From Calling to Eventing

Romain De Groote

24 September 2026

Step-by-step summaries of the SIN.04028 tutorial “From Calling to Eventing”, from direct method calls to an Event-Driven Architecture over MQTT. Process Control, given by Reto Konig, University of Fribourg, autumn semester 2026.

Table of contents

About these notes **4**
How to read the diagrams . . . 4
Building this book . . . 4

1 Step 0 — What the tutorial is about **5**
1.1 The three questions . . . 5

2 Part I — Start with what you already know **6**
2.1 Step 1 — Rock bottom: just call it . . . 6
2.2 Step 2 — Three egoistic capabilities . . . 7
2.3 Step 3 — Give every capability a mailbox . . . 7
2.4 Steps 4–6 — The three capabilities become egoistic . . . 7
2.5 Step 7 — Enter the mediator . . . 8
2.6 Step 8 — Ask again: who knows whom? . . . 9
2.7 Step 9 — Why virtual threads? . . . 10
2.8 Step 10 — Keep the services egoistic . . . 10
2.9 What part I moved . . . 10

3 Part II — Cross the process boundary **11**
3.1 Step 11 — Separate the components . . . 11
3.2 Steps 12–13 — A directory becomes a mailbox . . . 11
3.3 Step 14 — The first distributed-systems problem: publication . . . 12
3.4 Step 15 — Observe a mailbox with `WatchService` . . . 12
3.5 Steps 16–17 — Services and mediator become processes . . . 12
3.6 Step 18 — Run the fully external local version . . . 13
3.7 What part II moved . . . 14

About these notes

These notes summarise the tutorial **"SIN.04028 — From Calling to Eventing"**, worked through step by step. The tutorial itself lives next to this book, in `SIN.04028-From-Calling-to-Eventing/README.md`, and belongs to the *Process Control* course given by Reto Konig at the University of Fribourg (autumn semester 2026).

The program under construction never changes: you type a line, it comes back rotated by 13 letters.

```
keyboard → ROT13 → console
```

What changes is everything around it. The tutorial rebuilds the same system again and again while asking three questions, in this order:

> **Who knows whom?**
>
> **Who owns what?**
>
> **Who decides how the system is wired together?**

One summary is added here each time a step is finished, so the book grows at the pace of the work. Each summary says what the step did, what was observed when it ran, and what moved architecturally — usually a piece of knowledge that changed hands.

How to read the diagrams

Every diagram exists once as a drawing and three times as a language. In the web version, a diagram that carries a message flow has a small play bar under it: **Play** walks one message through the whole chain, **One hop** advances a single step, and **Reset** puts it back. Boxes light up as the message reaches them, so you can see which component is waiting and which one is working.

In the PDF, and in a browser with JavaScript switched off, the same picture is there — just still. Nothing in the notes depends on the animation.

Building this book

The three books come from one source:

```
./render-all.sh              # figures, then html + pdf for en, fr, de
./render-all.sh --serve      # …and serve _book/ on http://localhost:8000
```

`_book/` holds the English book, `_book/fr/` the French one and `_book/de/` the German one; the language menu in the sidebar switches between them and remembers the choice. A single `quarto render` builds English only and empties `_book/`, which removes the other two — so use the script.

1 Step 0 — What the tutorial is about

Step 0 contains no code. It states the subject: not ROT13, but the architecture around it. The functionality stays almost unchanged for 81 steps, which is what makes the changes around it visible.

Two rules are set here and hold for the whole tutorial:

Meet each solution only after you have experienced the problem it solves.

That is why the tutorial does not start with MQTT, and not with dynamic subscriptions either. It first builds a system with deliberately awkward mediators, puts it into operation, and only then asks whether they have earned their place.

The second rule is about vocabulary: *Intent*, *Status/State* and *Event* are named only after they have been used. They appear in the tutorial around step 42, not before.

The planned route, from step 0:

```
direct method calls → egoistic services → mediators → separate processes
→ filesystem as transport → Syncthing → "why all these mediators?" → rewind
→ services subscribe themselves → dynamic wiring → across processes
→ across Syncthing → when a mediator IS useful → MQTT
```

1.1 The three questions

The wording of the three questions matters, and it is easy to replace it with something weaker:

The tutorial asks	A weaker version
Who **knows** whom?	who *communicates* with whom
Who **owns** what?	which *data* each part holds
Who **decides** how the system is wired?	which part *connects* everything

"Knowing" is not "communicating": two components can exchange messages all day without either one naming the other. "Deciding the wiring" is not "connecting": the part that forwards messages need not be the part that chose the route. Both distinctions become concrete later — the first already in step 1, the second in part III.

2 Part I — Start with what you already know

Steps 1 to 10 stay inside one JVM. The system is built three times: as direct calls, as three capabilities that refuse to know each other, and as those same capabilities glued together by mediators. Nothing is distributed yet, and that is the point: the only thing that changes is who knows whom.

2.1 Step 1 — Rock bottom: just call it

`ProceduralRot13.java` is the whole system in one compact source file: a loop that reads a line, and one statement that prints the rotated text.

```
void main() {
    while (true) {
        var text = IO.readln();

        if (text == null)
            return;

        IO.println(rot13(text));
    }
}
```

`IO.readln()` stands on its own line, but `rot13(text)` is not a step of its own: it is called *inside* the argument of `IO.println`. Neither `rot13` nor `IO.println` knows anything about the other. The one that knows both is `main()`, which names them and fixes their order.

What `main()` knows about `rot13`: that it exists, how to invoke it (a bare name, both live in the same implicitly generated class), what argument it takes, what it returns, that it is available *right now*, and that execution continues after it returns. It does not know how `rot13` is implemented — that is the one thing it can ignore.

The availability guarantee is worth naming, because it disappears later: `main()` and `rot13` are two methods of one already-loaded class, called on the same thread, in the same process. Nothing has to be started first, no message travels anywhere, no listener has to be subscribed.

Observed: typing a line prints it rotated; `Ctrl+D` ends the program silently, because `IO.readln()` returns `null` and `main()` returns.



Figure 2.1: The procedural version. The wiring is written into the code: `main()` names all three parts and fixes their order.

2.2 Step 2 — Three egoistic capabilities

An **egoistic** component says: I know what I can do, what data I accept and what data I produce — and I would rather not know the business process around me. Concretely, `ConsoleIn` must not contain `rot13(...)`, and `ROT13` must not contain `consoleOut(...)`.

Removing those calls immediately creates a problem, which the tutorial states as a question: if nobody knows anybody, **who makes this into a system?** Note what the question does *not* say: the knowledge is not gone, it has to go somewhere else. Two connections are missing, not one — `ConsoleIn` ↔ `ROT13` and `ROT13` ↔ `ConsoleOut`.

2.3 Step 3 — Give every capability a mailbox

The smallest mechanism for the missing connections: a message and a named queue.

```
record Message(String type, String value) {}

Map<String, BlockingQueue<Message>> mailboxes = new ConcurrentHashMap<>();

BlockingQueue<Message> mailbox(String name) {
    return mailboxes.computeIfAbsent(name, _ -> new LinkedBlockingQueue<>());
}
```

The tutorial insists on calling these **mailboxes**, not topics, and gives the reason: if two readers call `take()` on the same mailbox, they compete — each message goes to exactly *one* of them and is not copied to both. A topic, in Publish/Subscribe, would hand every subscriber its own copy. One reader per message is enough for now, and the limitation is kept on purpose: it is what makes it clear later what Publish/Subscribe actually adds.

2.4 Steps 4–6 — The three capabilities become egoistic

Each capability now touches only its own mailbox names:

Capability	reads	writes
ConsoleIn	the keyboard	`consoleIn.outbox`
ROT13	`rot13.inbox`	`rot13.outbox`
ConsoleOut	`consoleOut.inbox`	the console

No capability contains the name of another. The tutorial's own summary of the result: "three nicely isolated capabilities. And they do absolutely nothing together." Nothing moves a message from `consoleIn.outbox` to `rot13.inbox`, or from `rot13.outbox` to `consoleOut.inbox`.

2.5 Step 7 — Enter the mediator

The smallest possible glue: a loop that takes from one mailbox and puts into another, knowing nothing about what travels through it.

```
void mediate(String from, String to) {
    while (true)
        mailbox(to).put(mailbox(from).take());
}
```

`main()` starts five virtual threads: the three capabilities, plus two mediators whose routes are given as arguments — `consoleIn.outbox` → `rot13.inbox` and `rot13.outbox` → `consoleOut.inbox`.

Observed: `Hello World` in, `Uryyb Jbeyq` out, exactly as in step 1. But `Ctrl+D` no longer ends the program. ConsoleIn's thread *returns* (it does not loop), and the other four threads sit blocked in `take()` on empty mailboxes, using no CPU. What actually keeps the process alive is the last line of `main()`, `Thread.currentThread().join()`: the main thread waits for itself, which can never finish. Virtual threads are daemon threads, so those four waiting threads would not have kept the JVM running on their own.

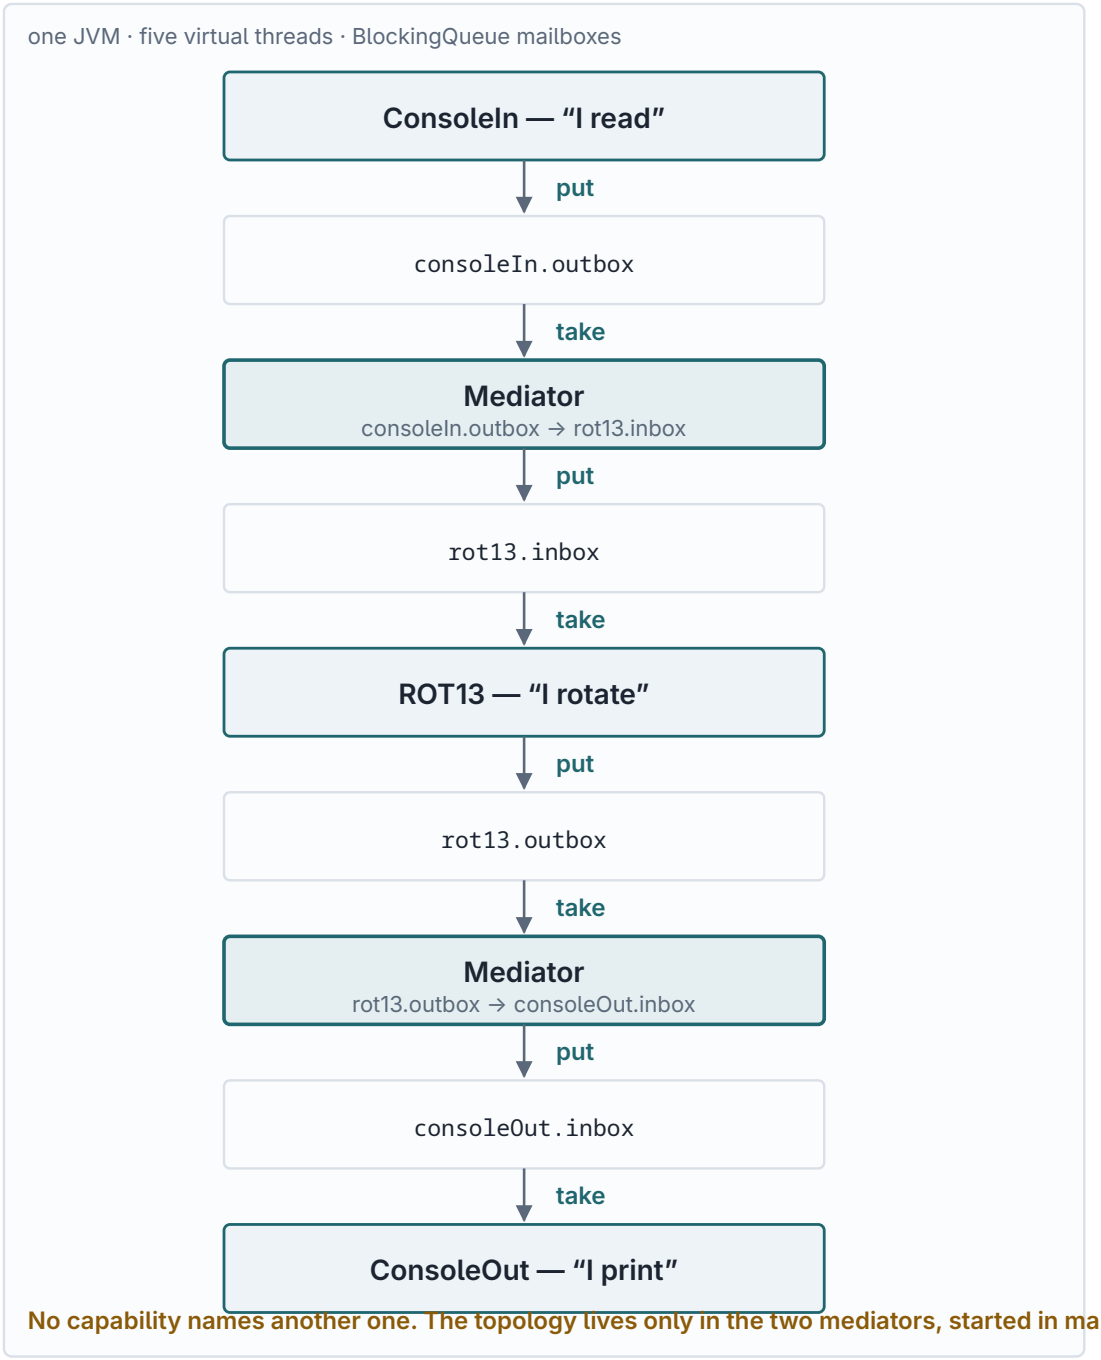


Figure 2.2: The mediated version inside one JVM. In the web version, **Play** walks one message through all eight hops.

2.6 Step 8 — Ask again: who knows whom?

Does ConsoleIn know ROT13? No. Does ROT13 know ConsoleIn or ConsoleOut? No. Has the topology disappeared? No — it sits in the two `mediate(...)` calls inside `main()`.

This is the tutorial's first major result:

> **Decoupling often does not make knowledge disappear. It moves knowledge to where it belongs.**

The capability "ROT13 transforms text" belongs to ROT13. The topology `ConsoleIn` → `ROT13` → `ConsoleOut` belongs to system composition. What changed since step 1 is not that the wiring left `main()` — it is still there — but that it is no longer written *inside* a capability. It is now data: two pairs of mailbox names. Rewiring touches those two lines and no capability at all.

2.7 Step 9 — Why virtual threads?

Five loops now wait indefinitely: the keyboard, `rot13.inbox`, `consoleOut.inbox` and the two mediators' sources. Virtual threads are simply the cheap Java way to let five waits happen at once inside one JVM.

They are an implementation convenience, not an architecture. The architecture — who knows whom, and where the wiring lives — is in the mailbox names, not in the threads. The tutorial's own evidence for the separation: the same architecture will run as separate processes in step 11, with no threads shared between the parts.

2.8 Step 10 — Keep the services egoistic

The obvious shortcut: the first mediator only copies messages unchanged, so why not delete it and let ROT13 `take()` from `consoleIn.outbox` directly?

Because ROT13 would then contain the name `consoleIn.outbox`, which means it would know that ConsoleIn exists and where it publishes. The topology would move *into* ROT13, and ROT13 would stop being an egoistic capability. The test that makes this concrete: suppose the input comes from a new `FileIn` capability publishing to `fileIn.outbox`.

	with the mediator	with the shortcut
what changes	one argument of the mediator	`runRot13` itself
who is touched	no capability	ROT13, whose job did not change

The tutorial is honest about the price: this design "may look indirect — even a little odd", and it promises a different way to keep services egoistic *without* a mediator in every message's path.

2.9 What part I moved

stage	where the topology lives	who could rewire it
step 1	inside the capability code (`IO.println(rot13(text))`)	whoever edits that code
steps 7–8	in `main()`, as pairs of mailbox names	whoever edits `main()`

Functionality: unchanged. Lines of code: more. Parts that know each other: none. That is the trade the rest of the tutorial keeps testing.

3 Part II — Cross the process boundary

Steps 11 to 18 take the architecture of part I and split it across five operating-system processes on one computer. The capabilities keep their mailbox names; only the transport underneath changes, from Java objects to files in directories. The first genuinely distributed problem shows up here, and it has nothing to do with ROT13.

3.1 Step 11 — Separate the components

Three processes cannot share a `BlockingQueue`: every `java` command starts its own JVM with its own memory, so the `mailboxes` map would exist three times and each copy would see only its own messages. An external transport is needed.

TCP would do, MQTT would do. The tutorial deliberately picks the most primitive option, the **filesystem**, and gives a one-line reason: *because you can see it.* Messages become files you can list with `ls` and read with `cat` while the system runs. Primitive here does not mean inferior; part VIII comes back to what the filesystem is genuinely good and bad at.

3.2 Steps 12–13 — A directory becomes a mailbox

The mapping is one line long:

```
directory = mailbox        file = message
publish   = create file    receive = watch directory
```

```
messageBoard/
├── consoleIn/outbox/
├── rot13/inbox/
├── rot13/outbox/
└── consoleOut/inbox/
```

The `Message` record gains an optional timestamp, and the file format stays two lines plus an optional third:

```
TEXT
Hello World
TIMESTAMP: 2026-09-22T10:42:12.123Z
```

A message written by hand may leave the timestamp out. The `TIMESTAMP:` marker keeps a payload line that happens to look like a date from being mistaken for metadata; in this tiny format, a final line starting with `TIMESTAMP:` is reserved.

Jackson could be introduced here and deliberately is not. The point of a ridiculously simple format is that almost nothing hides the architecture — the whole message is readable in any text editor, which is what keeps the attention on who talks to whom.

3.3 Step 14 — The first distributed-systems problem: publication

A reader watching a directory can discover a file *before the writer has finished writing it.* Hence a publication protocol:

1. create a private temporary file;
2. write the complete message;
3. close it;
4. atomically rename it to its final `.msg` name;
5. consumers react only to final files.

```
Files.writeString(temporary, content);
Files.move(temporary, published, StandardCopyOption.ATOMIC_MOVE);
```

Both halves are needed. The rename is atomic — a reader sees either no `.msg` file or the complete one — but it protects nothing if readers also pick up `.tmp` files. `.tmp` and `.msg` stay in the same directory so the move cannot cross filesystems, where atomicity would be lost.

> **Writing and publishing are not the same operation.**

3.4 Step 15 — Observe a mailbox with `WatchService`

Java's own filesystem observer supplies the consumer half: register a directory for `ENTRY_CREATE`, block in `watcher.take()`, and ignore every filename that does not end in `.msg`.

That also gives a low-tech injection tool. With `ActuatorService` running as the consumer:

```
printf 'TEXT\nHello from a file' > messageBoard/consoleOut/inbox/manual.tmp
mv messageBoard/consoleOut/inbox/manual.tmp messageBoard/consoleOut/inbox/manual.msg
```

Observed: nothing happens when the `.tmp` file appears; the line is printed the moment the rename creates the `.msg` name. Both accepted formats work, with and without the trailing timestamp.

The architectural lesson is the one the tutorial draws as `producer` → `filesystem` → `consumer`: the consumer could not tell a human typing `mv` from a real service. It knows its own inbox, not who writes into it. Egoism survived the process boundary.

3.5 Steps 16–17 — Services and mediator become processes

The three capabilities become three programs, and the mediator a fourth kind of program:

program	plays	reads	writes
`SensorService.java`	ConsoleIn	the keyboard	`consoleIn/outbox`
`ProcessorService.java`	ROT13	`rot13/inbox`	`rot13/outbox`
`ActuatorService.java`	ConsoleOut	`consoleOut/inbox`	the console
`Mediator.java <from> <to>`	the glue	`<from>`	`<to>`

The mediator is generic: its route is not compiled in, it arrives as two startup arguments. But the tutorial points at an asymmetry worth keeping in mind: the three capabilities receive their **work** as messages in a mailbox, while the mediator receives its **route** as startup configuration, and it has no name and no control mailbox of its own that anyone could write to.

The practical consequence: arguments are read once, so changing the route of a running mediator means stopping it and starting it again. A configuration file would be no different — it is still read at startup. The tutorial deliberately leaves this unfixed until the cost of the distributed mediators has been felt.

3.6 Step 18 — Run the fully external local version

Five terminals, five processes, one local filesystem between them:

```
java SensorService.java                                  # terminal 1
java Mediator.java consoleIn/outbox rot13/inbox          # terminal 2
java ProcessorService.java                               # terminal 3
java Mediator.java rot13/outbox consoleOut/inbox         # terminal 4
java ActuatorService.java                                # terminal 5
```

Observed: `Hello World` typed in terminal 1 appeared as `Uryyb Jbeyq` in terminal 5. The three middle terminals showed nothing new — each mediator prints one "Mediating: …" line at startup, and `ProcessorService` prints nothing at all. Afterwards each of the four mailbox directories held one `.msg` file: the same line, copied along the chain.

Count it: **five processes** for what step 1 did in one statement. The tutorial's instruction at this point is not to optimise: "We are not done making it painful."

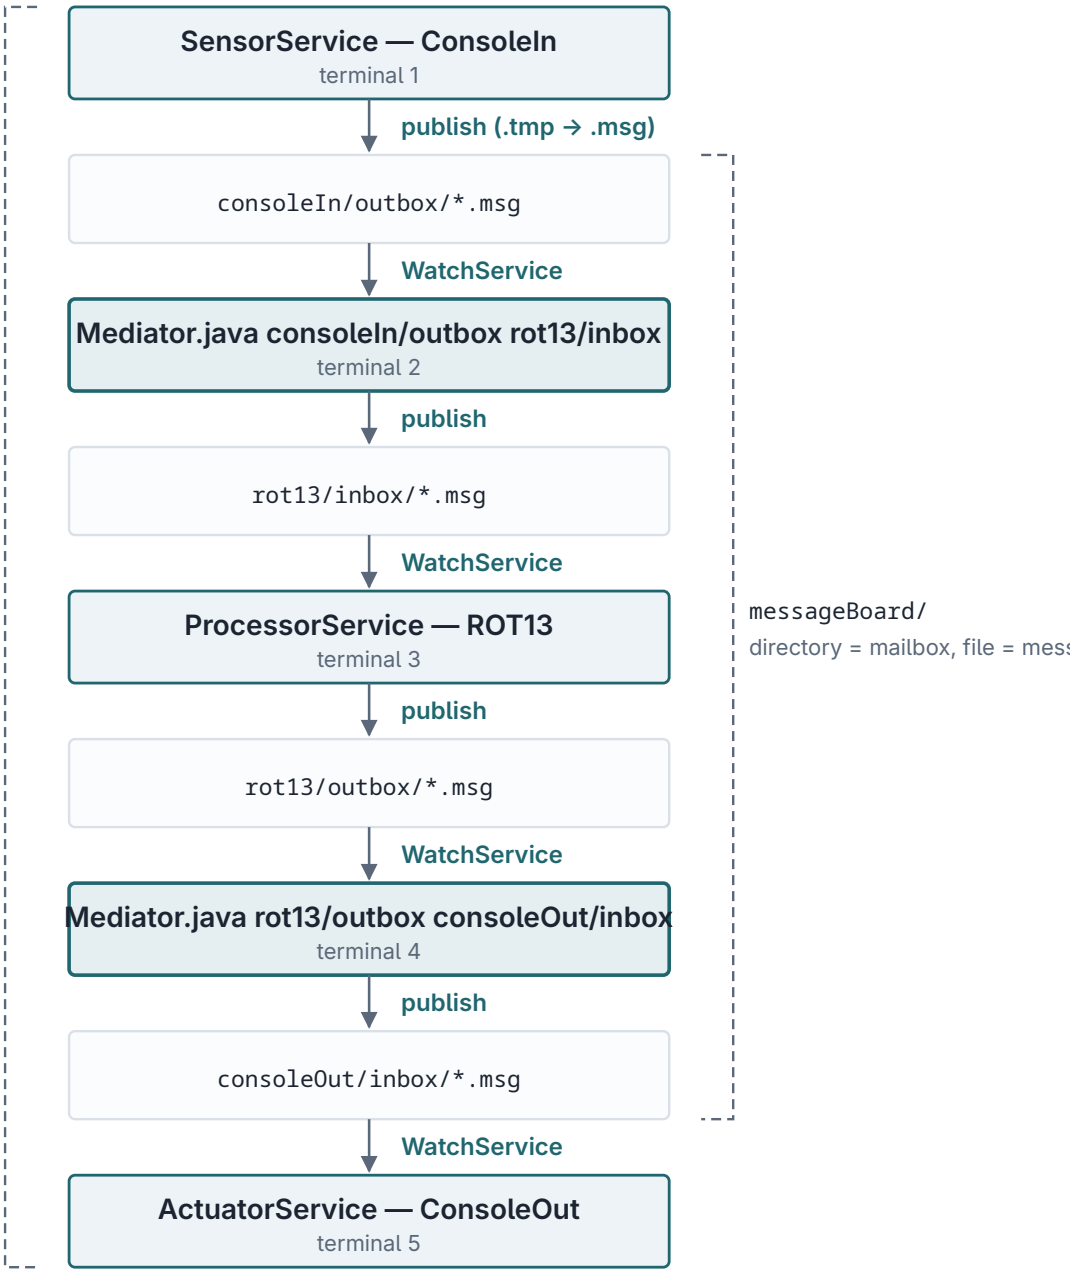


Figure 3.1: The five processes and the four mailbox directories. In the web version, **Play** follows one line of text from the keyboard to the console.

3.7 What part II moved

stage	where the topology lives	who could rewire it
step 1	inside the capability code	whoever edits that code
steps 7–8	in `main()`, as mailbox names	whoever edits `main()`
step 18	in the two `Mediator.java` `<from> <to>` command lines	whoever starts the processes

The topology has now left the source code altogether: no `.java` file contains the route. It lives in what the operator typed, which means a person is currently the orchestrator. What the five-process version bought is real — the wiring can change without touching any program — and what it cost is equally real: five processes, four copies of every message, and a program that no longer stops when you press `Ctrl+D`.