

Observability & Reliability Best Practices

SLOs, alerting, runbooks, and auto-remediation that actually reduce toil.

Define SLOs first

Start from the user's experience. Define service level objectives for the journeys that matter (login, search, checkout), then instrument to measure them. SLOs — not raw resource metrics — should drive alerting and prioritization.

Make alerting signal-rich

- Alert on symptoms (SLO burn) rather than every cause.
- Use dynamic baselines instead of static thresholds.
- Attach context to every alert: probable cause, owner, runbook.
- Continuously prune alerts that never lead to action.

Codify runbooks

Every recurring incident should have a runbook. Once a runbook is proven safe, promote it to automation. Codified response turns hours of manual triage into seconds of consistent action.

Automate remediation safely

- Start with low-risk actions: restart, recycle, scale.
- Add guardrails: rate limits, blast-radius checks, approvals for risky changes.
- Always validate recovery (synthetically) before closing.
- Record every action for audit and learning.

Run blameless incident reviews

Treat every incident as a chance to improve detection, correlation, or automation. Feed learnings back into SLOs, alerts, and runbooks.

Want this tailored to your environment? Book a demo at applicare.arcturustech.com/support.html