

AIAUGMENTED ORG PACKET: RACI, SLA MATRIX, COVERAGE GRID, AND ESCALATION TREE

A plain-English, copy-ready packet for nomad teams: roles + RACI, default SLAs, a UTC coverage grid with overlap math, a backup-reviewer rotation, a two-tier escalation tree, a blameless post-mortem template, and the Lisbon Test handoff packet. Duplicate it and keep work moving while you're in the air.

FROM EPISODE

[BUILD AN AI ORG CHART THAT WORKS WHILE YOU SLEEP](#)

CONTENTS

- 1) Roles + RACI (AI in the loop, no bystanders)
- 2) Default SLA matrix (publish and iterate)
- 3) UTC coverage grid + overlap calculator (one sheet, no guessing)
- 4) Backup Reviewer rotation (Reviewer Roulette)
- 5) Two-tier escalation tree (mapped to SEV/P levels)
- 6) Handoff packet (don't ship without it)
- 7) Blameless post-mortem (standard template)
- 8) 7-day rollout plan (pilot, then lock)

Duplicate this packet into Notion or Google Sheets and run the Lisbon Test: could your work keep moving for 24 hours while you're offline? This guide gives you the minimum viable scaffold-

ing—roles, SLAs, coverage, backups, escalation, post-mortems, and a handoff packet—so clients always know who's on point and what happens next.

1) ROLES + RACI (AI IN THE LOOP, NO BYSTANDERS)

Use these four seats across all work. Keep titles simple and stable across pods.

- Builder — ships the thing. Creates drafts, PRs, workflows, prompts.
- Operator — owns quality, schedules, budgets, and client comms. Can approve/rollback.
- Reviewer — independent check on risk/quality. Named backup is required.
- Agent/Dispatcher — routes work, maintains the coverage grid, pages on escalation. This can be a person, an on-call rotation, or an automation plus a human backstop.

RACI quick template (copy into Notion/Sheet):

- Task: Inbound lead triage → R: Agent · A: Operator · C: Builder · I: Reviewer
- Task: Content approval (ad/email/social) → R: Operator · A: Reviewer · C: Builder · I: Agent
- Task: Deployment/code review → R: Reviewer · A: Operator · C: Builder · I: Agent
- Task: Client weekly report → R: Operator · A: Operator · C: Builder · I: Reviewer

Guidelines:

- One A (Accountable) per task. If you have two, you have none.
- Reviewer must never be the same person as the Builder on P1/P2 work.
- Document decisions where the work lives (PR, task, or artifact). No private DMs for approvals.

2) DEFAULT SLA MATRIX (PUBLISH AND ITERATE)

Start with these defaults and tune by volume, risk, and staffing. Publish them where clients can see them.

Inbound leads (during sender's hours):

- Positive/qualified: first reply within 5–15 minutes. If no human ack in 15 minutes, auto-escalate to on-call. Next reply within 60 minutes until booked/qualified. Owner: Agent (R), Operator (A).
- Neutral/unknown: first reply within 1 business hour. Owner: Agent (R), Operator (A).

Content approvals (ads, email, landing pages):

- First decision within 24 hours for P1 campaigns; within 48 hours for P2–P3. On silence: auto-approve after 48 hours only if guardrails are met (style/policy checklist passed and a Reviewer pre-approved the template). Owner: Operator (R), Reviewer (A).
- Decision options: Approve, Approve with notes, Block (must include reason + fix path).

Code/automation PRs:

- First review within 4 business hours for small PRs (<150 LOC or low-risk automations). Within 8 business hours for medium; negotiate timeline for large/risky.
- Completion: 24–72 hours depending on size/risk. Owner: Reviewer (R), Operator (A).

Support/requests by priority:

- P1 (revenue-impacting outage or ad spend waste): ack ≤15 minutes, mitigation start ≤30 minutes.
- P2 (major degradation, deadline today): ack ≤30–60 minutes, plan posted ≤2 hours.
- P3 (normal work/blocker without revenue impact): ack same business day.

- P4 (nice-to-have/ideas): ack within 1 business day.

How to adapt fast:

- Track first-reply and next-reply separately. Quiet queues hide leaks.
- Tie PR size to SLAs (S/M/L) and publish the size rubric on your repo/automation board.
- Make SLAs visible on the client hub and inside your task template. If it's not visible, it doesn't exist.

3) UTC COVERAGE GRID + OVERLAP CALCULATOR (ONE SHEET, NO GUESSING)

Build a single UTC view so you can see gaps before they bite you.

Sheet columns (one row per person):

- Name · Role (B/O/R/A) · Home UTC offset (e.g., -03:00) · Local Work Start (HH:MM) · Local Work End (HH:MM) · PTO/OOO dates · Skills/Systems (e.g., FB Ads, Make.com, Python)

Convert to UTC:

- $\text{StartUTC} = \text{LocalStart} - \text{UTCOffset}$
- $\text{EndUTC} = \text{LocalEnd} - \text{UTCOffset}$

Pairwise overlap (hours) between Person A and B:

- $\text{Overlap} = \text{MAX}(0, \text{MIN}(\text{EndUTC}_A, \text{EndUTC}_B) - \text{MAX}(\text{StartUTC}_A, \text{StartUTC}_B))$

Coverage guardrails:

- Aim for ≥ 2 hours daily overlap between Builder and Reviewer; ≥ 3 hours between Operator and Agent.

- For follow-the-sun, ensure every 24-hour cycle has: (a) named on-call Agent, (b) at least one Reviewer with ≥ 2 hours overlap with each Builder they backstop, (c) an emergency paging route.
- Publish today's on-call and the backup Reviewer at the top of your client hub.

Common failure points (and fixes):

- Hidden gaps around weekends or daylight-saving shifts → add "Local Weekends" and "Observes DST?" flags; render a weekly heatmap in UTC.
- Ownership confusion at handoff → require a handoff packet (next section) and make the receiving person comment "I own it" where the work lives.

4) BACKUP REVIEWER ROTATION (REVIEWER ROULETTE)

Prevent stalls by rotating secondary reviewers and balancing load.

Rotation setup:

- Maintain a Reviewer pool per pod (min 3 people). Mark skills and PTO.
- Rule: the Builder cannot be the Reviewer. For P1/P2 work, the Reviewer cannot be the Builder's direct backup on that artifact.

Assignment algorithm (simple and effective):

1. Sort open reviews by oldest first.
2. From the pool, pick the next available Reviewer who isn't the author and has ≥ 2 hours overlap in the next 24 hours.
3. If none, assign the on-call backup Reviewer and page the Operator to adjust scope/timing.

4. Rotate the pool pointer after each assignment (Roulette).

Implementation tips:

- Keep a “Reviewer Roulette” tab: columns for Date, PR/Task link, Assigned Reviewer, SLA due, Status. Freeze the pointer cell so the next person is obvious.
- Use CODEOWNERS or automation rules for default reviewers, but keep this manual roulette as the human override for load and PTO.
- Post the roulette result as a comment on the PR/task so the SLA clock is public.

5) TWOTIER ESCALATION TREE (MAPPED TO SEV/P LEVELS)

Two layers are enough for small teams. Tie actions to impact, not loudness.

Severity definitions:

- SEV1 / P1 — Known or likely revenue impact now (e.g., broken checkout, suspended ads, mis-spent budget).
- SEV2 / P2 — Major degradation or deadline today; revenue risk soon.
- SEV3 / P3 — Normal defects or questions; business-hours handling.

Escalation tree (ack = acknowledgement of ownership):

- Layer 1 (on-call Agent): ack in 5–15 minutes for SEV1; ≤30 minutes for SEV2. If no ack, auto-escalate.
- Layer 2 (Operator): if Layer 1 misses ack window, page Operator; Operator may re-route or roll back. For SEV1, Operator must ack within 15 minutes of page.
-

Advisory (Reviewer/Builder): pulled in by Operator for fixes or rollbacks; not paged for SEV3 after hours.

Paging routes:

- Primary: tool-based (on-call schedule) → mobile/email → fallback voice call.
- Evidence log: every page creates a timestamped note on the incident card/PR with owner, severity, and next checkpoint time.

Org patterns:

- Solo + contractors: You are Operator; name one on-call Agent per day. Keep a single phone escalation for SEV1.
- Pod: Agent pages pod Operator; Operator pulls the right Builder/Reviewer.
- Agency cell: Dispatcher (Agent) pages pod Operator; if two pods miss ack, escalate to Cell Lead.

Note: Treat the 15-minute window as sacred only for revenue-critical paths. Everything else can wait until business hours.

6) HANDOFF PACKET (DON'T SHIP WITHOUT IT)

Before you hand work across time zones, attach this packet to the task/PR and @-mention the receiver. If any field is blank, you don't have a handoff—delay or page.

Handoff packet fields (Lisbon Test):

- Context: what we're doing and for whom. Link the brief/spec.
- Constraints: deadlines, budgets, platform limits, legal/brand rules.
- Last good output: the most recent artifact or working state to copy/continue. Include screenshots.

- Budget left: hours/dollars remaining for this scope; include who can approve overage.
- Fallback: safe rollback or alternate plan if blocked in the next 12 hours.
- Owner on receipt: tag the human who now owns it; they must comment “I own it” and restate the next checkpoint time in UTC.

Example snippet you can paste:

```
Context: [Client/Product], [Goal], links → [Brief] [Jira/Notion] [Repo]
Constraints: [Deadline UTC], [Budget $/hrs], [Brand/Legal notes]
Last good output: [Link/Screenshot]; known issues: [list]
Budget left: [X hrs / $Y]; overage approver: [Name]
Fallback: [Rollback plan] or [Alternate copy v2]
Receiver: @[Name]; Next checkpoint: [YYYYMMDD HH:MM UTC]
```

7) BLAMELESS POSTMORTEM (STANDARD TEMPLATE)

Keep blame out; keep facts in. This is how your system learns.

Template (copy/paste and fill):

- Incident title: [What broke] — [Date UTC]
- Severity: [SEV1/SEV2/SEV3]
- Impact summary: [Who/what was affected; time window; rough \$\$\$/risk]
- Timeline (UTC):
 - [T0] Detected by [tool/person]
 - [T+15m] Acknowledged by [Agent]

- [T+30m] Mitigation started by [Operator]
- [T+Xh] Resolved; verification by [Reviewer]
- Root cause: [Systems + contributing factors]
- What worked: [Processes/tools that helped]
- What failed: [Where we lost time/context]
- Fixes (ranked):
 1. [Prevent class of issue] — owner [Name] — due [Date]
 2. [Reduce time to detect/acknowledge] — owner [Name] — due [Date]
 3. [Documentation/guardrail update] — owner [Name] — due [Date]
- Evidence: links to PRs, screenshots, logs, client comms
- Review date: [YYYY-MM-DD]; sign-off by [Operator]

Ritual:

- For SEV1/SEV2, run within 72 hours while details are fresh.
- Publish to your “Incidents” page and link from the artifact that failed.

8) 7DAY ROLLOUT PLAN (PILOT, THEN LOCK)

Timebox the rollout. Prove it on one pod/client, then scale.

Day 1 — Duplicate the packet; list your current offers and create one-page RACIs. Assign a named backup Reviewer for each.

Day 2 — Publish SLAs on your client hub and inside task/PR templates. Add the auto-approve rule for content with guardrails.

Day 3 — Build the UTC coverage grid; confirm ≥ 2 hours Builder $\leftrightarrow$ Reviewer overlap per pairing. Set the on-call Agent schedule.

Day 4 — Stand up Reviewer Roulette (sheet tab + pointer). Dry-run on two old PRs.

Day 5 — Implement the two-tier escalation tree in your paging tool (or manual protocol). Test the 15-minute SEV1 path end-to-end.

Day 6 — Run a red-team handoff overnight using the packet. The receiver must ship or escalate without meeting you.

Day 7 — Hold a 30-minute retro. Capture fixes and lock the process for the next 30 days.

Compliance readiness note: Most AI oversight provisions in the EU start applying on August 2, 2026. Clear human ownership (Operator/Reviewer), visible SLAs, and evidence logs help you show control without adding heavy process. This guide is operational, not legal advice.