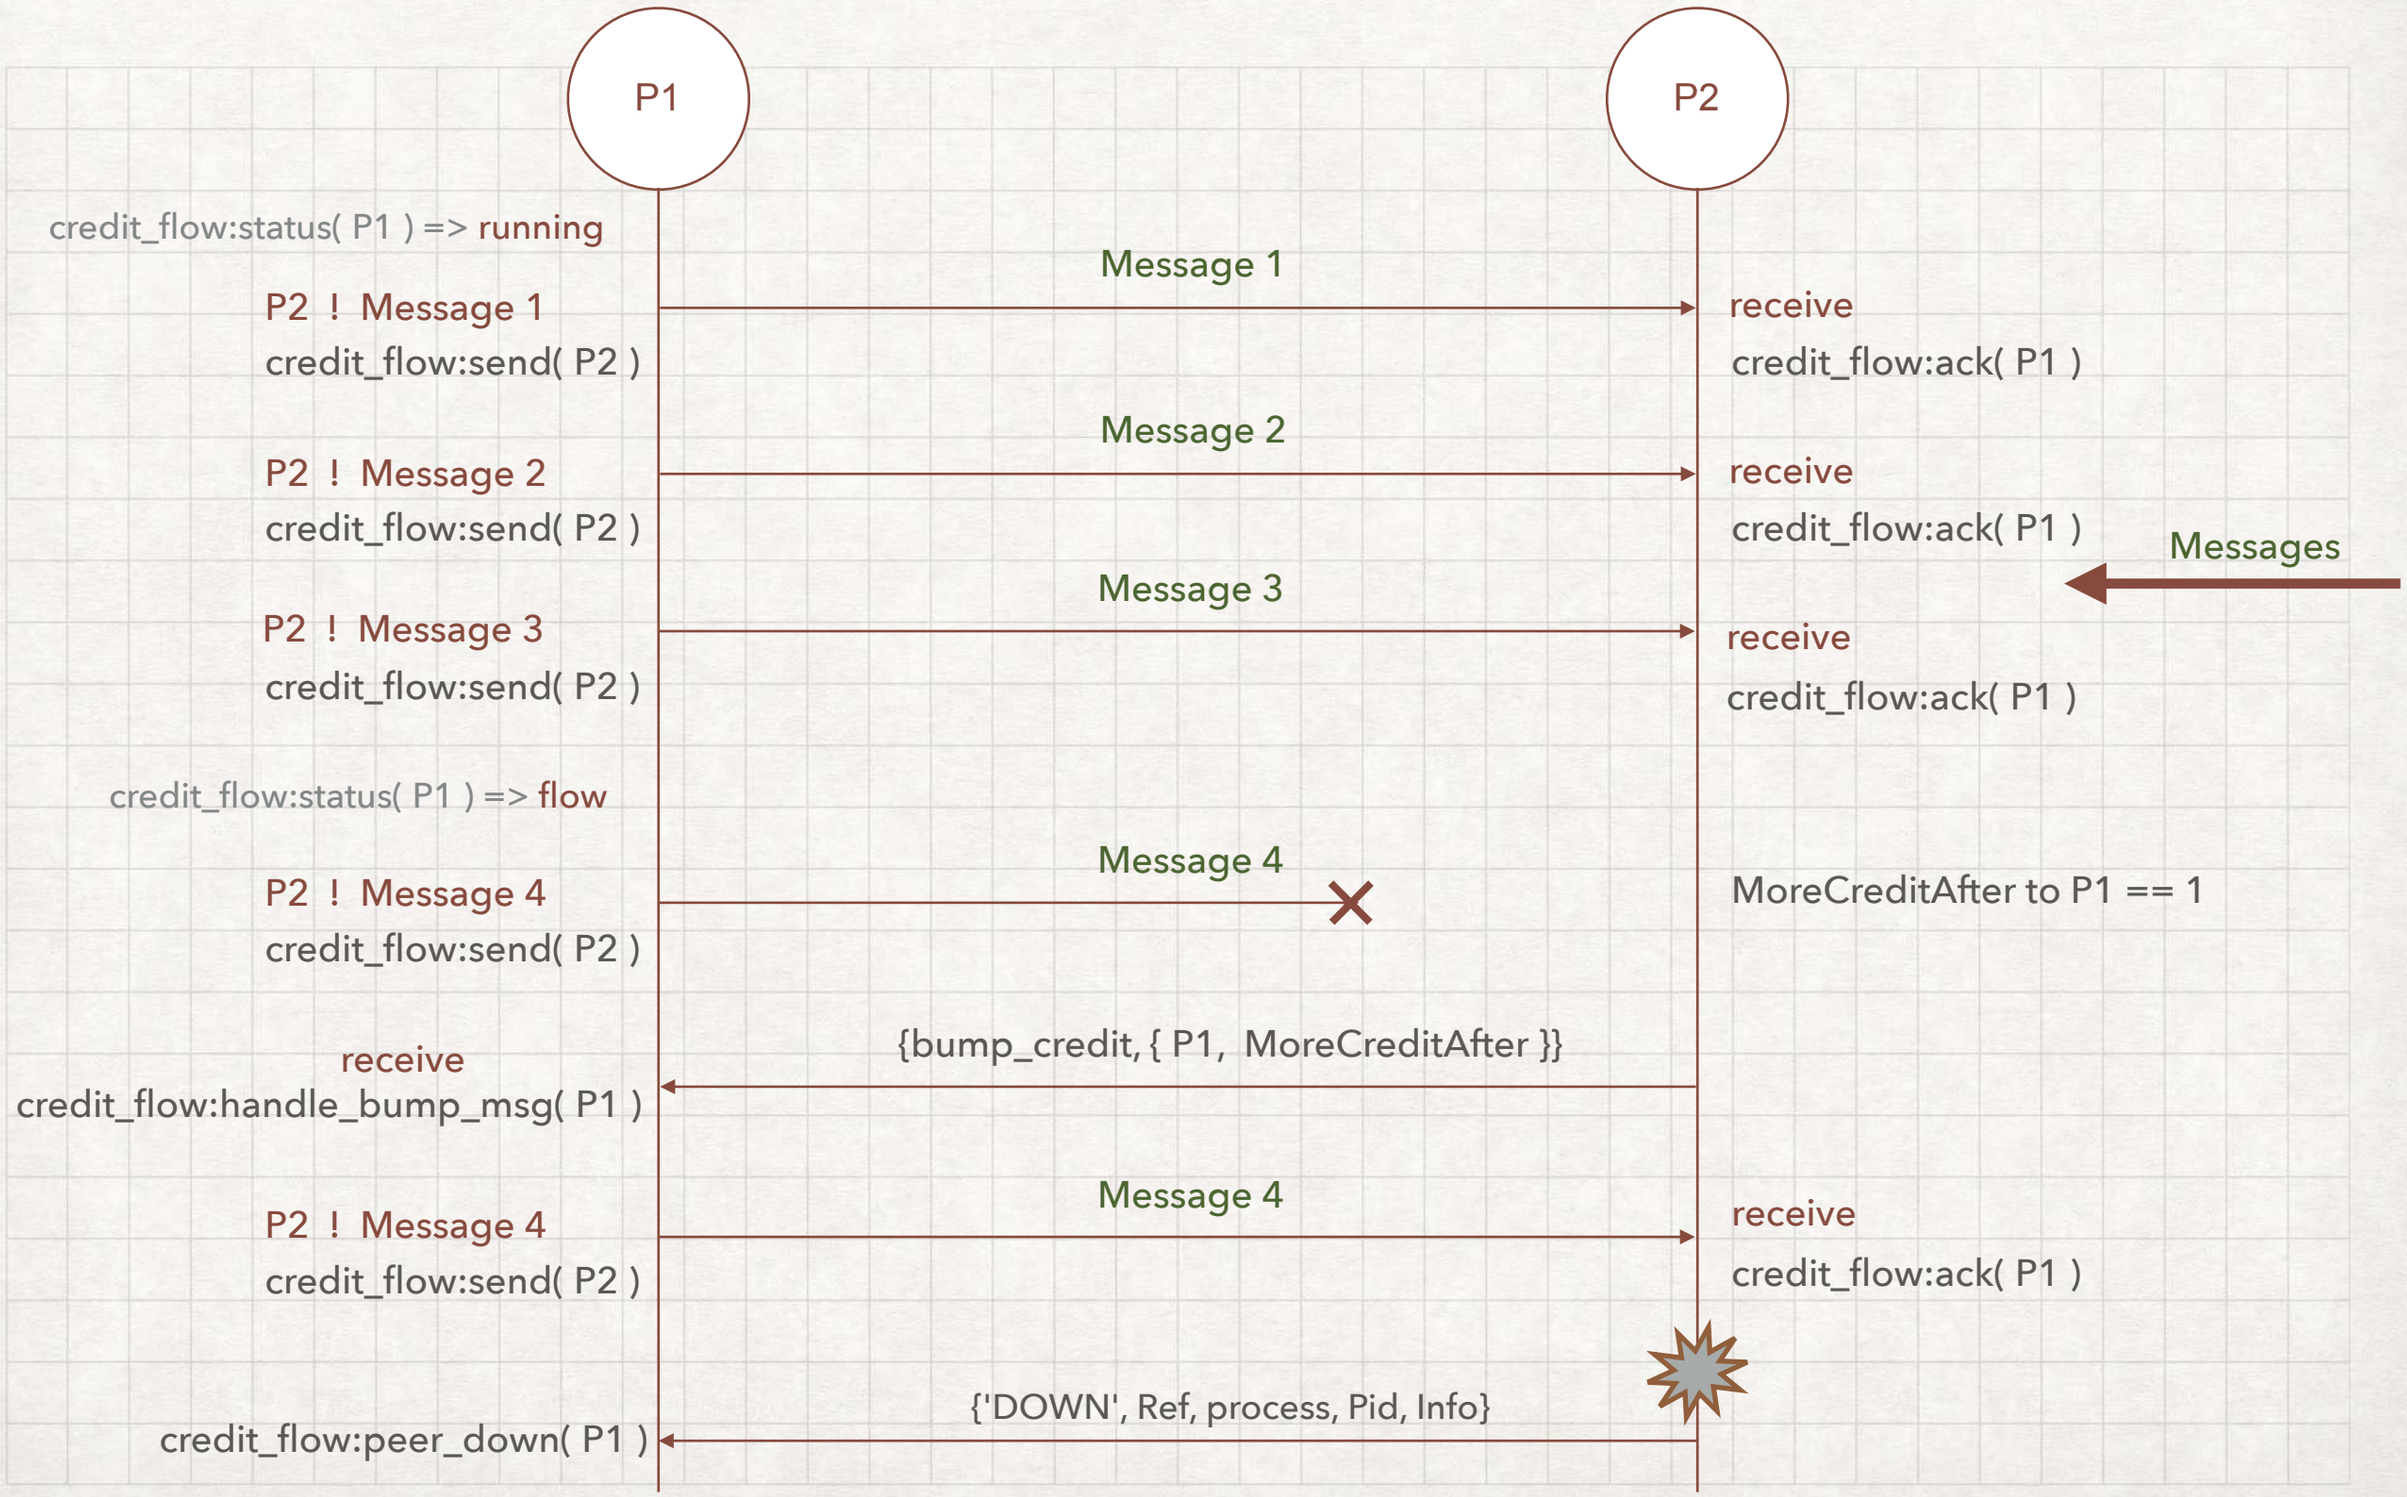

PROCESSES & BEHAVIOURS: CREDIT FLOW

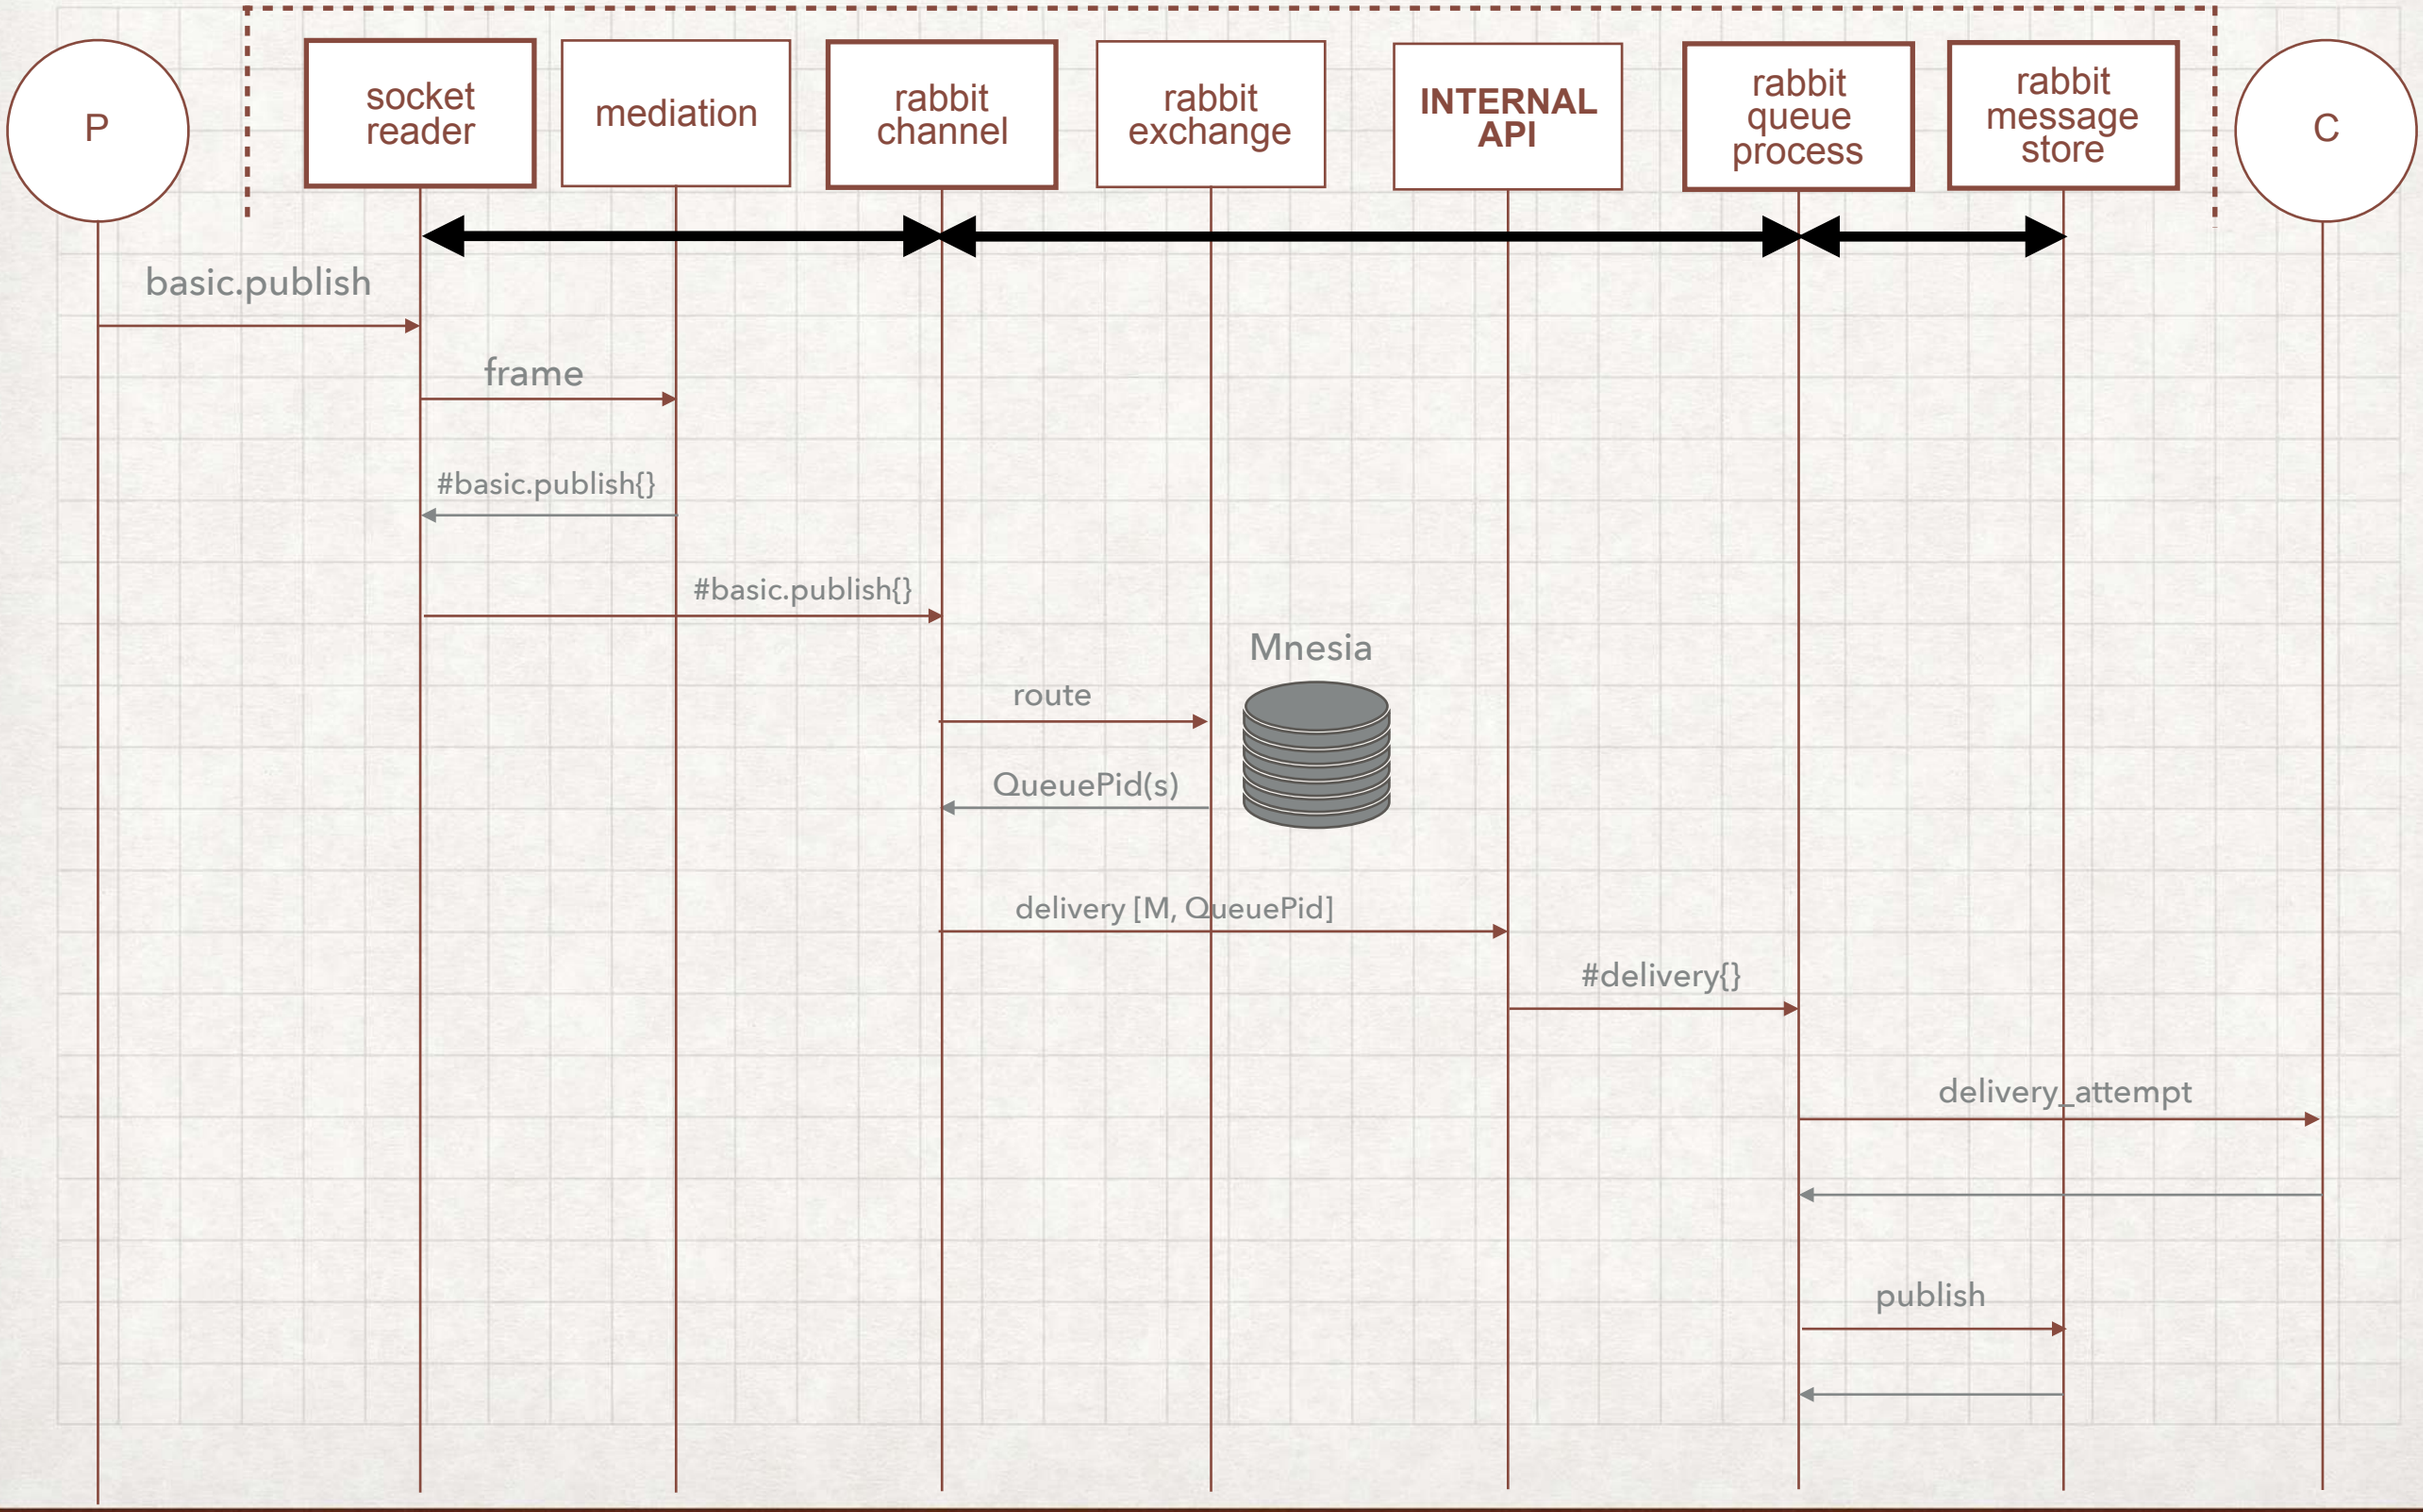

PROCESSES & BEHAVIOURS: DELEGATES

- Optmized internode communication
- Synchronous and Asynchronous operations
- Minimum blocking - configurable pool size
- Optmized process monitoring (on local node only)
- Low bandwidth usage on distribution links

PROCESSES & BEHAVIOURS: DELEGATES

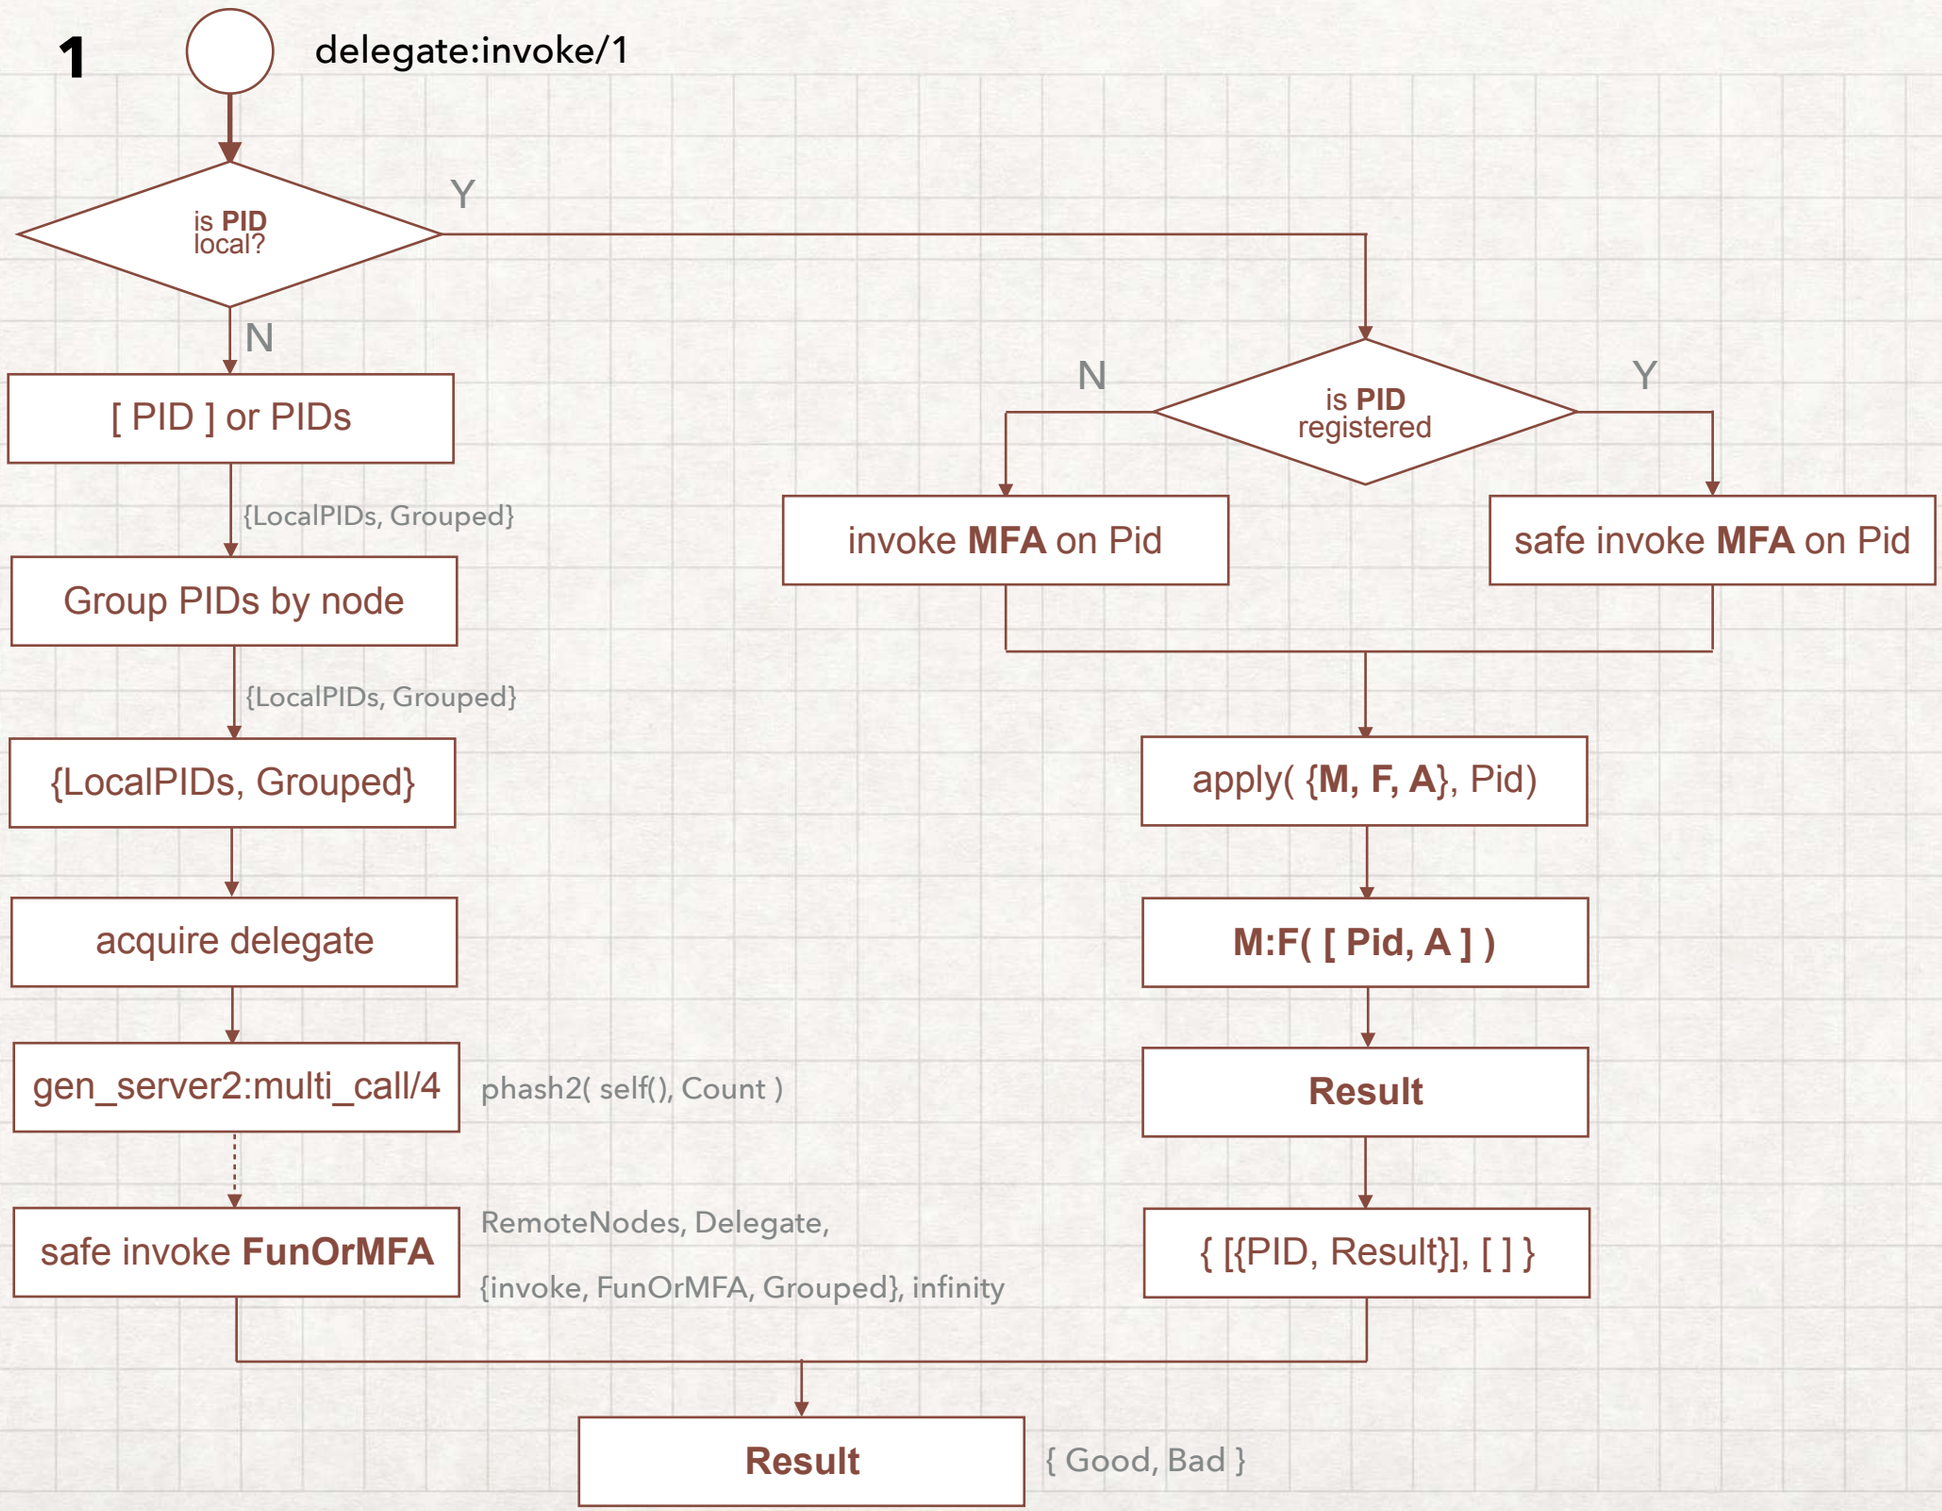

PROCESSES & BEHAVIOURS: DELEGATES

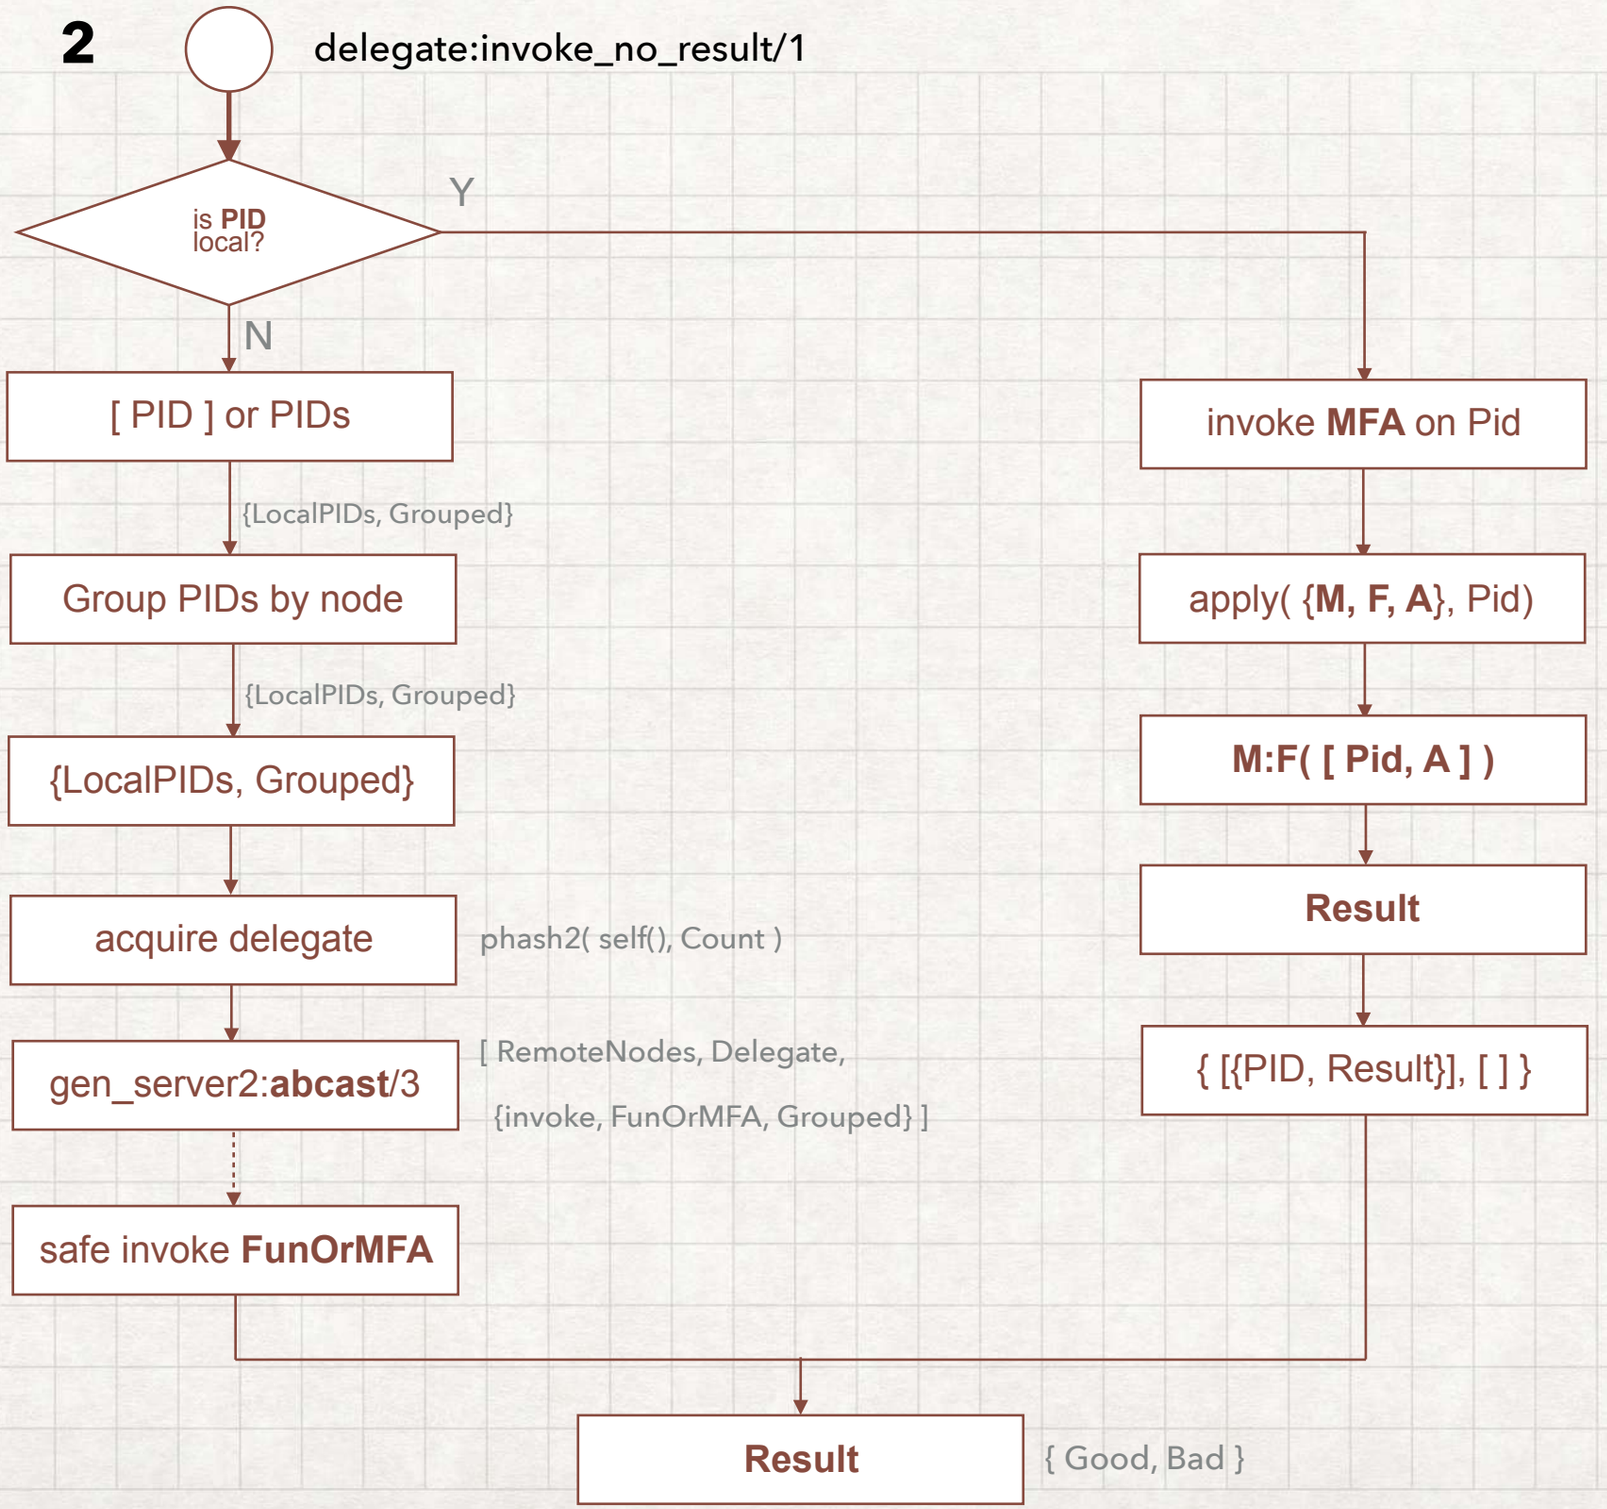

PROCESSES & BEHAVIOURS: DELEGATES

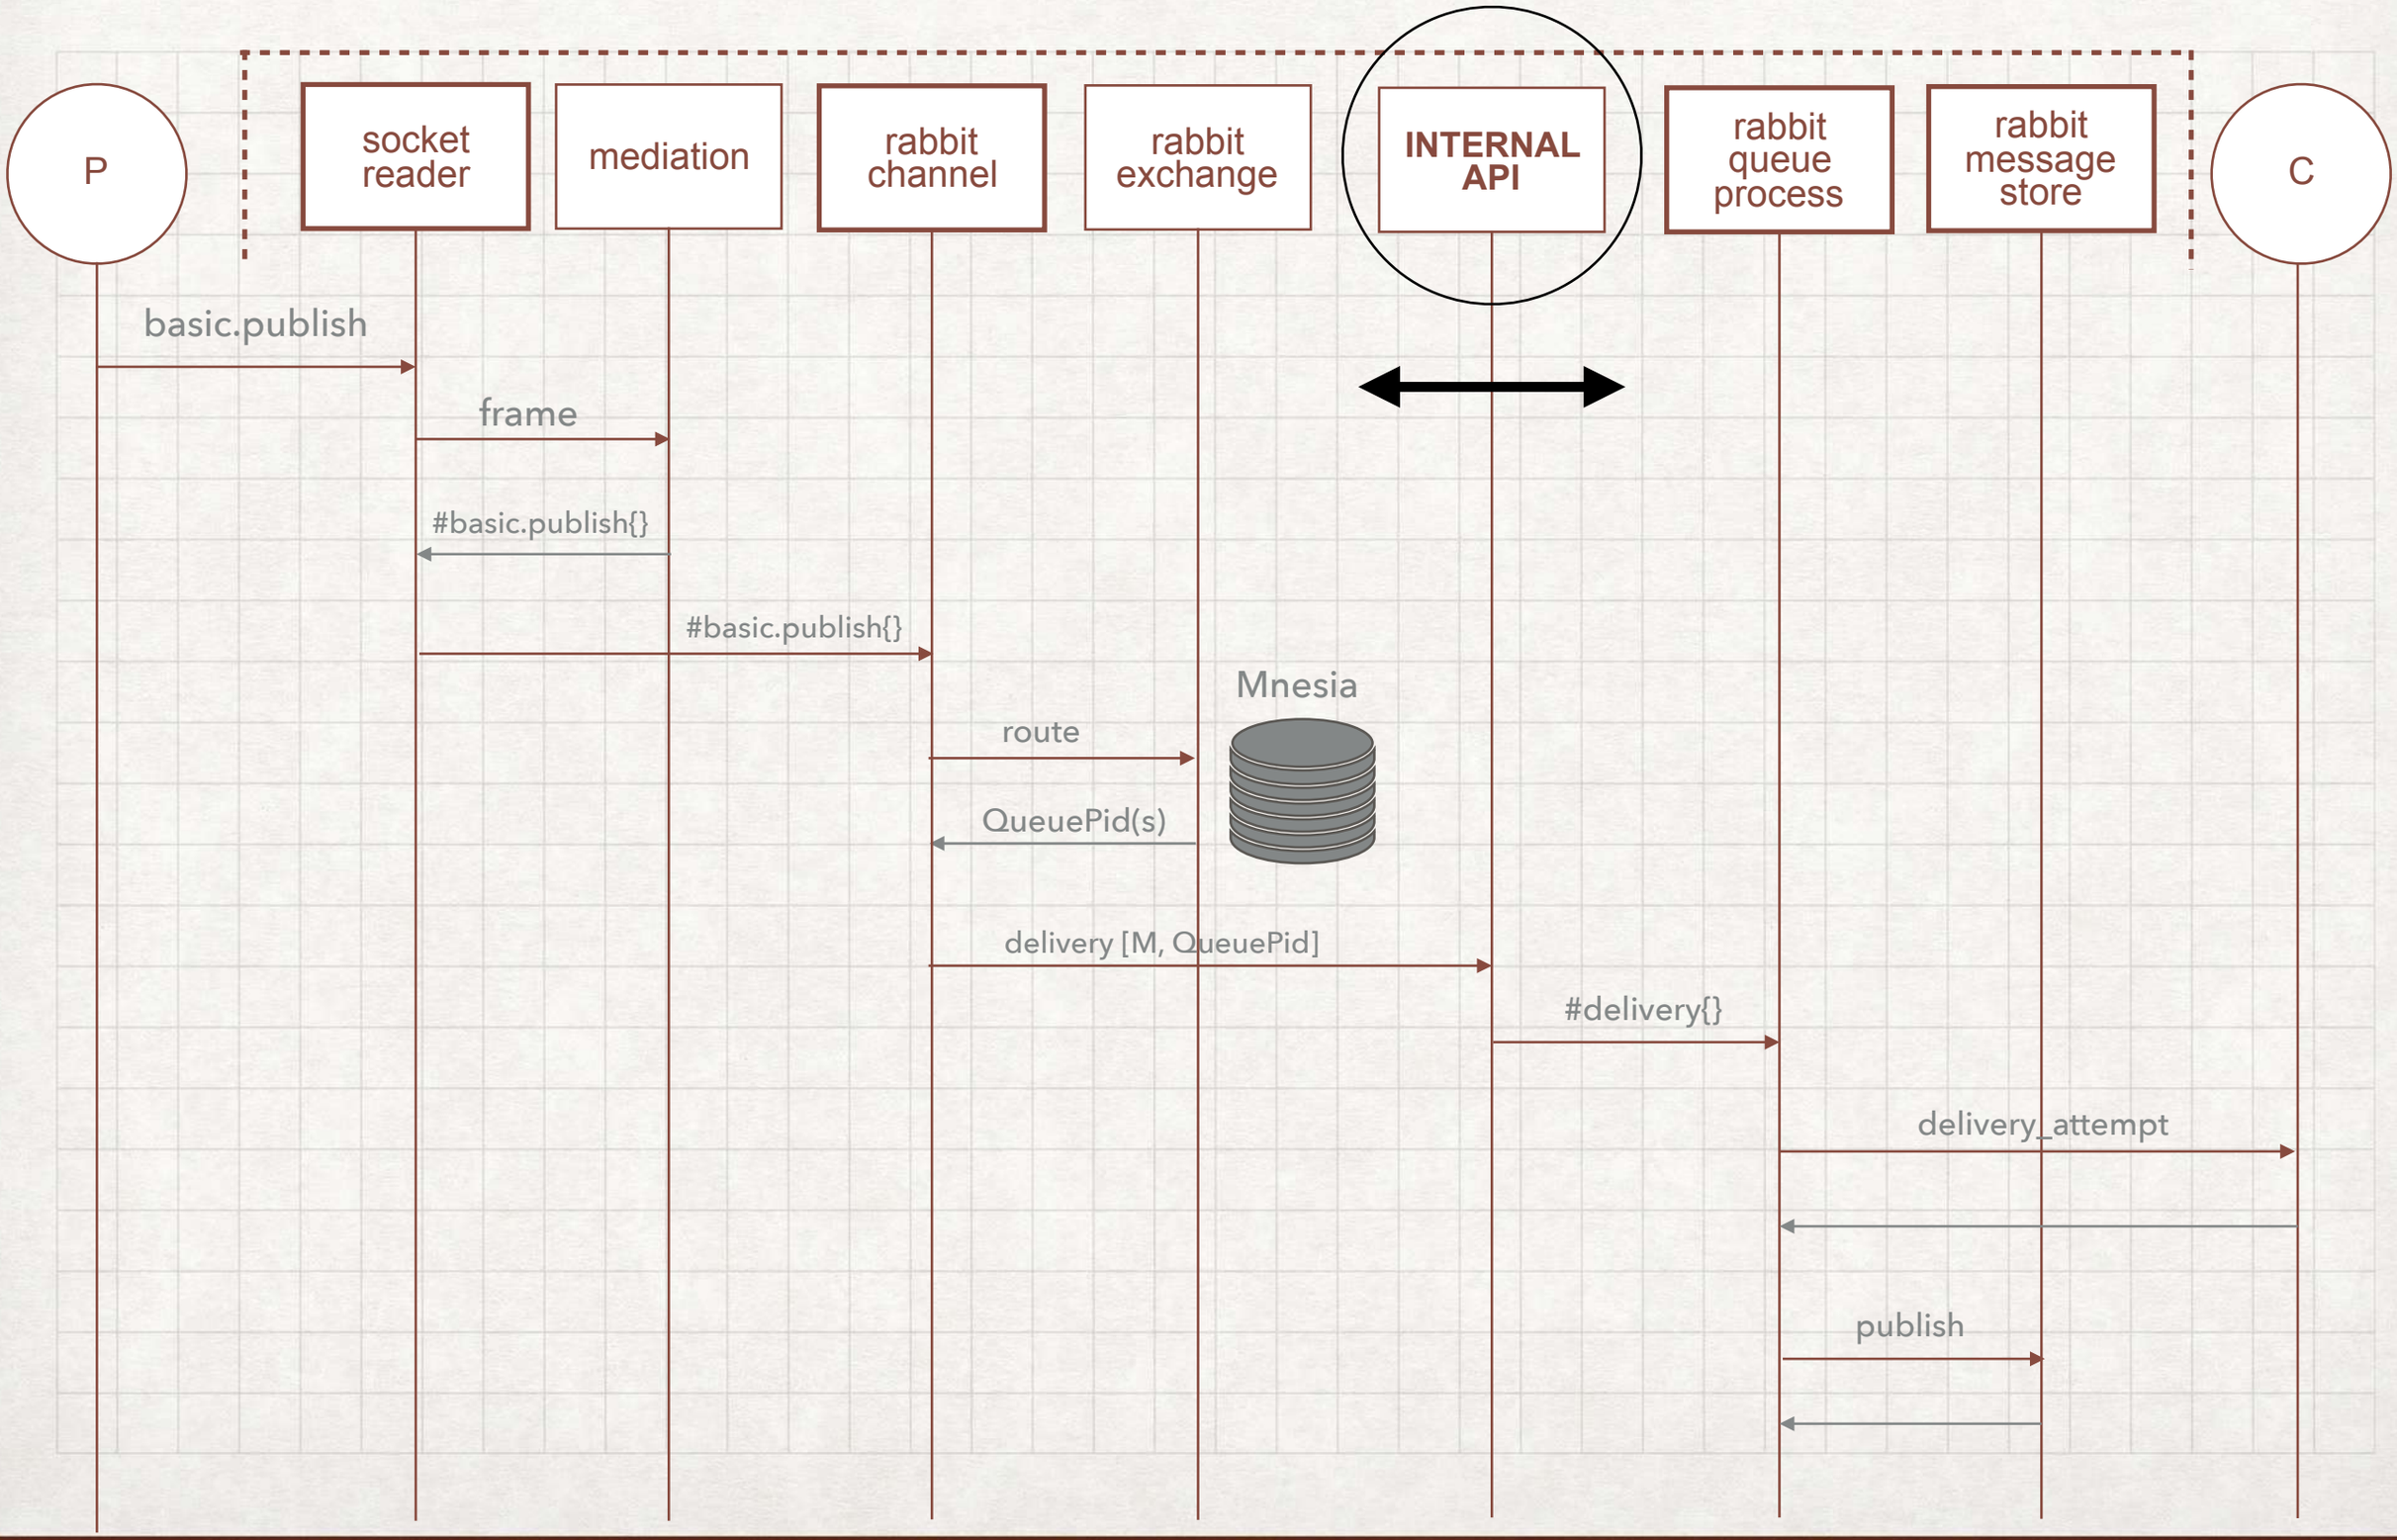

PROCESSES & BEHAVIOURS: DELEGATES

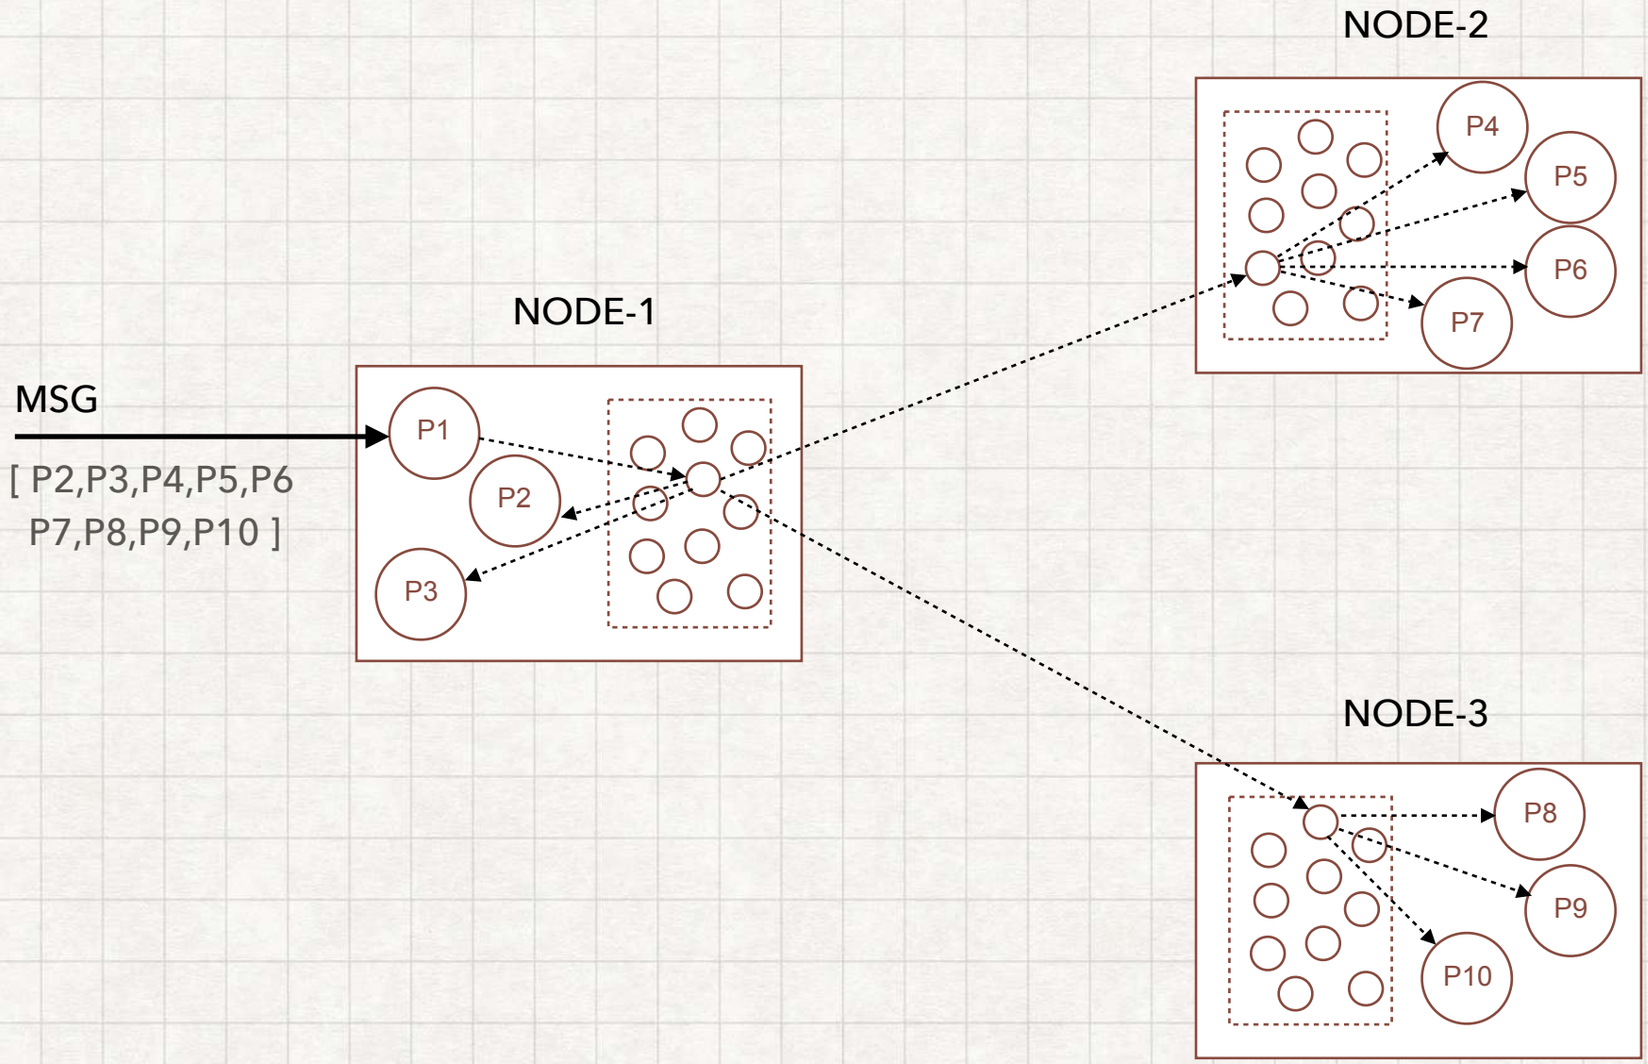

PROCESSES & BEHAVIOURS: MORE ...

- **Decorators**
  - Dynamic state updates of implementing processes
  - Facilitate OAM, CLI tools, e.g. live policy updates
- **GM (Guaranteed Multicast)**
  - Behaviour for attaining consensus on a group of processes
- **PMon**
  - Optimized monitors, at most 1 monitor per process
  - Querying capabilities, e.g. **is_monitored/1**

PLUGINS

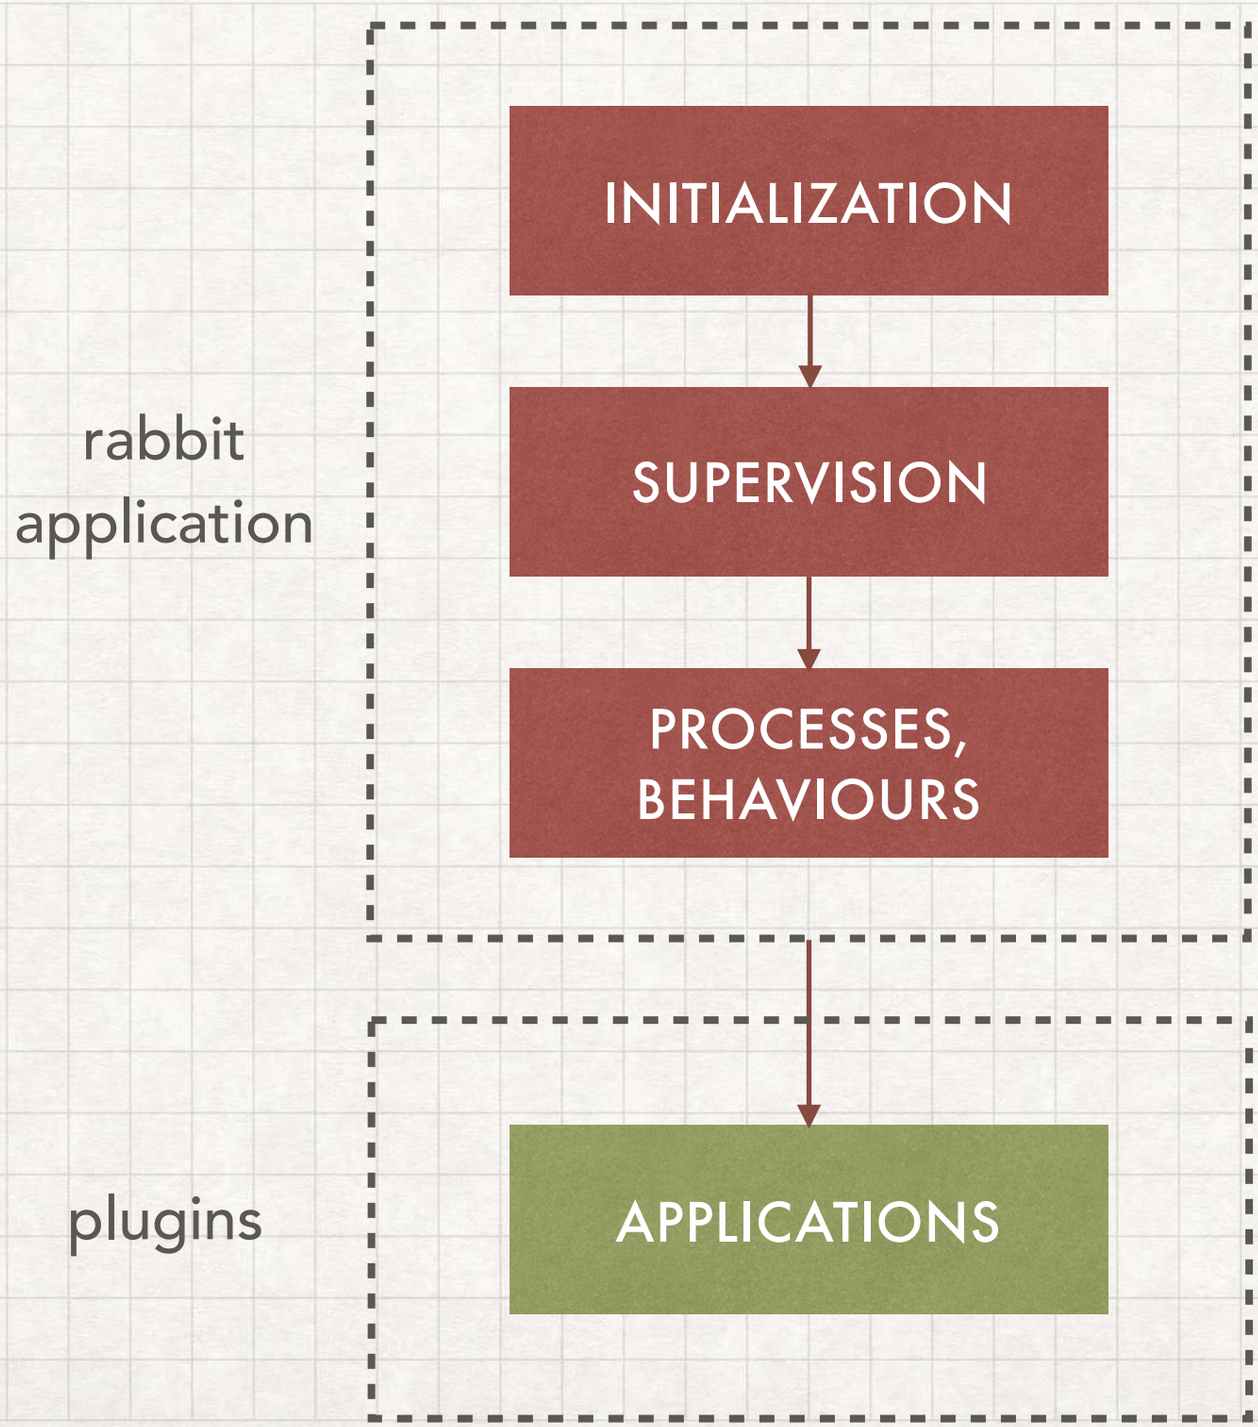
rabbit
application
INITIALIZATION
SUPERVISION
PROCESSES,
BEHAVIOURS
plugins
APPLICATIONS

APPLICATIONS: PLUGINS

- Plugin architecture “pattern”: highly extensible
- Plugins are simply OTP applications (zipped)
- Dynamically started/stopped (& expanded) via CLI
  - **rabbitmq-plugins enable** <PLUGIN/APP-NAME>
  - **rabbitmq-plugins disable** <PLUGIN/APP-NAME>
- Active plugins maintained in an **enabled_plugins** file
- Queried and updated on runtime
- Enabler for multi-protocol handling (MQTT, STOMP, ..)
- **Awesome** for abstracting Erlang/Elixir expertise!

CODEBASE: RABBITMQ

https://github.com/rabbitmq/rabbitmq-server

https://github.com/rabbitmq/rabbitmq-common

END: THANK YOU

QUESTIONS

Twitter: dube_aya

Github: Ayanda-D