ROKT

AI Playbook

Operations

How the Operations team automated the work that slowed them down

Built by the people who lived it.

Table of Contents

WORKFLOW AUTOMATION

AI-Assisted Creative Approvals

The Context / Challenge

Creative approvals were a major bottleneck in the campaign delivery pipeline. Every submission required a human reviewer to manually check for policy violations, image quality issues, and landing page problems. The volume of submissions meant reviewers were spending the majority of their time on routine checks, leaving less time for the nuanced judgment calls that actually benefit from human expertise. Routine issues that should have been caught before submission were regularly making it to human review, creating rework cycles and slowing turnaround times.

Our Approach

Rather than building a single AI check, the team designed a multi-layer review system. Hard-coded policy checks parse each character and detect major issues deterministically, then an AI layer handles subjective judgment calls that require contextual understanding. Results from both layers are aggregated and logged. Humans still make the final call on every submission, but AI handles the mechanical screening across multiple dimensions. The system posts structured findings directly to the relevant Jira ticket so reviewers can see exactly what was flagged and why.

Results

The multi-layer review system reduced the time reviewers spend on routine checks and catches deterministic issues before they reach human review. Reviewers now focus on the subjective judgment calls that actually require human expertise, with full visibility into what the automated layers have already checked. Structured findings posted directly to tickets eliminated the back-and-forth that previously slowed turnaround times.

TOOLING & ENABLEMENT

Ads Operations Hub

The Context / Challenge

As teams across the Ads organization built automation tools, dashboards, and AI-powered workflows, a new problem emerged: discovery. Tools were being built but not consistently adopted because people did not know they existed, could not find them, or were not sure which tool was the right one for their task. The tooling landscape was fragmented — some tools were well-maintained and actively used, others were built for one-off needs and forgotten. There was no central place to see what was available, how much each tool was being used, or where the gaps were.

Our Approach

The team built a centralized Ads Operations Hub that brings together automation tools, dashboards, product updates, and team-built projects in one platform. The hub includes tool discovery, usage analytics, and a feedback loop so the most impactful tools get maintained and improved. Each project page includes documentation, a walkthrough video, usage analytics, version history, and a feedback mechanism. A strong example: the Mock Generator tool reduced creative mock generation for 18 mocks from 30 minutes to 30 seconds.

Results

The Ads Operations Hub solved the discovery problem by giving the entire Ads organization a single place to find, use, and provide feedback on the tools being built across the team. Usage analytics now inform which tools get continued investment, and the feedback loop ensures the most impactful tools are maintained and improved. The model is being extended to the full Operations function.

METRICS & REPORTING

Operational Metrics Dashboard

The Context / Challenge

When leadership needed operational metrics, someone had to go pull the data from multiple systems, assemble it in a spreadsheet, and put together a report. This process was slow, manual, and meant that the numbers leadership saw were always slightly stale. Questions during meetings often required follow-up because the data was not at hand. Team leads were spending time compiling reports rather than acting on the insights. The metrics infrastructure was reactive: leadership had to ask for numbers instead of having them available on demand.

Our Approach

The team built a centralized leadership dashboard (the "Ops Command Center") that replaced manual data pulls with real-time views of team workload, project health, and operational performance. Behind the dashboard, n8n automations track all Jira requests, calculate metrics, and send results to a Google Sheets database. A separate workflow tracks client meeting time and communication patterns. The Replit-hosted app serves as the UI layer, giving leadership deep-dive views across approvals, QA, speed-to-launch, audits, anomalies, and more.

Results

The Ops Command Center replaced manual data pulls and spreadsheet-based reporting with real-time views that leadership can access on demand. Team leads no longer spend time compiling reports, and leadership no longer has to wait for numbers. Deep-dive views across approvals, QA, speed-to-launch, audits, and anomalies give leadership self-serve access to every operational metric without requesting a report.

TEAM DEVELOPMENT

End-to-End Onboarding Program

The Context / Challenge

Onboarding new team members into the Operations function was inconsistent and largely ad-hoc. The experience varied depending on who facilitated it, critical knowledge gaps went undetected until weeks into the role, and managers had limited visibility into whether their new hires were actually ramping effectively. There was no structured curriculum, no centralized resource library, and no way to assess whether someone had absorbed the material. As the team scaled, this created compounding problems: inconsistent quality, longer ramp times, and a growing dependency on individual tribal knowledge holders.

Our Approach

The team designed a structured, multi-track onboarding curriculum that replaced the ad-hoc approach with a consistent experience for every new team member. At its core is an AI-powered hub built with Claude that includes knowledge assessments, a resource library, and progress tracking for managers. The assessments test whether someone actually understands the material, not just whether they sat through the session. Managers can see exactly where each new hire stands in their ramp, which topics need reinforcement, and whether the program is working.

Results

The onboarding program replaced a process that varied by person with a structured, repeatable curriculum. New hires now follow the same multi-track program regardless of location or facilitator. Managers have visibility into ramp progress and can identify knowledge gaps early rather than discovering them weeks into the role. The AI-powered assessments ensure the program measures actual understanding, not just attendance.

WORKFLOW AUTOMATION

QA Alignment — Chat Integration

The Context / Challenge

Partner integrations are a complex, multi-step process with multiple stakeholders involved. When it comes time to QA the integration, it is critical to ensure that teams are aligned, have the right context, and are able to surface blockers, updates, and dependencies quickly to maintain momentum. As integrations scaled, Jira became a limiting factor: important comments could get buried, context was spread across multiple tickets, and staging QA, production QA, and post-launch work lacked a single cohesive thread connecting them.

Our Approach

The team built a chat space automation to centralize end-to-end context for the entire QA process. When a Jira ticket is assigned, a chat space is automatically created, adding relevant stakeholders at each stage, surfacing ticket context, notifying members of their assigned tasks, and unifying every step from staging through production to post-launch checks. The result is reduced context loss, full visibility for all stakeholders, and faster decision-making and communication — so the team can launch partners as quickly as possible without compromising quality.

Results

The automated chat space system replaced the scattered, ticket-comment-driven coordination that was slowing partner launches. Every QA process now has a single cohesive thread from staging through production to post-launch, with the right people added at the right time and full context carried forward between stages. The team moves faster, context loss is eliminated, and blockers surface immediately instead of getting buried in Jira comments.

QUALITY ASSURANCE

QA Automation — Prism

The Context / Challenge

Partner QA is a manual, human-executed process. An Ops team member must navigate a partner's checkout flow, capture the Rokt experience data from browser DevTools, verify the data integration, record screenshots along the way, and manually upload this information into a QA document. This workflow is time-intensive, prone to inconsistency, and creates a bottleneck at high-launch volume. It also takes time away from higher-leverage tasks such as analyzing the customer experience overall and identifying improvements to the Rokt–client setup.

Our Approach

The team initially set out to automate the process entirely using a fully autonomous browser agent. In controlled scenarios it worked well, but in practice the long tail of edge cases — CAPTCHA gates, mandatory account creation, bot detection, atypical checkout flows — required progressively more engineering investment for diminishing returns. Full automation was solving the wrong problem first. The team pivoted to a human–LLM collaboration model named Prism. An Ops team member drives the test purchase while a browser extension records everything: network requests, Rokt integration data, and UI screenshots at each key step. At the end of the session, the system validates all integration attributes against the checklist, validates UI/UX results against the customer policy handbook, and generates a standardised QA document with structured analysis.

Results

The result is the same output quality as a fully automated run — with much less investment required — and it works on every partner site without exception. The approach leverages the best of what humans and AI can do together: humans handle unpredictable checkout navigation, AI handles structured validation and documentation. The pivot from full autonomy to human–LLM collaboration is itself a key lesson in identifying where AI adds the most value. Built by Julian Mullins and Pep Pattamasaevi.