GROWTH TERMINAL

Growth Review Portal

User Guide

A → Z

A practitioner's walkthrough of the Growth Review Portal - every destination, every tab, every shipping feature.

Audience	Consultants, agencies, and operators using the portal.
Edition	v1.0 reflects the portal as it ships today
How to read	Sequential by section, or jump in via the Table of Contents.

CONTENTS

Table of contents

SECTION 01 REFERENCE

Welcome - what the portal is for

Growth Review Portal is a guided diagnostic platform for consultants, agencies and operators. It helps you find the single biggest constraint holding back a business, forecast the upside of removing it, ship a 90-day plan, and verify whether the predicted lift actually happened. Every figure on screen is traceable: confidence scores, evidence sources, and engine version are stamped on every saved review.

STEPS

01 Bring the business context (industry, current revenue, funnel metrics) plus optional live data via connectors.

02 The portal returns a primary constraint, a likely revenue range, a ranked 90-day plan, and a tracking workspace.

03 Verify by recording outcomes (or letting Stripe revenue auto-grade them) and comparing against the saved forecast.

> **NOTE**
> This guide covers the portal product only. The marketing site, public framework PDF, and admin Customer Dashboard are documented separately.

SECTION 02 ACCOUNT

Signing in & workspace basics

Authentication is handled by Clerk. New users come in through the marketing site or via a claim link from a paid checkout intent. Once signed in, every consultant gets a personal workspace; agencies can switch between client workspaces from the top-bar workspace switcher.

STEPS

01 Click "Sign in" from the marketing site or follow the link in your welcome email.

02 Pick (or create) the workspace you want to work in from the top-bar workspace switcher.

03 Use the sidebar (or the mobile tab bar) to navigate between the seven top-level destinations.

> TIP
> Press Cmd+K (Ctrl+K on Windows) anywhere in the portal to open the command palette and jump to any section, save a review, or load the sample data.

SECTION 03 REFERENCE

Information architecture - the seven destinations

The portal is organized around the Diagnose → Predict → Verify loop, plus an Account block. Six destinations live in the sidebar; Messages stays routable but is reached through the floating ASK button and notifications, not the rail; Clients is admin-only and opens the standalone Customer Dashboard.

DIAGNOSE

Overview

Single page (no sub-tabs)

DIAGNOSE

Data Sources

Integrations Live Integrations Business Inputs Metrics

DIAGNOSE

Diagnostic

Primary Constraint Transparency Sensitivity Decision Summary

PREDICT

Forecasts

Projected Outcome Financial Impact Scenario Simulator

VERIFY

90-Day Plan

Game Plan Progress Outcome Validation

ACCOUNT

Settings

Account Settings Automation Policy Referrals Billing Security & Trust Audit Log

ACCOUNT

Clients (admin)

Opens the Customer Dashboard

NOTE

Old URLs (e.g. /game-plan, /integrations, /billing) still resolve - they redirect into the right tab on the new top-level page so deep links from QBR recaps, emails, and procurement docs keep working.

SECTION 04 DIAGNOSE

Overview

Overview is the home page after sign-in. It surfaces current workspace status: the active constraint, the most recent review, outstanding plan moves, and any blockers raised from progress check-ins. It is the only top-level destination without sub-tabs - read it top-down and follow the suggested next step.

STEPS

01 Open Overview from the sidebar (it is the default route after sign-in).

02 Read the headline status card - it tells you what the engine currently believes the binding constraint is.

03 Scroll to "Next steps" and click the suggested action for today (start a diagnostic, finish intake, log progress, share a review).

SECTION 05 DIAGNOSE

Data Sources - inputs, integrations, metrics

Data Sources is where every fact the engine reasons about lives. It has four tabs: Integrations (catalog and connect), Live Integrations (sync status and last-pulled values), Business Inputs (the manual intake form), and Metrics (computed daily metrics derived from inputs and connectors).

STEPS

01 Open Data Sources → Business Inputs and fill in fields you can answer with reasonable confidence. Skip the rest - the engine handles partial data.

02 Switch to the Integrations tab to wire up live data sources, then check Live Integrations to confirm the most recent sync.

03 Open Metrics to see the computed daily metric series the engine uses (revenue, leads, conversion rate, response time, etc.).

> **TIP**
> Connector evidence is timestamped. If a sync is stale, the diagnostic call-out tells you which integration to refresh before re-running.

SECTION 06 DIAGNOSE

Connections & integrations

The portal connects to a curated set of revenue, ad, CRM, and accounting platforms. Each connector contributes weighted evidence to the diagnostic and, where applicable, feeds the auto-grade loop. The current catalog is below; connectors marked "Connected - last synced ..." are live for evidence.

Stripe
payments

Shopify
ecommerce

Google Analytics 4
analytics

Meta Ads
ads

Google Ads
ads

HubSpot
crm

GoHighLevel
crm

QuickBooks
accounting

STEPS

01 Open Data Sources → Integrations and click "Connect" next to the platform you want to enable.

02 Complete the OAuth or API-key flow. Connector credentials are encrypted at rest server-side via AES-GCM with a key derived (scrypt) from the CONNECTIONS_ENCRYPTION_KEY secret.

03 Check Data Sources → Live Integrations to confirm the connector shows "Connected - last synced ..." before relying on it.

HEADS UP

Disconnecting an integration revokes its evidence weight in the next diagnostic run. Saved historical reviews are unaffected - they remain pinned to the engine version and evidence available at save time.

SECTION 07 DIAGNOSE

Diagnostic - constraint, transparency, sensitivity, summary

The Diagnostic destination is where the constraint engine runs and explains itself. Primary Constraint shows the chosen lever with ranked alternatives. Transparency exposes the input weights and evidence sources behind the choice. Sensitivity shows how the ranking moves when individual inputs are perturbed. Decision Summary is the consultant-friendly write-up.

STEPS

01 Open Diagnostic → Primary Constraint to see the ranked candidate constraints and the chosen primary.

02 Open Transparency to read the per-candidate evidence weights and the inputs / connectors that drove the score.

03 Open Sensitivity to inspect which inputs are pivotal vs robust before taking the recommendation to a client.

04 Open Decision Summary for the prose write-up; this is the surface most consultants screenshot or paste into a deck.

> **TIP**
> The Evidence panel under each candidate is the single most useful surface for explaining the diagnostic to a client - it cites the exact metrics and connectors that moved the score.

SECTION 08 PREDICT

Forecasts - projected outcome, financial impact, scenarios

Forecasts contains the engine's revenue prediction, the dollarized financial impact view, and the Scenario Simulator. Predictions saved here are what the calibration loop later grades against actual results.

STEPS

01 Open Forecasts → Projected Outcome to see the predicted revenue range and the engine's confidence score.

02 Switch to Financial Impact to translate the range into cumulative dollars over the chosen horizon.

03 Open Scenario Simulator to flex assumptions (lead volume, conversion, AOV) and see the projected outcome update live.

04 Click "Save prediction" from Projected Outcome to lock in the current forecast for later grading.

SECTION 09 VERIFY

90-Day Plan - moves, progress, outcome validation

The 90-Day Plan destination owns the execution and verification half of the loop. Game Plan shows the ranked moves generated for the chosen constraint. Progress is the weekly check-in surface and feeds the blocker list. Outcome Validation is where measured results meet saved predictions.

STEPS

01 Open 90-Day Plan → Game Plan to see the ordered list of moves with effort, impact, and certainty signals.

02 Click any move to open its Playbook - scripts, email templates, checklists, and SMS copy tailored to the constraint and industry.

03 Use Progress to log weekly status on each move; blockers surface back into Overview.

04 Open Outcome Validation to compare measured results (manual or auto-graded from Stripe) against saved predictions.

> NOTE
> Playbooks are deterministic and bundled with the portal - they do not call any AI service.

SECTION 10 VERIFY

Outcomes & auto-grading from Stripe

Outcome Validation is where you record (or let the system record) what actually happened so the engine can grade its own predictions. If a workspace has Stripe connected, eligible predictions are graded automatically against measured revenue.

STEPS

01 Open 90-Day Plan → Outcome Validation to see graded and pending predictions.

02 For workspaces without a live revenue integration, log the actual value manually using "Record outcome".

03 For Stripe-connected workspaces, the auto-grade loop runs in the background and writes a result row when enough graded data exists.

> **HEADS UP**
> Auto-grading only triggers for predictions that include a measurable revenue target and a time horizon. Soft predictions (e.g. "improve response time") still need a manual outcome entry.

SECTION 11 VERIFY

Calibration & accuracy

Calibration aggregates every graded prediction in the workspace into a single accuracy view: how often the actual result fell inside the predicted range, the systematic bias (if any), and the per-tenant conformal quantile used to widen or tighten future intervals.

STEPS

01 Calibration shows the running accuracy summary across saved predictions.

02 Filter by engine version to compare accuracy across releases.

03 Read the bias multiplier - values above 1.0 mean the engine has been under-forecasting; below 1.0 means over-forecasting.

SECTION 12 PREDICT

Personalized Review Narrative

A Personalized Review Narrative is an AI-generated prose layer that wraps the constraint rationale and game plan in a tone matched to the client's industry. It renders in-line on Decision Summary and never replaces the deterministic engine output.

STEPS

01 Open Diagnostic → Decision Summary.

02 Click "Personalize narrative" to expand the prose layer above the engine output.

03 Toggle it off if you prefer the raw engine view (your preference is remembered per workspace).

> **NOTE**
> Narrative generation respects the per-workspace daily token cap and the AI_NARRATIVE_DISABLED kill switch. When either trips, the surface gracefully falls back to the deterministic write-up.

SECTION 13 PREDICT

Client Update - AI plain-English summary

Client Update generates a plain-English status summary the consultant can paste into an email or QBR deck. It is served by the API server's /client-update/generate route, which streams completions from OpenAI via the Replit AI Integrations proxy. It is a separate AI feature from Personalized Review Narrative (which uses Anthropic Claude) and does not share the narrative token cap.

STEPS

01 Open the active review and click "Generate Client Update".

02 Edit the produced text inline before sending - the portal never sends emails on your behalf.

03 Re-generate if you want a different draft; the call requires a paid workspace and an authenticated session.

SECTION 14 DIAGNOSE

Analyze My Data - one-click optimizer insights

Analyze My Data runs the optimizer over the most recent connected dataset and surfaces what changed since the last review: new constraint candidates, shifted confidence, fresh integration evidence.

STEPS

01 Open Overview or Diagnostic.

02 Click "Analyze my data" in the top-right action bar.

03 Read the "What changed" panel - it lists the deltas in plain English with links straight to the affected sections.

SECTION 15 VERIFY

Sharing reviews with clients

Reviews can be shared with a client via a tokenized share link or via the in-portal QBR recap. Share links are read-only, expire on a configurable date, and respect the audience layer (Client / Consultant) you choose.

STEPS

01 From a saved review, click "Share".

02 Pick the audience layer, set an expiry, and copy the link.

03 Open the QBR recap from the floating callout that appears after a snapshot save to walk the client through the period-over-period change.

> **TIP**
> Share links open in a stripped-down read-only shell - clients never see the consultant sidebar, admin tools, or other workspaces.

SECTION 16 ACCOUNT

Messages, ASK, and the AI Assistant

Workspace owners can message each other inside the portal (no email round-trip). The Messages route stays alive for deep links, but its front door is the floating ASK button bottom-right and the notifications area in the top bar - not a sidebar row. The AI Assistant lives in the same surface and answers questions about the active review using the workspace's own data.

STEPS

01 Click the floating ASK button (bottom-right) on any portal page.

02 Pick a conversation, or start a new one with another workspace owner who has shared a review with you.

03 Type "@assistant ..." to ask a question about the current review - the assistant cites the section it pulled the answer from.

> **NOTE**
> Messages does not appear in the sidebar by design - it is hiddenFromNav in the IA model. Use ASK, the command palette (Cmd+K), or notifications to reach it.

SECTION 17 ACCOUNT

Settings - account, automation, referrals, billing, security, audit

Settings consolidates every account-level surface into one tabbed destination. It contains: Account Settings (profile and workspace), Automation Policy (auto-grade and AI behaviour), Referrals (the consultant referral program), Billing (Stripe-managed subscription), Security & Trust (formerly the /trust page), and Audit Log (formerly /audit-log).

STEPS

01 Open Settings → Account Settings to update profile and workspace details.

02 Open Settings → Automation Policy to tune the auto-grade loop and AI feature behaviour.

03 Open Settings → Referrals to copy your unique referral link and track accepted referrals, billed revenue, and pending payouts.

04 Open Settings → Billing for the Stripe-managed subscription portal.

05 Open Settings → Security & Trust for the trust posture overview, and Settings → Audit Log for the chronological change log.

> NOTE
>
> The /trust and /audit-log routes still resolve directly for procurement deep-links and SOC2 audit trails, even though the sidebar surfaces them as Settings tabs.

SECTION 18 ACCOUNT

Audit log & rollback history

Two related-but-separate surfaces. Settings → Audit Log is a read-only chronological feed of workspace changes (saves, edits, integration changes, role changes, deletions) with filters and CSV export. Rollback of executed actions lives in the dedicated admin page at /admin/rollback-history, where workspace admins can review and revert eligible actions.

STEPS

01 Open Settings → Audit Log to browse the chronological change feed.

02 Filter by actor, resource type, or date range, then use Download CSV to export the filtered view (the same filters are applied to the export).

03 For rollbacks, navigate to /admin/rollback-history (admin-only). Each eligible action exposes review and revert affordances; status transitions are tracked as pending / executed / reverted / failed.

> **NOTE**
> Audit Log itself is intentionally read-only - it does not contain a "Revert" button per row. Reverts always happen from the rollback admin page so the action and its outcome are themselves audited.

SECTION 19 ACCOUNT

Admin & Clients (cross-artifact)

Admin surfaces are visible only to workspace owners and staff. The Clients sidebar entry is admin-only and opens the standalone Customer Dashboard artifact (a separate app at /customers/), not an in-portal page. Activation funnels, page-view funnels, and external API key issuance all live in that dashboard.

STEPS

01 Click Clients in the sidebar (admin only) - the portal performs a full document load to /customers/.

02 Use the Customer Dashboard for activation analytics, page-view funnels, and external API key management.

03 Return to the portal via the workspace switcher or by navigating back.

SECTION 20 REFERENCE

Troubleshooting & where to get help

Most issues fall into three buckets: integration sync errors, AI feature limits, and stale cached data. The portal surfaces the actionable error in-line; this section explains the most common ones.

STEPS

01 Integration shows "Disconnected": re-run the OAuth flow from Data Sources → Integrations. Saved reviews remain intact.

02 Client Update / Narrative greyed out: your workspace has hit the daily AI token cap or AI features are disabled. The button tooltip shows which.

03 A page looks empty after a recent change: hard-refresh (Cmd+Shift+R / Ctrl+Shift+R). The portal sets short cache headers in development to make this rare.

04 Still stuck: open ASK and ask the in-portal AI Assistant, or email support from the help center on the marketing site.

SECTION 21 REFERENCE

Glossary

Short definitions for the terms used throughout the portal.

STEPS

01 Constraint - the single growth lever currently limiting revenue. The engine picks one primary constraint per diagnostic run.

02 Evidence - the inputs and integrations that contributed to a constraint score, weighted by source quality.

03 Confidence score - the engine's self-rated certainty in a prediction, expressed 0-100.

04 Likely range - the lower / upper bounds of the predicted revenue lift (90% interval by default).

05 Engine version - the code revision of the constraint-agent and interval engine that produced a saved row. Stamped for auditability.

06 Calibration / conformal quantile - the per-workspace correction factor learned from graded outcomes; used to widen or tighten future intervals.

07 Game plan - the ordered list of moves generated for the active constraint.

08 Progress check-in - a short weekly status update on each plan move; feeds the blocker list and the QBR recap.

09 QBR recap - the period-over-period summary surfaced after a snapshot save and shareable with clients.

10 Workspace - the per-business container for inputs, integrations, predictions, outcomes, and reviews.

END OF GUIDE

That's the portal, end to end.

Need a hand? Open ASK from any page to reach the in-portal AI Assistant, or email support from the help center on the marketing site.