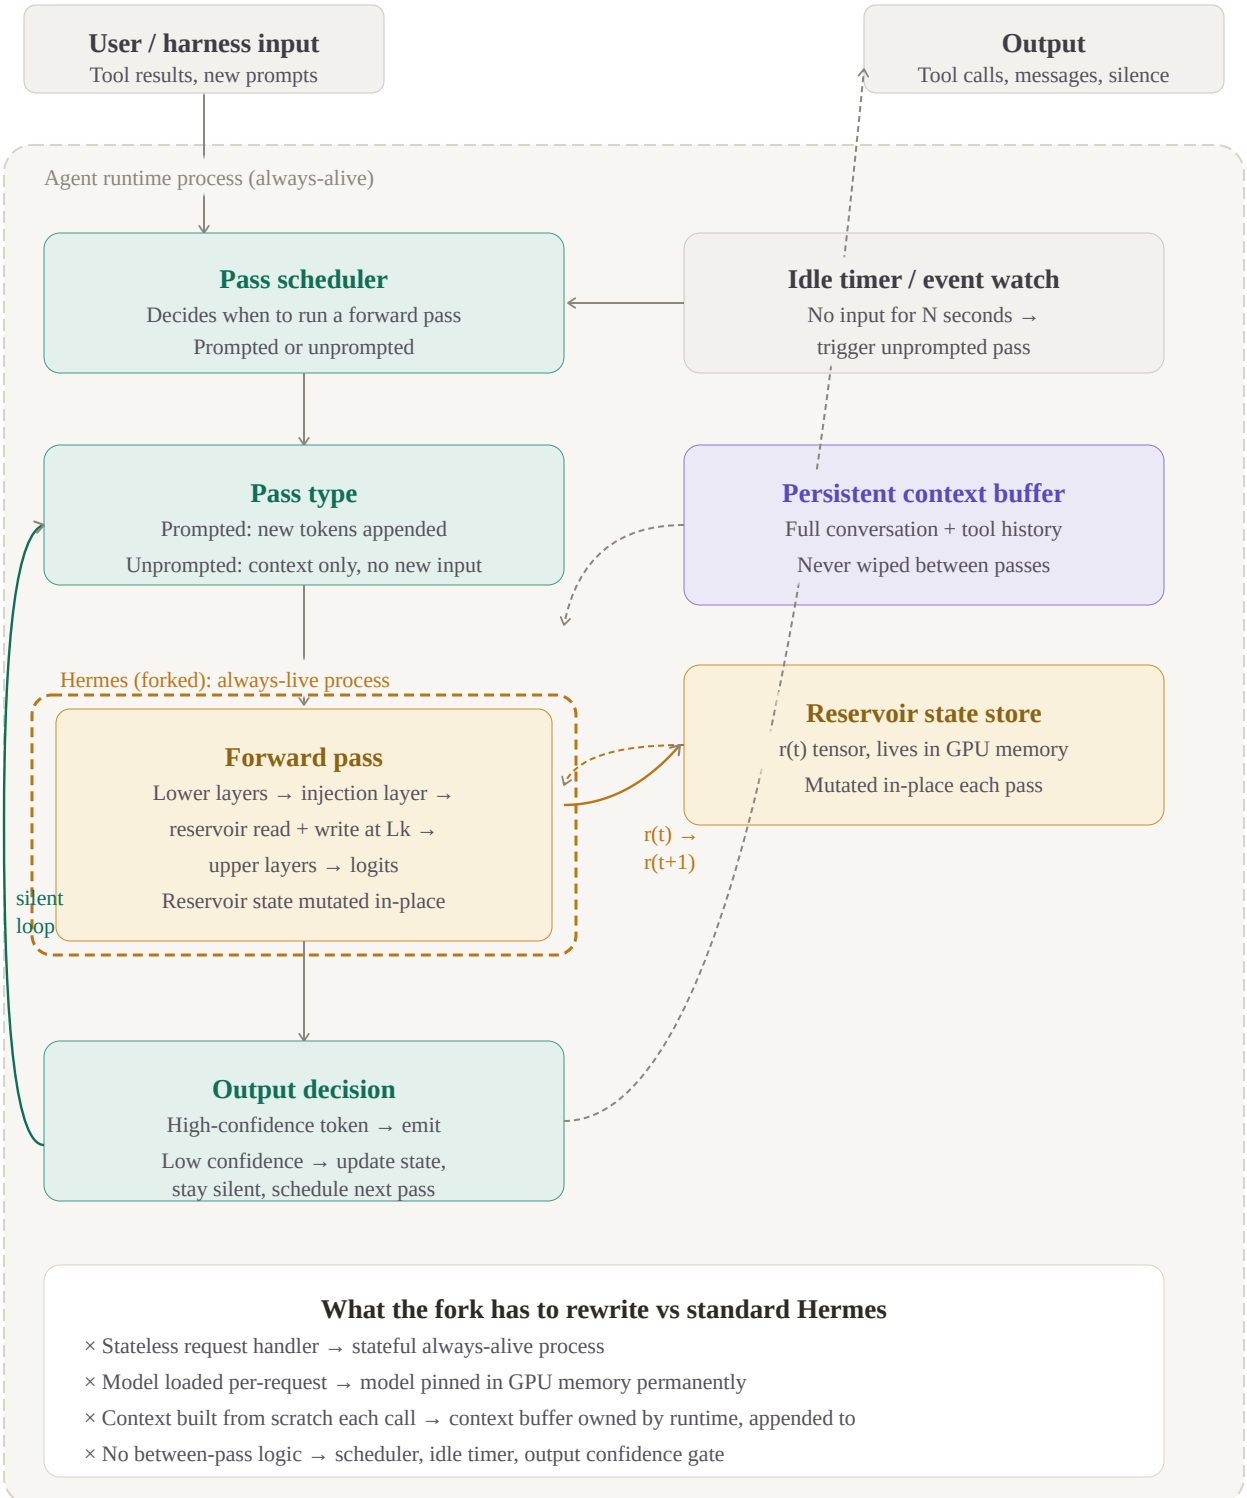


What the fork has to rewrite vs standard Hermes

- × Stateless request handler → stateful always-alive process
- × Model loaded per-request → model pinned in GPU memory permanently
- × Context built from scratch each call → context buffer owned by runtime, appended to
- × No between-pass logic → scheduler, idle timer, output confidence gate