PORTFOLIO INTELLIGENCE PLATFORM

AI Portfolio Rationalization Advisor

Centaur Governance Model · Q1 2026

User Guide & Reference Documentation

Version 1.0 · May 2026

Prepared by K. Turner, Senior AI Product Manager

About this document

This guide covers all six screens of the AI Portfolio Rationalization Advisor, a Centaur governance model demo built to showcase human-AI decision making in enterprise application portfolio management. It is intended for interviewers, business stakeholders, and technical reviewers.

Access Information

Live URL: ai-advisor-stage.streamlit.app
Screens 1–5: Fully public, no login required
Screen 6 (Semantic CMDB Search): Access key required — contact demo owner

Contents

1 The Centaur Governance Model

The Centaur model is a human-AI collaboration framework for enterprise decision making. Named after the mythological figure combining human and horse, it pairs AI analytical capability with human judgment, accountability, and contextual reasoning.

Core Principle

AI recommends — humans decide. No AI output in this platform is final until reviewed and approved by a designated Application Portfolio Manager (APM). Every recommendation includes an explainability layer showing exactly which signals drove the output.

The Three Governance Layers

Layer	Role	Who Acts
AI Engine	Analyzes 6 signals per application, generates recommendation with confidence score	Automated
APM Review	Validates recommendation against business context, approves or overrides	K. Turner — APM
CIO Escalation	Reviews edge cases, high-stakes decisions, and overrides	Executive Sponsor

> Every decision — approval, override, or escalation — is captured in the Decision Audit Log with timestamp, reviewer identity, and written rationale. This creates a full governance trail.

2 Getting Started

Accessing the Application

The application is hosted on Streamlit Community Cloud and accessible via any modern web browser. No installation is required.

Step	Action	Notes
1	Open your browser and navigate to the app URL	Chrome, Firefox, Edge, or Safari
2	Wait for the app to load (10–20 seconds on first visit)	Free tier may sleep between visits
3	Use the left navigation panel to move between screens	Screens 1–5 require no login
4	For Screen 6, enter the access key when prompted	Contact demo owner for key

Navigation

The left sidebar contains the primary navigation. The currently active screen is highlighted in navy blue. The Validation Queue item displays a red dot indicator when pending decisions require APM review.

Data Mode Indicator

A status badge at the top of the sidebar indicates the current data source:

- Yellow badge: Demo mode — CSV data is loaded locally
- Green badge: Live mode — Connected to ServiceNow PDI

> For interview demonstrations, the app runs in Demo mode with realistic CSV data representing a 15-application enterprise portfolio. All metrics, recommendations, and governance workflows are fully functional in Demo mode.

3 Screen 1 — Portfolio Overview

The Portfolio Overview is the primary dashboard, providing an executive-level summary of the application portfolio and its current rationalization cycle status.

Key Metrics

Metric	Description
At-Risk Apps	Applications flagged for immediate attention based on combined risk signals
Portfolio TCO	Total Cost of Ownership across all 15 applications in the current cycle
Pending Decisions	Applications awaiting APM review before the cycle close date
Cycle Accuracy	Percentage of AI recommendations validated as correct by APM review

Filters

Three multi-select filter bars allow slicing the portfolio by Lifecycle stage, AI Recommendation type, and Business Unit. All filters interact in real time — the application table updates immediately on every selection change.

Application Table

The table lists all applications matching the active filters with the following columns:

- Application
- Business Unit
- Criticality
- Lifecycle
- Annual Cost
- Incidents (12M)
- Tech Debt
- AI Recommendation
- Status

4 Screen 2 — AI Analysis & Signals

This screen provides full explainability for every AI recommendation. Select any application from the dropdown to see the complete signal breakdown that drove the AI output.

Application Summary Card

The top card shows the application's key attributes — business unit, lifecycle stage, annual cost, and incident volume — alongside the AI recommendation badge and confidence score.

Frontier Position

Each application is classified against the Jagged Frontier — a concept describing where AI performs reliably versus where human judgment is essential:

Position	Meaning	Action Required
Inside	Strong data signals, high AI reliability	APM can approve with confidence
Edge	Competing signals reduce reliability	APM should validate carefully
Outside	Insufficient data for reliable recommendation	APM must exercise independent judgment

Signal Breakdown

A horizontal bar chart displays the six signals used by the AI engine, each scored 0–100 with a severity badge (High/Medium/Low). The signals are:

- **Incident trend** — 12-month incident volume and direction
- **Duplicate functionality** — Overlap with other portfolio applications
- **Technical debt index** — Code quality, age, and maintenance burden score
- **Lifecycle status** — Current/Aging/End-of-Life classification
- **12-month incident volume** — Raw incident count over the trailing year
- **Stability inverse** — Inverse of application stability score
- **Replacement candidate exists** — Whether a viable replacement has been identified

> The 'Why this recommendation?' panel on the right provides a plain-English explanation of the top contributing signals, making the AI reasoning accessible to non-technical stakeholders.

5 Screen 3 — Validation Queue

The Validation Queue is the APM's primary workspace. It lists all applications with AI recommendations that have not yet been reviewed in the current cycle.

Reviewer Identity

The reviewer name field at the top of the screen captures the APM identity. This name is stamped on every decision made during the session and recorded in the audit log.

Decision Actions

Each application card presents three action options:

Action	When to Use	Audit Outcome
Approve	AI recommendation is correct and appropriate given business context	Logged as Approved with reviewer identity and timestamp
Override	APM disagrees with AI recommendation — must provide written rationale	Logged as Overridden with rationale captured
Escalate to CIO	Decision requires executive input — high risk, high cost, or ambiguous	Logged as Escalated, flagged for CIO review

> Override requires a written rationale of at least one sentence. This is a deliberate governance control — it forces the APM to articulate their reasoning, creating accountability and enabling future model improvement.

Recommendation Drift Alert

If an AI recommendation has changed since the previous cycle, a drift alert is displayed on the application card. This flags potential model instability and signals that extra scrutiny is warranted before approving.

6 Screen 4 — Decision Audit Log

The Decision Audit Log provides a complete, immutable record of every governance decision made in the current cycle. It is the primary accountability mechanism of the Centaur model.

Log Entries

Each audit entry captures the following:

- Application name and ID
- Decision type (Approved / Overridden / Escalated)
- AI recommendation at time of decision
- Reviewer name and timestamp
- Written rationale (for overrides)
- Confidence score at time of review

> The audit log is session-based in Demo mode. In a production deployment, entries would persist to a database and be exportable for compliance reporting.

Override Analysis

A summary panel shows the override rate for the current cycle — the percentage of AI recommendations that the APM disagreed with. A high override rate indicates model drift or data quality issues requiring investigation.

7 Screen 5 — Governance Dashboard

The Governance Dashboard provides portfolio-level analytics on the rationalization cycle, designed for senior stakeholder reporting and model performance monitoring.

Cycle Summary

The top section shows cycle-level KPIs including total decisions made, approval rate, override rate, escalation rate, and overall cycle accuracy against historical benchmarks.

Recommendation Distribution

A breakdown of AI recommendations across the five categories:

- **Retain** — Application is performing well, no action required
- **Modernize** — Application requires investment to remain viable
- **Retire** — Application should be decommissioned
- **Replace** — Application should be replaced with an alternative
- **Evaluate** — Insufficient data — further assessment required

Frontier Analysis

A chart showing the distribution of applications across Inside, Edge, and Outside frontier positions — providing insight into where the AI model is most and least reliable across the current portfolio.

8 Screen 6 — Semantic CMDB Search

> Screen 6 requires an access key. Contact the demo owner to receive the key before your session. The key is valid for the duration of your browser session.

The Semantic CMDB Search is an AI-powered conversational interface that allows natural language queries against the full portfolio dataset. It is powered by Claude (Anthropic) via the Anthropic API.

Access Gate

On first visit, Screen 6 displays an access key prompt. Enter the key provided by the demo owner. Once authenticated, the gate does not reappear for the remainder of your browser session.

Query Quota

Each session is limited to 20 AI queries to manage API costs. A quota counter in the top-right corner shows remaining queries. The counter resets when you open a new browser session or clear your browser data.

Suggested Prompts

Eight suggested prompts are displayed when the conversation is empty. Click any prompt to run it immediately. Example queries include:

- Which apps should we retire first and in what order?
- What's the total cost of all End of Life applications?
- Show me all high-criticality apps that are aging or end of life
- Which Finance apps carry the most security risk?
- Explain the dependency risk if we retire LegacyDB
- Compare ERP Core and DataWarehouse — which is a higher priority to modernize?
- What would the annual savings be if we retired all recommended apps?
- Which apps are most likely to have an outage in the next 6 months?

Portfolio Quick Reference

An expandable panel at the bottom of the screen shows a quick reference table of all 15 applications with their key attributes — useful for formulating precise queries.

9 Data Sources & Demo Mode

The application supports two data modes that are detected automatically at startup.

Demo Mode (CSV)

When no ServiceNow credentials are configured or the PDI is unreachable, the application loads from five CSV files representing a realistic 15-application enterprise portfolio. All screens are fully functional in Demo mode.

File	Contents
01_cmdb_ci_business_app.csv	Application inventory with metadata
02_cmdb_rel_ci.csv	Application dependency relationships
03_sn_vul_vulnerable_item.csv	Security vulnerability records
04_business_capability.csv	Business capability mappings
05_incident_summary.csv	12-month incident summary by application

Live Mode (ServiceNow)

When ServiceNow credentials are present in the application secrets, the app attempts to connect to a ServiceNow Personal Developer Instance (PDI) at startup. If successful, the sidebar badge changes to green and live CMDB data is loaded. If the connection fails, the app silently falls back to CSV data.

10 FAQ & Troubleshooting

Q: The app is taking a long time to load.
Streamlit Community Cloud free tier apps sleep after inactivity. The first load after a sleep period takes 15–20 seconds. This is normal — wait for the spinner to complete and the app will load fully.

Q: Screen 6 shows an access key prompt.
Screen 6 requires an access key to control API usage costs. Contact the demo owner (K. Turner) to receive the key. Enter it in the prompt and press Enter.

Q: I've used all 20 queries on Screen 6.
The per-session query limit has been reached. Open a new browser tab or window to start a fresh session with a new quota.

Q: The sidebar shows 'Demo mode · CSV data'.
This is expected behavior. The app is running with realistic CSV data. All screens are fully functional — the demo mode label indicates data source only.

Q: A decision I made isn't showing in the Audit Log.
The audit log is session-based in Demo mode. If you refreshed the browser, session state is reset. In production, decisions would persist to a database.

Q: The AI recommendation seems wrong for a specific application.
This is intentional — some applications are placed at the Edge or Outside frontier where AI reliability is lower. Use Screen 2 to review the signal breakdown and exercise APM judgment on the Validation Queue.

Q: Can I export the portfolio data or decisions?
Export functionality is not included in the Demo version. In a production deployment, CSV export and PDF reporting would be standard features.

For additional questions or to schedule a live walkthrough, contact K. Turner, Senior AI Product Manager.

This document is confidential and intended for interview and evaluation purposes only.