

Enterprise Data Platform

Operational Runbook & DAMA-DMBOK Teaching Companion

Ramachandran Balasubrahmanian — Data Management Leader
Version 1.5 · 14 June 2026 · Classification: Internal – Training Use

Table of Contents

Document Classification: Internal – Training Use **Product:** Enterprise Data Platform (codename *Pipeline Pulse*) **Live environment:** <https://pipeline-pulse-79.lovable.app/> **Built with:** Lovable (frontend) · OpenAI Codex + Claude Cowork (build pipeline) · TanStack Router · Supabase mirror
Document owner: Ramachandran Balasubrahmanian — Data Management Leader **Aligned to:** DAMA-DMBOK, 2nd Edition (DAMA International, 2017)

How to read this runbook. This document serves two audiences at once. As an **operational runbook** it tells you how to run, navigate, and maintain the platform end to end. As a **teaching artifact** it maps every screen to the DAMA Data Management Body of Knowledge so students can see governance theory expressed in working software. Audience callouts mark whose content a passage is for:

- **Operator** — day-to-day users running jobs, checks, and reviews.
 - **Student** — learners connecting the UI to DMBOK concepts.
 - **Admin** — platform owners configuring access, tenants, and integrations.
-

List of Figures

- Figure 2.1 — Platform architecture, captured live from the Architecture Viewer
- Figure 3.1 — First-time navigation: the sidebar's five groups and the Overview canvas
- Figure 4.1 — Overview command centre (animated tour, top → freshness scoreboard)
- Figure 4.2 — Executive Governance dashboard (animated tour)
- Figure 4.3 — Value Realization (animated tour)
- Figure 4.4 — DAMA-DMBOK Control Tower (animated tour)
- Figure 4.5 — DAMA Evidence Hub (animated tour)
- Figure 4.6 — Demo Video page
- Figure 4.7 — Pipeline Maps lineage graph
- Figure 4.8 — Architecture Viewer (animated tour)
- Figure 4.9 — Projects board
- Figure 4.10 — Data Product Marketplace
- Figure 4.11 — Job Manager
- Figure 4.12 — Data Contract & Schema Registry
- Figure 4.13 — Data Reliability Engineering
- Figure 4.14 — Data Catalog & Business Glossary
- Figure 4.15 — Data Modeling Studio
- Figure 4.16 — Master & Reference Data Hub
- Figure 4.17 — DQ Control Center (full animated workflow — Feed → Metadata Rules → Validate → Results → Reconciliation → Analytics)
- Figure 4.18 — Data Governance Control Center
- Figure 4.19 — AI Governance Control Center
- Figure 4.20 — Data Ethics & Responsible Use Board
- Figure 4.21 — Document & Content Governance
- Figure 4.22 — Data Lifecycle & Retention Manager
- Figure 4.23 — Stewardship Workbench
- Figure 4.24 — Access Control (RBAC + ABAC)
- Figure 4.25 — Billing & Usage
- Figure 4.26 — Incidents
- Figure 4.27 — Audit Trail
- Figure 4.28 — Tenants directory
- Figure 4.29 — Settings

Section 1 — Document Control

1.1 Cover & Ownership

Field	Value
Document title	Enterprise Data Platform — Operational Runbook & DAMA-DMBOK Teaching Companion
Document owner	Ramachandran Balasubrahmanian, Data Management Leader
Classification	Internal – Training Use
Version	1.5 (current)
Status	Released
Date of issue	14 June 2026
Review cadence	Quarterly, aligned to DAMA review cycle
Source environment	https://pipeline-pulse-79.lovable.app/
Capture method	Live screen capture of the running application (29 pages), recorded as animated GIFs with static frames extracted for print

1.2 Distribution List

Recipient group	Purpose
Data Management Office (CDO org)	Ownership, governance, sign-off
Data Stewards & Custodians	Day-to-day operation
Data Engineering & DataOps	Pipeline maintenance & reliability
Internal Audit & Compliance	Evidence and control review
Data Management students / trainees	Teaching and certification preparation

1.3 Version History

The platform evolved across an initial build and five formal reviews from March to June 2026. Each row records what changed and which DAMA-DMBOK knowledge area was most affected.

Version	Date	Author / Reviewer	Summary of changes	Sections affected	DMBOK area impacted
1.0	02 Mar 2026	R. Balasubrahmanian (Author)	Initial build: Overview command centre, Job Manager, Projects, Pipeline Maps, Architecture. Core DataOps surface established.	2, 3, 4.1, 4.7–4.11	Data Storage & Operations; Data Architecture

1.1	24 Mar 2026	L. Mendes (DQ Lead)	Added Metadata-Driven Data Quality Control Center with rule engine, reconciliation, and quarantine workflow.	4.17, 5	Data Quality; Metadata Management
1.2	16 Apr 2026	S. Khan (Integration Lead)	Lineage enhancements (AI impact analysis), Data Contract & Schema Registry , breaking-change prevention.	4.7, 4.12	Data Integration & Interoperability
1.3	11 May 2026	T. Brooks (CISO) / G. Patel (Steward)	Governance tokenization/masking, Access (RBAC/ABAC) , Audit Trail anomaly detection, Evidence Hub .	4.5, 4.18, 4.24, 4.27	Data Security; Data Governance
1.4	29 May 2026	N. Osei (AI Gov) / Ethics Board	AI Governance model registry, Data Ethics Board , MDM golden records, Catalog & Glossary harvesting.	4.14, 4.16, 4.19, 4.20	Reference & Master Data; AI & Big Data Governance
1.5	14 Jun 2026	R. Balasubrahmanian (Owner)	DAMA Control Tower , Value Realization , Reliability SLO/SLI, Lifecycle & Retention , Document & Content. Full DAMBOK coverage map finalised; runbook published.	4.3, 4.4, 4.13, 4.21, 4.22, 6	All 11 knowledge areas (consolidation)

DMBOK Principle Spotlight — Data management is lifecycle management. Notice how the version history itself is a governance artifact: it records *who* changed *what*, *when*, and *why*, mapped to a knowledge area. DAMBOK treats change control and traceability as first-class governance practices, not afterthoughts.

Section 2 — Executive Summary & Architecture

2.1 Purpose of the Platform

The Enterprise Data Platform is a **DAMA-DMBOK-aligned control plane** for an enterprise data estate. It unifies, in a single web application, the operational and governance surfaces that are usually scattered across a dozen tools: pipeline operations, data quality, lineage, metadata and cataloguing, master/reference data, AI model risk, reliability engineering, security and access, retention, and executive evidence. The running system describes itself as a *“DAMA DM-BOK2 aligned control plane for governance, quality, lineage, AI risk, reliability, stewardship, and executive evidence.”*

All data in the environment is **synthetic demo / portfolio-simulation data** — figures illustrate the operating model, not real financials or customers.

2.2 Intended Users & Business Value

User	What they get	Business value
Operators / Stewards	One place to run quality checks, work queues, trace lineage, resolve incidents	Faster cycle times; fewer escalations
Admins / Platform owners	Tenants, access, integrations, billing, settings	Controlled, auditable operations
Executives (CDO/CIO/VP)	Maturity, risk, value, and quarterly priorities	Evidence-based investment decisions
Students	A working reference implementation of DMBOK	Concepts made concrete

Representative value signals reported in-app (synthetic): **+\$2.4M total realized value, 95% incident reduction, 70% faster tenant onboarding, 99.7% PII protection rate, 92% audit readiness.**

2.3 High-Level Architecture




Figure 2.1 — Architecture Viewer: Data Flow, C4 Container, and Semantic Layer blueprints rendered live in the platform.

Figure 2.1 — Platform architecture, captured live from the Architecture Viewer. Alt text: Three stacked architecture diagrams — a left-to-right Data Flow (Source → Extract → Transform → Load → BI/Analytics), a C4 Container diagram inside a production VPC, and a Semantic Layer blueprint.

The application is a single-page React frontend (built in Lovable) served over HTTPS with JWT auth, talking to an Express API gateway that fronts an execution engine, a fleet of 24 compute nodes, a Postgres metadata store, and S3 object storage. Identity is issued as short-lived tokens via a Firebase-style auth layer with RBAC. The build pipeline uses OpenAI Codex and Claude Cowork; a Supabase mirror provides shared persistence.

Technology stack at a glance

Layer	Technology
Frontend	React SPA generated in Lovable ; TanStack Router; client-side route tree (29 routes)
Build pipeline	OpenAI Codex + Claude Cowork
API / services	Express API gateway → Execution engine
Compute	24 compute nodes (auto-scaled)
Metadata store	Postgres
Object storage	Amazon S3 (lakehouse tiers)
Auth & identity	Firebase-style auth, RBAC + ABAC, short-lived JWTs
Persistence mirror	Supabase
External integrations	Slack, PagerDuty, ServiceNow, Datadog, GitHub, Okta, Snowflake, BigQuery

2.4 DMBOK Coverage Map

The eleven DAMA-DMBOK knowledge areas and the platform components that address each. (Maturity levels in parentheses are the in-app self-assessed scores from the DAMA Control Tower, on a 1–5 scale.)

#	DMBOK Knowledge Area	Primary components in the platform	In-app maturity
1	Data Governance	DAMA Control Tower, Executive Governance, Data Governance (tokenization/policy), Evidence Hub, Stewardship, Audit Trail	L4
2	Data Architecture	Architecture Viewer (Data Flow, C4, Semantic Layer), Pipeline Maps	L4
3	Data Modeling & Design	Modeling Studio (Conceptual / Logical / Physical / Standards)	L3
4	Data Storage & Operations	Job Manager, Projects, Reliability, Billing & Usage, Incidents	L4
5	Data Security	Access (RBAC/ABAC), Governance (masking/tokenization), Audit Trail	L5
6	Data Integration & Interoperability	Pipeline Maps (lineage), Data Contract & Schema Registry, Settings (integrations)	L4
7	Document & Content Management	Document & Content Governance (repos, PII scanner, OCR, legal hold)	L2
8	Reference & Master Data	MDM & Reference Hub (golden records, match/merge, survivorship, code sets)	L3
9	Data Warehousing & BI	Overview, Value Realization, Data Product Marketplace	L4
10	Metadata Management	Catalog & Glossary, DQ Metadata Rules, Pipeline Maps inventory	L3
11	Data Quality	DQ Control Center, Reliability, Stewardship queues	L4
+	<i>Extensions: AI & Big Data Governance, Data Ethics, Lifecycle &</i>	AI Governance, Data Ethics Board, Lifecycle & Retention, DAMA	L2–L3

	Retention, Data Management Maturity	Control Tower	
--	--	---------------	--

DMBOK Principle Spotlight — The DAMA Wheel. DMBOK arranges the knowledge areas around **Data Governance at the hub**, with the other ten areas as spokes. This platform is laid out the same way: the *Executive Layer* and *DAMA Control Tower* sit at the centre, and the operational spokes (quality, lineage, security, MDM, metadata...) radiate outward. See Section 6 for the full DAMA Wheel discussion.

Section 3 — Getting Started / Environment Setup

3.1 Prerequisites

Requirement	Detail
Browser	Any modern desktop browser (Chrome, Edge, Safari, Firefox)
URL	https://pipeline-pulse-79.lovable.app/
Account	Demo personas are pre-provisioned; no sign-up required for the training environment
Data	All data is synthetic — safe for classroom use
Screen	1280 px width or wider recommended (the layout is a fixed left sidebar + content canvas)

3.2 Access, Roles & Permissions

The platform ships with five demo personas spanning the role spectrum. Persona and tenant can be switched from the **user menu at the bottom-left of the sidebar** (avatar → *Switch persona / Switch tenant*).

Persona	Role	Clearance	Typical use
Ram Balasubrahmanian	Platform Admin	Restricted	Full admin, all tenants
Gita Patel	Data Steward	Confidential	Stewardship, glossary, approvals
Liam Mendes	Data Engineer	Confidential	Pipelines, jobs, reliability
Aisha Rivera	Business Analyst	Internal	Reporting, products, value
Thomas Brooks	Auditor	Restricted	Evidence, audit trail, read-only review

Admin note. Roles drive what each persona can see and do; the **Access (RBAC/ABAC)** screen (Section 4.24) exposes the role-permission matrix and a live policy-decision simulator. Switching persona is the fastest way to demonstrate least-privilege access in a class.

3.3 First-Time Navigation

The left sidebar groups all 29 screens into five collapsible sections. Learning this map is the single most useful orientation step.




Figure 3.1 — The Overview command centre, showing the five-group navigation sidebar and the live KPI canvas.

Figure 3.1 — First-time navigation: the sidebar’s five groups and the Overview canvas. *Alt text: Application shell with a dark left sidebar listing grouped navigation links and a light content area showing KPI cards and charts.*

The five navigation groups

Group	Screens
Executive Layer	Overview · Executive Governance · Value Realization · DAMA Control Tower · DAMA Evidence Hub · Demo Video
Data Platform	Pipeline Maps · Architecture · Projects · Data Product Marketplace · Job Manager · Data Contract Registry
Data Governance & Quality	Data Reliability · Catalog & Glossary · Modeling Studio · MDM & Reference · DQ Control Center · Data Governance
AI, Ethics & Compliance	AI Governance · Data Ethics Board · Document & Content · Lifecycle & Retention · Stewardship
Operations & Admin	Access (RBAC/ABAC) · Billing & Usage · Incidents · Audit Trail · Tenants · Settings

Common top-bar controls (present on every screen): global search (“Search jobs, projects, or incidents...”), the active **Tenant** chip, a **Notifications** bell, and **Settings**. A floating **Demo Guide** button (bottom-right) launches an in-app walkthrough.

Recommended first session (5 steps)

1. Open the **Overview** to read platform health at a glance.
2. Open the **DAMA Control Tower** to see knowledge-area maturity.
3. Open the **DQ Control Center** and click *Run Data Quality Validation* to watch the rule engine work end to end.

4. Open **Pipeline Maps** and click a node for AI impact analysis.
 5. Switch persona to *Auditor* and open the **Evidence Hub** to see read-only, audit-ready evidence.
-

Section 4 — Component-by-Component Walkthrough

Every screen below follows the same seven-part structure: **screenshot/GIF + annotations**, **Purpose**, **Why it is needed**, **Step-by-step usage**, **DMBOK mapping**, **Student Insight**, and **Troubleshooting**. Figures are animated GIFs in the HTML and Markdown editions; the Word/PDF editions show a representative frame with a link to the live page.

4.1 Overview — Command Centre (/)

Nav group: Executive Layer · **Audience:** Operator · Executives




Figure 4.1 — Overview command centre: KPI cards, Job Execution Trend, System Pulse, and the Data Freshness Scoreboard.

Figure 4.1 — Overview command centre (animated tour, top → freshness scoreboard). *Alt text: Dashboard with four KPI cards, a 30-day completed-vs-failed line chart, a system-pulse panel, and a per-source freshness table colour-coded healthy/warning/breached.*

Annotations — What you're seeing

Marker	UI element	What it does
1	KPI row	Total Jobs Run 14,823 (+12%), Success Rate 98.2% , Active Pipelines 47 (3 critical), Avg Execution 4.2 min
2	Job Execution Trend	30-day completed-vs-failed line chart
3	System Pulse	Compute Nodes Online 24/24 , Queue Depth 3 waiting , Network I/O 2.4 GB/s
4	Data Freshness Scoreboard	SLA status per source: Salesforce 99.1% (healthy), SAP ERP 97.3%

		(healthy), AWS S3 88.2% (warning), Oracle DB 61% (breached) Breached
5	Visitor tracking + <i>Start Demo</i>	Live session metrics and the guided-demo launcher

Purpose. A real-time command centre that summarises platform health across all tenants and pipelines on one screen.

Why it's needed. Executives and operators need a single, trustworthy “is everything OK?” view. Surfacing a **breached freshness SLA** (Oracle, 61%) on the home page turns a latent data-trust problem into a visible, actionable one.

Step-by-step usage. (1) Scan the KPI row for throughput and success rate. (2) Check System Pulse for capacity pressure (queue depth, nodes). (3) Read the Freshness Scoreboard top-to-bottom; anything amber/red is a candidate for the Incidents or Reliability screens. (4) Click *Start Demo* for the guided tour.

DMBOK mapping. *Data Warehousing & BI* (executive dashboards), *Data Quality* (freshness/SLA scoreboard), *Data Governance* (cross-tenant oversight). **Principle (paraphrased):** trustworthy BI depends on timely, monitored data — freshness and SLA adherence are quality dimensions, not just ops metrics.

Student Insight. This screen demonstrates that a dashboard is itself a *data product* with quality obligations. *Reflection: which two freshness states here would you escalate first, and to which other screen in this platform would you route each?*

Troubleshooting. Numbers look static → it is synthetic demo data refreshed per session. A breached source → open **Data Reliability** (4.13) for the SLO/error-budget detail and **Incidents** (4.26) for any linked RCA.

4.2 Executive Governance Dashboard (/executive)

Nav group: Executive Layer · **Audience:** Executives (CDO/CIO/VP)




Figure 4.2 — Executive Governance: top enterprise risks, next-quarter actions, and the Governance ROI summary.

Figure 4.2 — Executive Governance dashboard (animated tour). Alt text: Risk list with High/Medium severity chips, a numbered next-quarter action list, and an ROI summary totalling \$2.4M.

Annotations — What you’re seeing

Marker	UI element	What it does
1	Top 5 Enterprise Risks	Schema drift (High), Unreviewed high-risk AI models (High), Duplicate customer records (Medium), Catalog adoption under target (Medium), Unstructured PII sprawl (High) — each with a mitigation
2	Top 5 Actions — Next Quarter	Approve DAMA roadmap, stand up AI Ethics Board, promote MDM to prod, fund catalog adoption, finish zero-trust onboarding
3	Governance ROI Summary	Avoided incident cost \$1.1M; faster onboarding \$680K; compute optimization \$184K; audit prep saved \$420K; Total \$2.4M

Purpose. A CIO/CDO-level portfolio view of data trust, risk, value, and the quarter’s priorities.

Why it’s needed. Governance competes for funding like any program. Pairing **risks with mitigations** and **value with dollars** lets leaders make defensible investment decisions.

Step-by-step usage. (1) Review each risk and its mitigation owner. (2) Confirm the next-quarter actions align to the risks. (3) Use the ROI summary in budget conversations. (4) Drill into any risk via its related operational screen (e.g., schema drift → Contract Registry).

DMBOK mapping. *Data Governance* (strategy, risk, value). **Principle (paraphrased):** governance must demonstrate business value and manage risk explicitly; a governance program is steered by leadership with measurable outcomes.

Student Insight. Risks here map one-to-one onto operational screens elsewhere in the app — governance is not a slide, it is wired to the controls that mitigate it. *Reflection: pick one risk and name the exact screen that provides the evidence it is being mitigated.*

Troubleshooting. ROI figures are synthetic (“Figures illustrate the operating model — not real financials”). Risks not updating → they are curated portfolio items, not live telemetry.

4.3 Executive Value Realization (/value)

Nav group: Executive Layer · **Audience:** Executives · Student




Figure 4.3 — Value Realization: outcome KPIs, before/after metrics, and risk reduction by DAMA area.

Figure 4.3 — Value Realization (animated tour). Alt text: Eight outcome KPI tiles, a before-vs-after table, a risk-reduction-by-DAMA-area bar list, and a business-value-by-domain table.

Annotations — What you're seeing

Marker	UI element	What it does
1	Outcome KPIs	Incident Reduction 95%, Onboarding Time –70%, Audit Readiness 92%, PII Protection 99.7%, DQ Improvement +34%, Cost Optimization 75%, SLA 99.9%, Reusable Data Products 7
2	Before vs After	e.g., MTTD 47m → 3m, MTTR 6h12m → 11m, Pipeline failures/wk 38 → 2, Cost/1M records \$1.84 → \$0.46

3	Risk Reduction by DAMA Area	Security 62 → 12, Quality 58 → 18, AI Gov 71 → 24, Lifecycle 49 → 17, Metadata 55 → 22
4	Business Value by Domain	Customer +28% campaign ROI, Finance -3-day close, Risk +12% detection, Claims \$4.2M leakage prevented, Privacy 0 consent incidents

Purpose. Translate platform activity into quantified business outcomes by domain and DAMA area.

Why it's needed. It closes the loop between governance effort and measurable value — the evidence base for continued investment.

Step-by-step usage. (1) Read outcome KPIs for headline movement. (2) Use Before/After to show trajectory. (3) Map risk reduction back to the DAMA areas in the Control Tower. (4) Cite domain value in business reviews.

DMBOK mapping. *Data Governance* (value realization), *Metadata Management* and *Data Quality* (the measured dimensions). **Principle (paraphrased):** managed data is an enterprise asset; its value should be measured and reported like any asset's return.

Student Insight. “Before vs After” is the data-management business case in miniature. *Reflection: which single before/after metric would most convince a sceptical CFO, and why?*

Troubleshooting. All metrics synthetic. Values should reconcile with Reliability (MTTR/MTTD) and DQ (improvement %) — note the cross-screen consistency.

4.4 DAMA-DMBOK Control Tower (/dama)

Nav group: Executive Layer · **Audience:** Student · Governance leads




Figure 4.4 — DAMA Control Tower: knowledge-area maturity heatmap and the detailed coverage table.

Figure 4.4 — DAMA-DMBOK Control Tower (animated tour). *Alt text: A grid of knowledge-area tiles colour-graded L1–L5, above a table listing maturity, status, evidence, key risk, next action, owner, and steward per area.*

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Overall Maturity 3.4/5 , Implemented 7, Partial 7, Open Gaps 0, Critical Risks 2, Audit Evidence 47 artefacts
2	Knowledge Area Heatmap	L-levels per area: Security L5 , Governance/Architecture/Storage/Integration/Warehousing/Quality L4 , Modeling/Reference/Metadata/AI L3 , Doc & Content / Ethics L2
3	Coverage table	Per area: maturity bar, status, linked module (e.g., /governance, /dq), evidence, key risk, next action, owner, steward, last reviewed
4	Export controls	<i>Export DAMA Maturity Report and JSON</i>

Purpose. A single board that scores the estate against all DAMA-DMBOK knowledge areas and links each to its owner, evidence, and remediation.

Why it's needed. Maturity assessment (DAMA-DMM/DCAM style) is how organisations know where to invest. The heatmap makes gaps unmissable.

Step-by-step usage. (1) Read the heatmap for weak areas (here: Doc & Content L2, Ethics L2). (2) Open the table row to see the key risk and next action. (3) Note the owner/steward accountable. (4) Export the report for the governance council.

DMBOK mapping. *All eleven knowledge areas, plus Data Management Maturity Assessment.*

Principle (paraphrased): DMBOK is organised as knowledge areas governed from a central hub; maturity is assessed per area to prioritise improvement.

DMBOK Principle Spotlight — Maturity is per-area, not global. A single “3.4/5” hides the story; the heatmap shows Security at L5 but Ethics at L2. Improvement programs target the *lowest, highest-risk* areas first.

Student Insight. This is the DAMA Wheel rendered as an operations board. *Reflection: given two areas at L2 (Document & Content, Ethics), which would you fund first and what evidence would prove progress?*

Troubleshooting. Levels are self-assessed demo scores. The linked module column is the fastest way to jump from “what’s weak” to “where to fix it.”

4.5 DAMA Evidence Hub (/evidence)

Nav group: Executive Layer · **Audience:** Thomas Brooks (Auditor) · Compliance

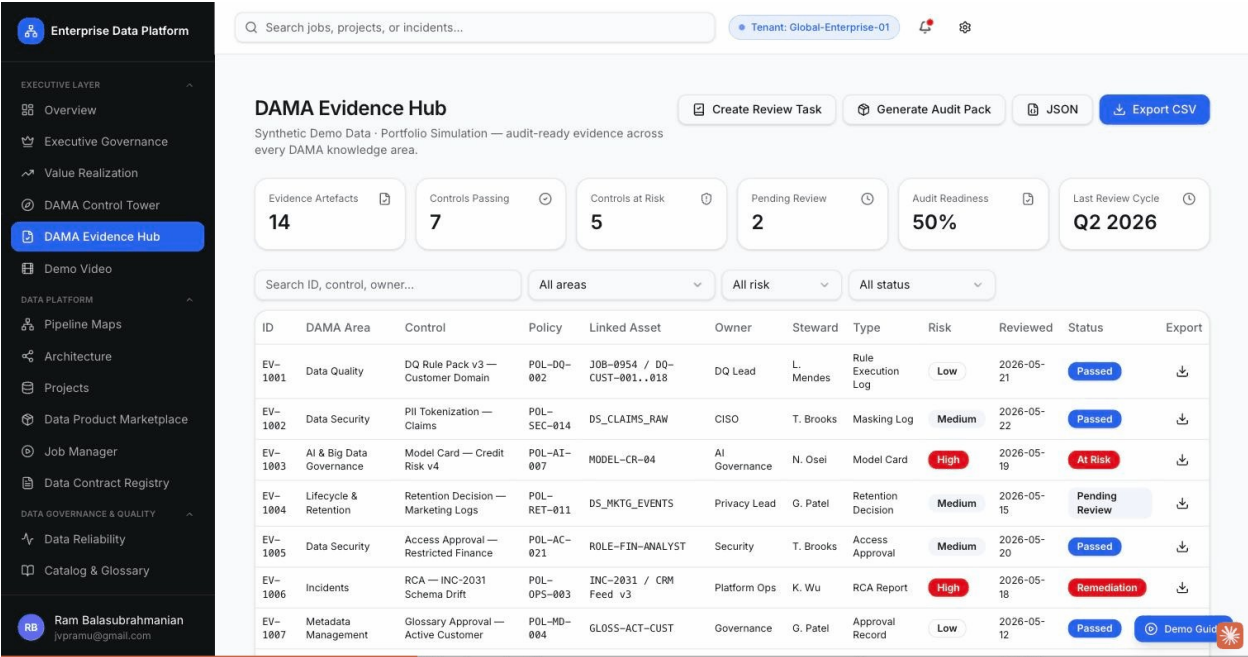


Figure 4.5 — Evidence Hub: audit-ready evidence artefacts mapped to DAMA areas, controls, and status.

Figure 4.5 — DAMA Evidence Hub (animated tour). Alt text: A filterable evidence table (EV-1001...EV-1014) with DAMA area, control, policy, linked asset, owner, steward, type, risk, status, and per-row export.

Annotations — What you’re seeing

Marker	UI element	What it does
1	Summary tiles	Evidence Artefacts 14, Controls Passing 7, Controls at Risk 5, Pending Review 2, Audit Readiness 50%, Last Review Q2 2026
2	Evidence table	EV-1001...EV-1014 mapping each control to a DAMA area, policy ID, linked asset, owner, steward, risk, and status (Passed / At Risk / Pending / Remediation)
3	Actions	Create Review Task, Generate Audit Pack, JSON, Export CSV

Purpose. A defensible, exportable evidence store proving each control operates — organised by DAMA knowledge area.

Why it’s needed. Audits fail on missing or slow evidence. Pre-assembling evidence (and an “Audit Pack”) turns days of scramble into minutes.

Step-by-step usage. (1) Filter by area/risk/status. (2) Open an At Risk item to see the failing control. (3) Create Review Task to assign remediation. (4) Generate Audit Pack / Export CSV for auditors.

DMBOK mapping. *Data Governance* (control evidence, auditability), touching every area it documents. **Principle (paraphrased):** governance must be demonstrable; controls require evidence that they are designed and operating effectively.

Student Insight. Evidence is the bridge between *policy* (intent) and *audit* (proof). *Reflection: for control EV-1003 (AI model card, At Risk), what single action would move it to Passed, and which screen performs it?*

Troubleshooting. Audit readiness at 50% reflects 5 controls at risk + 2 pending — work the queue from the Stewardship screen. Export produces synthetic CSV/JSON.

4.6 Demo Video (/demo)

Nav group: Executive Layer · **Audience:** Student · New users

The screenshot displays the 'Enterprise Data Platform' interface. On the left is a dark sidebar with navigation options: EXECUTIVE LAYER (Overview, Executive Governance, Value Realization, DAMA Control Tower, DAMA Evidence Hub, Demo Video), DATA PLATFORM (Pipeline Maps, Architecture, Projects, Data Product Marketplace, Job Manager, Data Contract Registry), and DATA GOVERNANCE & QUALITY (Data Reliability, Catalog & Glossary). The main area shows a table of controls with columns for ID, Category, Description, Policy ID, Lineage, Owner, Action, Status, Date, and a 'Passed' button. The table lists 14 controls (EV-1002 to EV-1014) with various statuses like 'Passed', 'At Risk', 'Pending Review', and 'Remediation'. A 'Demo Guide' button is visible in the bottom right corner.

ID	Category	Description	Policy ID	Lineage	Owner	Action	Status	Date	Passed	
EV-1002	Data Security	PII Tokenization — Claims	POL-SEC-014	DS_CLAIMS_RAW	CISO	T. Brooks	Masking Log	Medium	2026-05-22	Passed
EV-1003	AI & Big Data Governance	Model Card — Credit Risk v4	POL-AI-007	MODEL-CR-04	AI Governance	N. Osei	Model Card	High	2026-05-19	At Risk
EV-1004	Lifecycle & Retention	Retention Decision — Marketing Logs	POL-RET-011	DS_MKTG_EVENTS	Privacy Lead	G. Patel	Retention Decision	Medium	2026-05-15	Pending Review
EV-1005	Data Security	Access Approval — Restricted Finance	POL-AC-021	ROLE-FIN-ANALYST	Security	T. Brooks	Access Approval	Medium	2026-05-20	Passed
EV-1006	Incidents	RCA — INC-2031 Schema Drift	POL-OPS-003	INC-2031 / CRM Feed v3	Platform Ops	K. Wu	RCA Report	High	2026-05-18	Remediation
EV-1007	Metadata Management	Glossary Approval — Active Customer	POL-MD-004	GLOSS-ACT-CUST	Governance	G. Patel	Approval Record	Low	2026-05-12	Passed
EV-1008	Reference & Master Data	MDM Merge Approval — CUST-000981	POL-MDM-002	MDM-3041	MDM Lead	J. Tan	Merge Approval	Medium	2026-05-22	Passed
EV-1009	Data Architecture	Architecture Review — Lakehouse Tier 2	POL-ARCH-001	ARCH-LH-T2	Chief Architect	M. Lee	Design Review	Low	2026-04-28	Passed
EV-1010	Data Integration	Contract Validation — Billing Events v2	POL-INT-006	CONTRACT-BILL-V2	Integration Lead	S. Khan	Contract Test	Medium	2026-05-10	At Risk
EV-1011	Data Ethics	Responsible Use Review — Churn Model	POL-ETH-001	USECASE-CHURN-02	CDO Office	G. Patel	Ethics Decision	High	2026-05-19	Pending Review
EV-1012	Document & Content Mgmt	Unstructured PII Scan — Claims PDFs	POL-DOC-003	REPO-CLAIMS-PDF	Privacy Lead	G. Patel	Scan Report	High	2026-05-09	Remediation
EV-1013	Data Quality	Reconciliation — GL vs Billing	POL-DQ-005	JOB-0961	Finance Data	R. Costa	Recon Report	Low	2026-05-21	Passed
EV-1014	Data Warehousing & BI	BI Freshness SLA — Exec Dashboard	POL-BI-002	DASH-EXEC-001	Analytics Lead	R. Costa	SLA Evidence	Medium	2026-05-08	At Risk

Figure 4.6 — Demo Video: an embedded 1080p product walkthrough.

Figure 4.6 — Demo Video page. Alt text: A media page with an embedded 1080p MP4 player titled “Pipeline Pulse demo”, with Open-raw and Download controls.

Annotations — What you’re seeing

Marker	UI element	What it does
1	Video card	“Pipeline Pulse demo — 1080p MP4 · AI jobs, lineage, anomaly detection, and monitoring”
2	Controls	Open raw and Download the source file

Purpose. A self-contained, contest-ready walkthrough for onboarding and presentations.

Why it’s needed. Video lowers the barrier for first-time users and gives trainers a ready asset.

Step-by-step usage. (1) Play inline. (2) *Open raw* for full screen. (3) *Download* to embed in a course module.

DMBOK mapping. *Document & Content Management* (managed media asset) and training enablement. **Principle (paraphrased):** content and knowledge assets are governed deliverables with ownership and access, like structured data.

Student Insight. Even a demo video is a content asset with a lifecycle. *Reflection: what retention and access rules should govern an internal product-demo recording?*

Troubleshooting. Video won't autoplay → use the play control; some browsers block autoplay. Use *Open raw* if the embed is blocked by a corporate proxy.

4.7 Pipeline Maps — Lineage (/lineage)

Nav group: Data Platform · **Audience:** Operator · Student




Figure 4.7 — Pipeline Maps: end-to-end lineage graph with topology-health panel.

Figure 4.7 — Pipeline Maps lineage graph. *Alt text: A left-to-right lineage graph of source systems, transform stages, and destinations, beside a topology-health panel and a metadata inventory table.*

Annotations — What you're seeing

Marker	UI element	What it does
1	Lineage Graph	Sources (Salesforce CRM 2.4M/day, SAP ERP 890K/day, AWS S3 1.2M/day) → Transforms (Parallel Extractors, Schema Validation, Spark Clusters, dbt Core, Semantic Layer) → Destinations (BigQuery, Snowflake)
2	Topology Health	12 sync links, 10 healthy, 1 degraded, 1 failed; redundancy 3x; avg latency 120 ms; last full sync 2 min ago

3	Search + node click	“Click any node for AI impact analysis”; search by asset
4	Metadata Inventory	Per-system type and records/day with active status

Purpose. Visualise how data flows source-to-destination and analyse the downstream impact of any node.

Why it's needed. You cannot govern what you cannot trace. Lineage answers “if this breaks, what else breaks?” and underpins impact analysis, root cause, and audit.

Step-by-step usage. (1) Read the graph left-to-right. (2) Check Topology Health for degraded/failed links. (3) Click a node for AI impact analysis. (4) Use the inventory to confirm volumes and status.

DMBOK mapping. *Data Integration & Interoperability* (data movement), *Metadata Management* (lineage is metadata), *Data Architecture*. **Principle (paraphrased):** lineage is core technical metadata; managed integration requires knowing the provenance and flow of every asset.

Student Insight. Lineage is “metadata about movement” and the backbone of trust. *Reflection: the graph shows one failed link — trace which destinations are at risk if the Schema Validation transform fails.*

Troubleshooting. A failed sync link → correlate with Reliability (4.13) and Contract Registry (4.12) for schema drift. Node click does nothing → ensure the graph has finished rendering.

4.8 Architecture Viewer (/architecture)

Nav group: Data Platform · **Audience:** Admin · Architects · Security reviewers




Figure 4.8 — Architecture Viewer: Data Flow, C4 Container, and Semantic Layer diagrams with SVG export.

Figure 4.8 — Architecture Viewer (animated tour). Alt text: Three diagrams — Data Flow pipeline, C4 container architecture inside a production VPC, and a Semantic Layer blueprint mapping raw data to business metrics.

Annotations — What you're seeing

Marker	UI element	What it does
1	Data Flow Diagram	Source → Extract → Transform → Load → BI/Analytics, with SVG download
2	C4 — Container Architecture	Production VPC: React Frontend, Express API Gateway, Execution Engine, Compute Nodes ×24, Postgres Metadata Store, Object Storage (S3); Firebase Auth · RBAC issues short-lived tokens
3	Semantic Layer Blueprint	Raw Data → Metadata Mapping → Business Logic → MRR / Churn Rate / CLV → Enterprise API
4	Export	<i>Export Narrated PDF, Download SVG per diagram</i>

Purpose. System diagrams for architects and security reviewers, exportable for design reviews.

Why it's needed. Architecture is the blueprint that aligns technology to data strategy and exposes trust boundaries for security review.

Step-by-step usage. (1) Use the Data Flow for the conceptual movement story. (2) Use C4 for service/identity boundaries in a security review. (3) Use the Semantic Layer to see how raw data becomes governed business metrics. (4) Export SVG/PDF for documentation.

DMBOK mapping. *Data Architecture* (primary), *Data Security* (trust boundaries), *Metadata Management* (semantic mapping). **Principle (paraphrased):** data architecture aligns data assets with business strategy through standard, shared blueprints.

Student Insight. C4 + Semantic Layer together show the two faces of architecture: *technical* containers and *business* meaning. *Reflection: identify the trust boundary where short-lived tokens are issued, and why it matters for least privilege.*

Troubleshooting. SVG download blocked → check browser pop-up settings. Diagram overflow → the C4 panel scrolls horizontally on narrow screens.

4.9 Projects (/projects)

Nav group: Data Platform · **Audience:** Operator · Liam Mendes (Data Engineer)




Figure 4.9 — Projects: pipeline projects across tenants with status and job counts.

Figure 4.9 — Projects board. *Alt text: A card grid of six pipeline projects, each with a status chip (Active/Paused/Archived), owning tenant, job count, and last-run time.*

Annotations — What you're seeing

Marker	UI element	What it does
1	Project cards	Salesforce → BigQuery Sync (Active, 234 jobs), SAP Financial Consolidation (Active, 89), Marketing Attribution Pipeline (Active, 156), Customer 360 Enrichment (Paused, 42), Real-time Inventory Sync (Active, 312), Legacy Oracle Migration (Archived, 18)
2	Status chips	Active / Paused / Archived
3	<i>New Project</i>	Create a new pipeline project

Purpose. Organise pipelines into named, ownable projects with status and run history per tenant.

Why it's needed. Projects give pipelines an accountable home — ownership, lifecycle state, and a unit for governance and cost.

Step-by-step usage. (1) Scan cards for status. (2) Open a project to view its jobs. (3) Pause/archive obsolete work. (4) Use *New Project* to onboard a pipeline.

DMBOK mapping. *Data Storage & Operations / DataOps.* **Principle (paraphrased):** operational data work must be organised, owned, and lifecycle-managed — not run as anonymous one-off jobs.

Student Insight. “Paused” and “Archived” are lifecycle states — DataOps is lifecycle management for pipelines. *Reflection: what governance checks should gate a project before it moves from Paused to Active in production?*

Troubleshooting. Job counts not changing → demo data. Archived projects are read-only by design.

4.10 Data Product Marketplace (/products)

Nav group: Data Platform · **Audience:** Aisha Rivera (Analyst) · Data consumers




Figure 4.10 — Data Product Marketplace: governed, reusable data products with DQ, sensitivity, and access controls.

Figure 4.10 — Data Product Marketplace. Alt text: A catalogue of data products as cards, each with certification badge, DQ score, freshness, sensitivity, consumer count, cost, and a Request Access button.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Published 7, Certified 6, Active Consumers 100, Avg DQ 95%, Access Requests 7, Monthly Usage 7.1M
2	Product cards	Customer 360 View (Certified, DQ 96%), Credit Risk Feature Set (94%), Claims Fraud Signals (91%), Revenue Performance Mart (98%), Consent Registry Product (99%), Billing Events Product (Pilot, 89%)
3	Per-product detail	Domain, owner, DQ, freshness, sensitivity, consumers, cost, <i>Details / Request Access</i>

Purpose. A governed catalogue of reusable **data products** with ownership, quality, sensitivity, and access controls.

Why it's needed. Data-as-a-product reduces duplication and enforces quality/sensitivity at the point of consumption — the consumer sees certification and DQ before requesting access.

Step-by-step usage. (1) Browse certified products. (2) Open *Details* for DQ/sensitivity. (3) *Request Access* to start the approval flow. (4) Track usage and consumers per product.

DMBOK mapping. *Data Warehousing & BI* (curated marts/products), *Data Governance* (ownership/access), *Metadata Management* (product metadata). **Principle (paraphrased):** treating data as a managed product — with an owner, SLAs, and quality — is how reuse scales safely.

Student Insight. Certification + DQ score + sensitivity on each card is governance-at-point-of-use. *Reflection: why might a “Pilot” product (Billing Events, 89% DQ) be restricted from certain consumers?*

Troubleshooting. *Request Access* won't grant immediately → it routes to an approval queue (see Stewardship / Access). Costs are illustrative (\$/1K reads).

4.11 Job Manager (/jobs)

Nav group: Data Platform · **Audience:** Operator · DataOps




Figure 4.11 — Job Manager: all jobs with status, tenant, hardware, threads, duration, and start time.

Figure 4.11 — Job Manager. Alt text: A filterable jobs table with status chips (Running, Completed, Errored Out, Queueing, Initialising, Cancelled), tenant, hardware, threads, duration, and start time.

Annotations — What you're seeing

Marker	UI element	What it does
1	Header	“38 of 38 jobs visible”; <i>Export CSV, Upload Config, Create Job</i>
2	Filters	Status, tenant, date range, free-text job search
3	Jobs table	Job name + ID, status, tenant,

		hardware (e.g., 32 CPU / 128 GB), threads, duration, started
4	Status variety	Running, Completed, Errored Out, Queueing, Initialising, Cancelled

Purpose. Operate and observe every batch/stream job across tenants.

Why it's needed. Jobs are where data operations actually happen; operators need to see, filter, and act on run state and resource usage.

Step-by-step usage. (1) Filter by status/tenant/date. (2) Inspect an *Errored Out* job. (3) *Create Job* or *Upload Config* to schedule work. (4) *Export CSV* for run reporting.

DMBOK mapping. *Data Storage & Operations*. **Principle (paraphrased):** operations management covers the technology that runs data processing — monitored, resourced, and recoverable.

Student Insight. Hardware + threads per job links operations to cost and capacity. *Reflection: an “Errored Out” job and a “Queueing” job need different responses — describe each.*

Troubleshooting. Errored job → open **Incidents** (4.26) for AI-assisted RCA. Queue building up → check Overview System Pulse and Billing utilisation.

4.12 Data Contract & Schema Registry (/contracts)

Nav group: Data Platform · **Audience:** Operator · Liam Mendes (Engineer) · S. Khan (Integration)

Data Contract & Schema Registry

Synthetic Demo Data - Portfolio Simulation — producer/consumer governance, schema change control, and breaking-change prevention.

Active Contracts: 5 | Breaking Blocked: 1 | Compatibility: 67% | Producers: 6 | Consumers: 6 | Pending Approval: 1

Registry | Schema Drift | Impact Analysis | Test Results | Consumer Notifications

ID	Producer	Consumer	Asset / Version	Compat	SLA	Owner	Approval	Last Validated	Actions
CT-001	CRM (SaaS)	Customer 360	CRM Customer Feed v3 - v3.2.0	Compatible	99.0% / 15m	Integration Lead	Approved	2026-05-22	ⓘ ⚙️ 🚩
CT-002	Billing Stream	Finance Mart	Billing Events Stream v2 - v2.4.1	Compatible	99.95% / 1m	Billing Eng	Approved	2026-05-21	ⓘ ⚙️ 🚩
CT-003	SAP ECC	GL Mart	SAP GL Extract v4 - v4.0.0	Breaking	99.5% / 1h	Finance Data	Pending	2026-05-20	ⓘ ⚙️ 🚩
CT-004	Claims Core	Fraud Signals	Claims Processing Feed v2 - v2.1.3	Compatible	99.9% / 5m	Claims Eng	Approved	2026-05-19	ⓘ ⚙️ 🚩
CT-005	Customer 360	Marketing Activation	Customer 360 Enrichment Contract - v1.5.0	Forward-only	99.5% / 1h	CDO Office	Approved	2026-05-18	ⓘ ⚙️ 🚩
CT-006	Consent Service	All Downstream	Consent Events v1 - v1.2.0	Compatible	99.99% / real-time	Privacy Lead	Approved	2026-05-22	ⓘ ⚙️ 🚩

Figure 4.12 — Contract Registry: producer/consumer contracts, compatibility, and breaking-change prevention.

Figure 4.12 — Data Contract & Schema Registry. Alt text: A summary of contracts and a registry table of producer → consumer contracts with compatibility chips (Compatible / Breaking / Forward-only), SLA, owner, and approval status; tabs for Schema Drift, Impact Analysis, Test Results, Schema Gates, Consumer Notifications.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Active Contracts 5, Breaking Blocked 1, Compatibility 67%, Producers 6, Consumers 6, Pending Approval 1
2	Tabs	Registry · Schema Drift · Impact Analysis · Test Results · Schema Gates · Consumer Notifications
3	Registry table	CT-001...CT-006: producer → consumer, asset/version, compatibility (Compatible/Breaking/Forward-only), SLA, owner, approval, last validated
4	Actions	<i>Run Compatibility, Export Evidence</i>

Purpose. Govern the contracts between data producers and consumers and prevent breaking schema changes from reaching production.

Why it's needed. Most integration outages are silent schema breaks. A contract registry with **schema gates** blocks a breaking change (here, 1 blocked) before it propagates.

Step-by-step usage. (1) Review the registry for compatibility. (2) Open *Schema Drift / Impact Analysis* for a flagged contract. (3) *Run Compatibility* before promoting a change. (4) Notify consumers; *Export Evidence* for the audit trail.

DMBOK mapping. *Data Integration & Interoperability* (primary), *Data Quality* (schema conformance), *Metadata Management* (schemas). **Principle (paraphrased):** interoperability requires explicit, versioned interface agreements and change control between systems.

Student Insight. A data contract is an SLA for *shape and meaning*, enforced automatically. *Reflection: contract CT-003 is "Breaking, Pending" — what are the producer's options, and which consumers does Impact Analysis say are affected?*

Troubleshooting. Compatibility at 67% → one breaking contract pending; resolve via Impact Analysis. A change reached prod unexpectedly → confirm Schema Gates were enabled for that contract.

4.13 Data Reliability Engineering (/reliability)

Nav group: Data Governance & Quality · **Audience:** Operator · SRE/DataOps

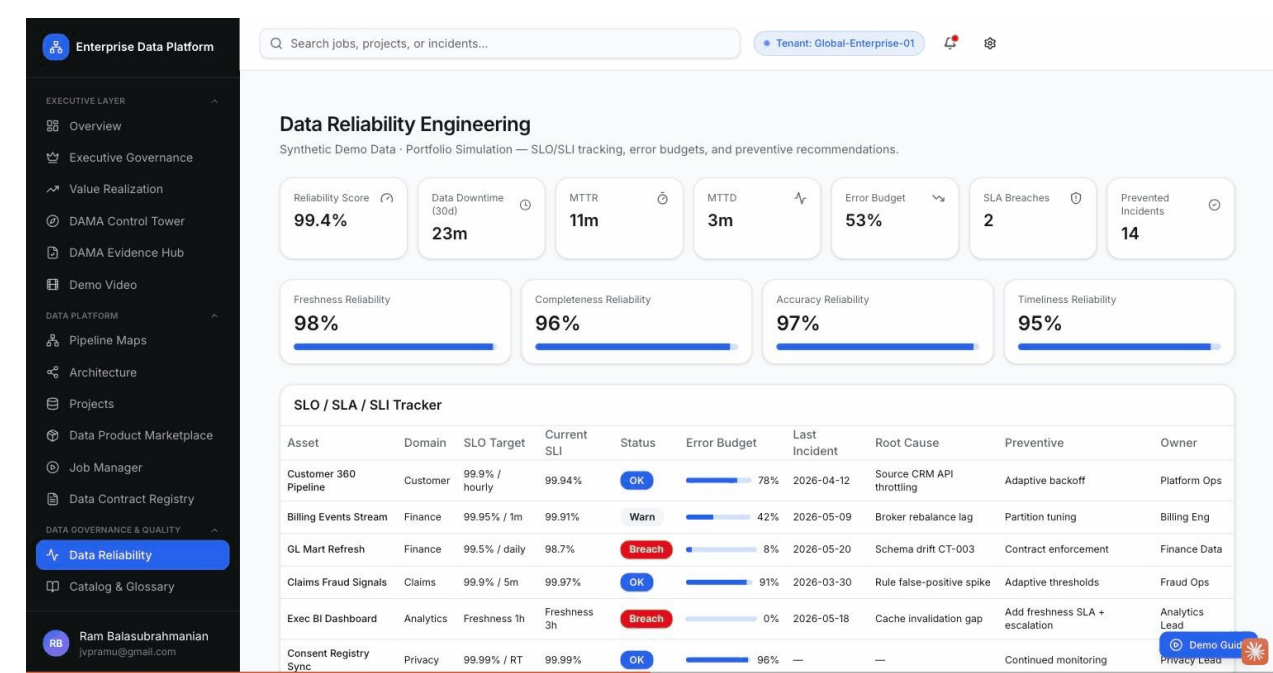


Figure 4.13 — Data Reliability: SLO/SLI tracking, error budgets, and preventive recommendations. **Figure 4.13 — Data Reliability Engineering.** Alt text: Reliability KPI tiles, four reliability-dimension gauges, and an SLO/SLA/SLI tracker table with error budgets and root causes.

Annotations — What you’re seeing

Marker	UI element	What it does
1	KPI tiles	Reliability Score 99.4%, Data Downtime (30d) 23m, MTTR 11m, MTTD 3m, Error Budget 53%, SLA Breaches 2, Prevented Incidents 14
2	Dimension gauges	Freshness 98% · Completeness 96% · Accuracy 97% · Timeliness 95%
3	SLO/SLA/SLI Tracker	Per asset: SLO target, current SLI, status (OK/Warn/Breach), error budget, last incident, root cause, preventive action, owner

Purpose. Apply SRE discipline (SLO/SLI, error budgets) to data, with preventive recommendations.

Why it’s needed. Reliability quantifies trust over time. Error budgets convert “is it reliable?” into a number you can spend and defend.

Step-by-step usage. (1) Read the reliability score and dimension gauges. (2) Find Breach/Warn rows in the tracker. (3) Review the root cause and preventive action. (4) Feed breaches into Incidents and the DAMA risk view.

DMBOK mapping. *Data Quality* (the four dimensions), *Data Storage & Operations* (availability/recovery). **Principle (paraphrased):** quality has measurable dimensions — completeness, accuracy, timeliness, consistency — that must be monitored against targets.

Student Insight. Freshness/Completeness/Accuracy/Timeliness are literally DMBOK's data-quality dimensions, instrumented. *Reflection: a breached “Exec BI Dashboard” SLO has a 0% error budget — what does that mean operationally?*

Troubleshooting. Breach with no incident → create one from the row's preventive action. Error budget exhausted → freeze risky changes to that asset.

4.14 Data Catalog & Business Glossary (/catalog)

Nav group: Data Governance & Quality · **Audience:** Gita Patel (Steward) · Student




Figure 4.14 — Catalog & Glossary: certified assets, glossary terms, completeness scoring, and harvesting.

Figure 4.14 — Data Catalog & Business Glossary. Alt text: Catalog KPIs, a completeness-formula panel, a tag cloud, and an asset table with owner, steward, sensitivity, DQ, policies, lineage links, completeness bar, and certification status; tabs for Assets, Glossary, Approval Workflow, Harvesting.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Catalogued Assets 10, Certified 7, Glossary Terms 10, Avg Completeness 88%
2	Tabs	Assets · Glossary · Approval Workflow · Harvesting
3	Completeness Formula	Weighted score across owner, steward, glossary link, sensitivity, lineage, DQ rules, policy tags, downstream consumers, certification
4	Asset table	customer_master,

		customer_transactions, policy_claims, billing_events, crm_contacts... with owner, steward, sensitivity, DQ, policies, lineage, completeness, status
--	--	---

Purpose. The searchable inventory of certified assets and the business glossary that gives them shared meaning.

Why it's needed. A catalogue makes data findable and trustworthy; a glossary makes terms *mean the same thing* across the enterprise. Together they are the heart of metadata management.

Step-by-step usage. (1) Search/filter assets by domain or sensitivity. (2) Open an asset to see completeness drivers. (3) Use *Glossary* to align definitions. (4) Use *Approval Workflow* to certify; *Harvesting* to ingest technical metadata automatically.

DMBOK mapping. *Metadata Management* (primary), *Data Governance* (certification/approval), *Data Quality* (completeness). **Principle (paraphrased):** metadata is “data about data”; a catalogue and glossary are the core deliverables that make data understandable and governable.

DMBOK Principle Spotlight — Business glossary ≠ data dictionary. The glossary captures *business meaning* (what “Active Customer” means to the business); the catalogue/dictionary captures *technical structure*. The platform links them so a term resolves to real assets.

Student Insight. “Completeness” here is a metadata-quality score, not a data-row score. *Reflection: which two completeness factors would you weight highest for a regulated asset, and why?*

Troubleshooting. Asset shows “Under Review” → finish the Approval Workflow. Low completeness → missing steward/glossary/lineage links are the usual cause.

4.15 Data Modeling Studio (/modeling)

Nav group: Data Governance & Quality · **Audience:** Architects · Student

Data Catalog & Business Glossary
Synthetic Demo Data - Search certified assets, glossary terms, owners, stewards, sensitivity, and lineage links.

Export

Catalogued Assets: 10 | Certified: 7 | Glossary Terms: 10 | Avg Completeness: 88%

Assets | Glossary | Approval Workflow | Harvesting

Completeness Formula
Completeness = weighted score across owner, steward, glossary link, sensitivity, lineage, DQ rules, policy tags, downstream consumers, and certification state.

Tag Cloud
PII, Customer 360, Revenue, AI Input, Certified, Lineage Covered, GDPR, DQ Controlled, Reference, High Value

Search assets... | All domains | All sensitivity

Asset	Domain	Owner	Steward	Sensitivity	DQ	Policies	Lineage	Downstream	Completeness	Status
customer_master	Customer	C. Adler	G. Patel	Confidential	12	6	9	BI, Risk Model, CRM	100%	Certified
customer_transactions	Finance	R. Costa	L. Mendes	Confidential	18	5	14	Revenue DW, Fraud Model	100%	Certified
policy_claims	Claims	S. Khan	A. Rivera	Restricted	22	8	11	Claims DW, Fraud Model	80%	Under Review
billing_events	Billing	K. Wu	L. Mendes	Confidential	9	4	7	Revenue DW	100%	Certified
crm_contacts	Sales	M. Lee	J. Tan	Confidential	7	3	5	Marketing Attribution	100%	Certified

Figure 4.15 — Modeling Studio: conceptual/logical/physical model views with standards validation and change impact.

Figure 4.15 — Data Modeling Studio. Alt text: An entity table (Customer, Account, Transaction, Product, Claim, Policy, Consent, Risk Score, Invoice, Payment) with attribute counts, primary keys, owners, stewards, sensitivity, DQ-rule counts, and downstream counts; tabs for Entities, Conceptual, Logical, Physical, Standards.

Annotations — What you're seeing

Marker	UI element	What it does
1	Tabs	Entities · Conceptual · Logical · Physical · Standards
2	Entity table	e.g., Customer (18 attrs, PK customer_id, Confidential, 12 DQ rules, 9 downstream), Transaction (22 attrs), Claim/Policy/Consent (Restricted)
3	<i>Analyze Change Impact</i>	Assess the blast radius of a model change

Purpose. Author and govern data models across the three classic abstraction levels, validated against naming/standards.

Why it's needed. Models are the shared design language of data. Standards validation prevents the “inconsistent naming” risk flagged in the Control Tower.

Step-by-step usage. (1) Browse entities and keys. (2) Move Conceptual → Logical → Physical to see refinement. (3) Run *Standards* validation. (4) Use *Analyze Change Impact* before altering an entity.

DMBOK mapping. *Data Modeling & Design* (primary), *Metadata Management*, *Data Quality* (rules per entity). **Principle (paraphrased):** modeling progresses conceptual → logical → physical, and naming/design standards are governance controls.

Student Insight. The same entity has different representations at each level — abstraction is deliberate. *Reflection: why does the Customer entity carry both a steward and a sensitivity label at the model level?*

Troubleshooting. Standards failures → usually naming-convention violations; adopt the enterprise naming guide (the Control Tower's stated next action). Change impact empty → ensure downstream links exist in lineage.

4.16 Master & Reference Data Hub (/mdm)

Nav group: Data Governance & Quality · **Audience:** J. Tan (MDM Lead) · Stewards




Figure 4.16 — MDM & Reference: golden records, duplicate detection, survivorship, reference code sets, and steward queues.

Figure 4.16 — Master & Reference Data Hub. *Alt text: MDM KPIs (golden records, duplicates, quality score, steward queue), a tokenized Customer Golden Record, and survivorship rules; tabs for Golden Record, Match/Merge, Reference Data, Steward Queue.*

Annotations — What you’re seeing

Marker	UI element	What it does
1	Summary tiles	Golden Records 12,481 , Duplicates Detected (30d) 317, MDM Quality 92%, Steward Queue 3
2	Tabs	Golden Record · Match/Merge · Reference Data · Steward Queue
3	Golden Record	Surviving customer record with tokenized PII (email/phone/address tokens, masked DOB), segment, consent, source
4	Survivorship Rules	Which source wins per attribute when records merge

Purpose. Create and govern the single, trusted version of master entities and the shared reference code sets.

Why it’s needed. Duplicate and conflicting master data (the “duplicate customer records” enterprise risk) corrupts analytics and customer experience. Golden records + survivorship resolve it deterministically.

Step-by-step usage. (1) Inspect a Golden Record and its tokenized fields. (2) Review *Match/Merge* candidates. (3) Confirm *Survivorship Rules*. (4) Work the *Steward Queue* to approve merges. (5) Manage *Reference Data* code sets.

DMBOK mapping. *Reference & Master Data* (primary), *Data Quality*, *Data Security* (tokenized PII).

Principle (paraphrased): master data management establishes an authoritative “golden” version via matching, merging, and survivorship governed by stewards.

Student Insight. Survivorship rules are governance encoded as logic — they decide truth when sources disagree. *Reflection: for a customer email, would you survive the most-recent or the most-trusted source, and what's the trade-off?*

Troubleshooting. Merge looks wrong → check survivorship rules and steward override. Rising duplicates → tune match thresholds in Match/Merge.

4.17 Metadata-Driven DQ Control Center (/dq) (KEY SCREEN)

Nav group: Data Governance & Quality · **Audience:** Operator · L. Mendes (DQ Lead) · Student




Figure 4.17 — DQ Control Center: an animated walkthrough from Incoming Feed and Metadata Rules through validation to Reconciliation and Analytics.

Figure 4.17 — DQ Control Center (full animated workflow — Feed → Metadata Rules → Validate → Results → Reconciliation → Analytics). Alt text: A multi-tab data-quality console; a metadata rule table; a live validation progress list; and result, reconciliation, and analytics panels showing 93 passed / 6 rejected / 1 quarantined of 100 records.

Annotations — What you're seeing (this screen has 14 tabs; the most important are annotated)

Marker	Tab / element	What it does
1	Incoming Feed	The file under validation: customer_transactions_may.csv from CRM-Core, ruleset DQ_RULESET_V3, 100 records,

		"Ready for Validation"
2	Metadata Rules	DQ-001...DQ-008: Mandatory (customer_id, Critical → Reject), Date Format, Regex (phone), Reference Check (ZIP), Range Check (amount, Critical), Email Format, Duplicate Check (→ Quarantine) , Domain Check
3	Run Data Quality Validation	Streams a 10-step pipeline (read → load rules → schema → mandatory → format → duplicates → status → reconciliation → audit log → rejected report)
4	Validation Results	100 → 93 passed, 6 rejected, 1 quarantined
5	Reconciliation	Pass 93% / Reject 6% / Quarantine 1%; Threshold 10%; Pipeline Decision: Continue
6	Analytics	Error Rate 6%, Pass Rate 93%, Duplicate Confidence 92%, Rules Executed 8; DQ Scorecard by Domain; Four-Week Quality Trend
7	Audit / Run History / Quarantine & Golden	Audit log, prior runs, quarantined records routed toward golden-record handling

Purpose. A reusable, **metadata-driven** quality engine that validates every record *before* it enters the pipeline — rejecting bad data, quarantining risky data, and producing audit-ready reconciliation in minutes.

Why it's needed. It replaces fragmented, manual, per-pipeline checks with one governed engine. As the screen states: *"move from manual validation and fragmented quality checks to a reusable metadata-driven governance engine."* Rules are configuration (metadata), not code, so they apply across feeds.

Step-by-step usage. (1) Confirm the Incoming Feed and ruleset. (2) Review/author rules in *Metadata Rules / Manage Rules* (try *AI Rule Assistant*). (3) Click **Run Data Quality Validation**. (4) Read *Validation Results* for the pass/reject/quarantine split. (5) Open *Reconciliation* for the pipeline decision against the threshold. (6) Use *Analytics* for trend and scorecard; *Audit* and *Run History* for evidence; resolve quarantined rows toward golden records.

DMBOK mapping. *Data Quality* (primary), *Metadata Management* (rules-as-metadata), *Data Governance* (policy enforcement + audit). **Principle (paraphrased):** quality is best enforced by *managed rules applied consistently*, with measurement and remediation, rather than ad-hoc checks — and rules themselves are governed metadata.

DMBOK Principle Spotlight — Prevent, don't inspect-after. Validating *before* the pipeline (reject/quarantine at the gate) embodies DMBOK's "design quality in" stance: cheaper to stop a bad record at ingestion than to chase it through every downstream mart.

Student Insight. "Metadata-driven" is the key phrase: change a rule's metadata and every feed inherits it — governance that scales. *Reflection: rule DQ-007 quarantines (not rejects) duplicates — why is quarantine the right action for suspected duplicates, and where do those records go next?*

Troubleshooting. Tabs missing → the result tabs (Reconciliation, Analytics, Audit...) appear *after* a validation run. Reject rate above threshold → the Reconciliation “Pipeline Decision” flips from Continue to hold; investigate the top failing rule in Analytics. Duplicate confidence low → tune the duplicate rule and re-run.

4.18 Data Governance Control Center (/governance)

Nav group: Data Governance & Quality · **Audience:** T. Brooks (CISO) · Admin




Figure 4.18 — Data Governance: synthetic classification, tokenization/masking, policy enforcement, and audit evidence.

Figure 4.18 — Data Governance Control Center. Alt text: A tokenization-and-masking configuration panel (SHA-256, custom salt, Base64) and a synthetic batch-processing demo with classification and quarantine metrics.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Total Synthetic Records 60, Processed Batches, PII Fields Tokenized, Protection Success Rate, Review Required, Quarantined Records
2	Tokenization & Masking Config	Hashing algorithm (SHA-256), custom salt, protection method Base64(Hash(value + salt)), visible-character policy, “no raw sensitive values stored/logged/exported”
3	Synthetic Batch Processing	Classify and protect a batch, with filters for batch/status/sensitivity
4	Actions	<i>Process Demo Batch, Reset</i>

		Demo
--	--	------

Purpose. Enforce data-protection controls — classification, tokenization, masking, policy — and produce the resulting audit evidence.

Why it's needed. Security and privacy controls must be *operational and provable*. Deterministic tokenization protects PII while preserving join-ability; the “never export raw” rule is policy made concrete.

Step-by-step usage. (1) Choose hashing algorithm and salt. (2) Decide whether protected tokens may be exported. (3) *Process Demo Batch* to classify and tokenize. (4) Review protection success and quarantined records; the run yields audit evidence.

DMBOK mapping. *Data Security* (primary), *Data Governance* (policy enforcement), *Data Quality* (classification). **Principle (paraphrased):** security applies controls — classification, masking/tokenization, and policy — proportionate to data sensitivity, with auditable enforcement.

Student Insight. Tokenization ≠ encryption: tokens are deterministic stand-ins enabling analytics without exposing raw PII. *Reflection: why does keeping the salt secret (and never exporting it) matter for re-identification risk?*

Troubleshooting. Protection rate < 100% → some fields failed classification; review quarantined records. Export disabled → “Include Full Protected Tokens in Export” is off by policy.

4.19 AI Governance Control Center (/ai-governance)

Nav group: AI, Ethics & Compliance · **Audience:** N. Osei (AI Gov) · Risk




Figure 4.19 — AI Governance: model registry, risk tier, PII exposure, drift, fairness, explainability, and human-in-the-loop.

Figure 4.19 — AI Governance Control Center. *Alt text: Model-governance KPIs and a model registry table with use case, type, risk tier, inputs, PII handling, approval, monitoring status, owner, and model-card links; tabs include Registry and Monitoring.*

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Registered Models 7, Critical/High Risk 3, Pending Review 1, Avg Explainability 82%
2	Tabs	Registry · Monitoring (and Overview, Training Data, Evaluation, Bias & Fairness, Governance)
3	Model registry	Customer Churn (Medium), Credit Risk Scoring (Critical), Claims Fraud (High), Revenue Forecasting (Low), PII Discovery Assistant (High) — with PII handling, approval, monitoring, owner
4	Model Card	Per-model documentation artefact

Purpose. Govern machine-learning models like data assets: registry, risk tier, PII exposure, drift/fairness monitoring, explainability, and human oversight.

Why it's needed. Unreviewed high-risk models are a named enterprise risk. A registry with risk tiers and model cards makes AI risk visible and auditable.

Step-by-step usage. (1) Review the registry and risk tiers. (2) Open a Model Card for a high-risk model. (3) Check Bias & Fairness / Evaluation. (4) Confirm approval and human-in-the-loop for Critical models.

DMBOK mapping. *AI & Big Data Governance / Data Science* (a DAMBOK Big-Data extension), *Data Security* (PII in features), *Data Governance*. **Principle (paraphrased):** analytics and AI must be governed — provenance of training data, model documentation, fairness, and human oversight are governance obligations.

Student Insight. A model card is to a model what a catalogue entry is to a dataset — managed metadata. *Reflection: Credit Risk Scoring is "Critical" with masked+tokenized PII — what additional controls would you require before approval?*

Troubleshooting. Pending review blocks deployment by design. Drift/fairness "Watch" → schedule re-evaluation; route to the Ethics Board if fairness is at risk.

4.20 Data Ethics & Responsible Use Board (/ethics)

Nav group: AI, Ethics & Compliance · **Audience:** Ethics Board · CDO Office

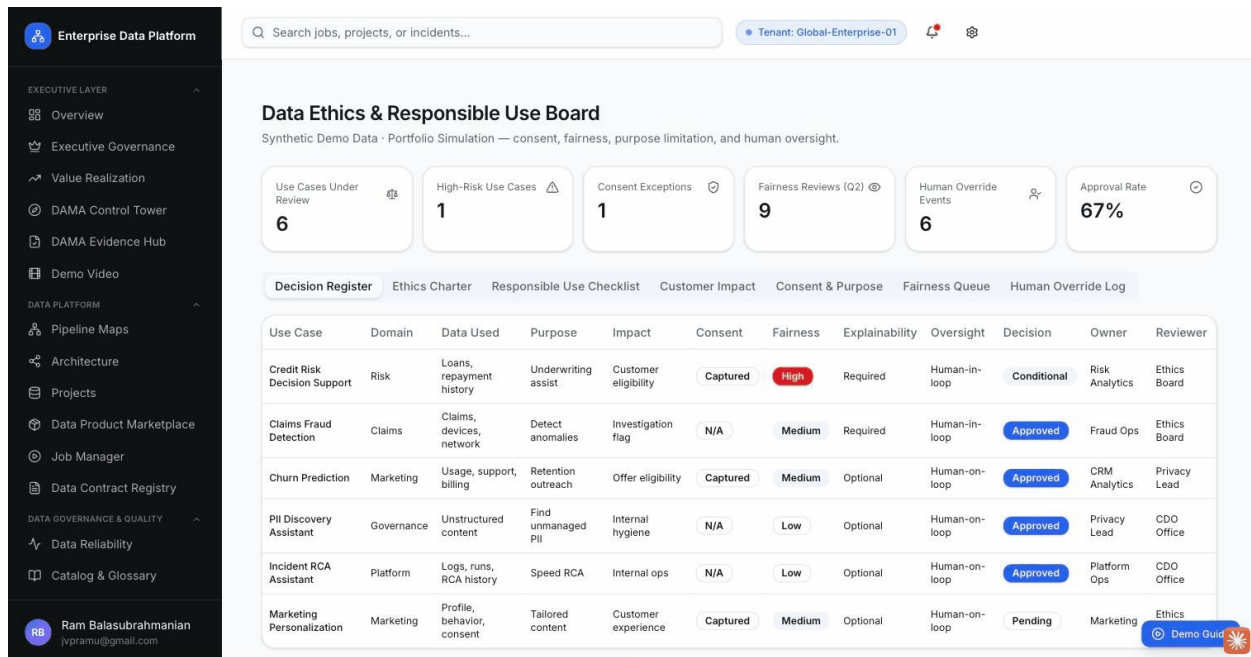


Figure 4.20 — Data Ethics Board: consent, fairness, purpose limitation, and human oversight.
Figure 4.20 — Data Ethics & Responsible Use Board. Alt text: Ethics KPIs and a decision register mapping use cases to data used, purpose, impact, consent, fairness, explainability, oversight, decision, owner, and reviewer; tabs for Ethics Charter, Responsible Use Checklist, Consent & Purpose, Fairness Queue, Human Override Log.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Use Cases Under Review 6, High-Risk 1, Consent Exceptions 1, Fairness Reviews (Q2) 9, Human Override Events 6, Approval Rate 67%
2	Tabs	Decision Register · Ethics Charter · Responsible Use Checklist · Customer Impact · Consent & Purpose · Fairness Queue · Human Override Log
3	Decision Register	Per use case: data used, purpose, impact, consent (Captured/N/A), fairness, oversight (human-in/on-loop), decision (Approved/Conditional/Pending)

Purpose. A governance body that reviews data and AI use cases against ethics principles: consent, fairness, purpose limitation, and human oversight.

Why it's needed. Legal compliance is necessary but not sufficient; the board adjudicates *should we*, not just *may we* — and records the decision.

Step-by-step usage. (1) Open the Decision Register. (2) For each case check consent, fairness, oversight. (3) Use Fairness Queue for pending reviews. (4) Record approve/conditional/reject; log any human overrides.

DMBOK mapping. *Data Ethics* (a DMBOK theme woven through Governance), *Data Governance*, *AI & Big Data Governance*. **Principle (paraphrased):** ethical handling of data — consent, purpose limitation, fairness, transparency — is a core governance responsibility, not optional.

Student Insight. “Conditional” approvals show ethics as a gate with conditions, not a yes/no stamp. *Reflection: the Credit Risk case is “Conditional” with High impact — write one condition you would attach.*

Troubleshooting. Approval rate 67% reflects conditional/pending cases — work the Fairness Queue. A human override should always create a Human Override Log entry.

4.21 Document & Content Governance (/documents)

Nav group: AI, Ethics & Compliance · Audience: Privacy Lead · Legal · Stewards

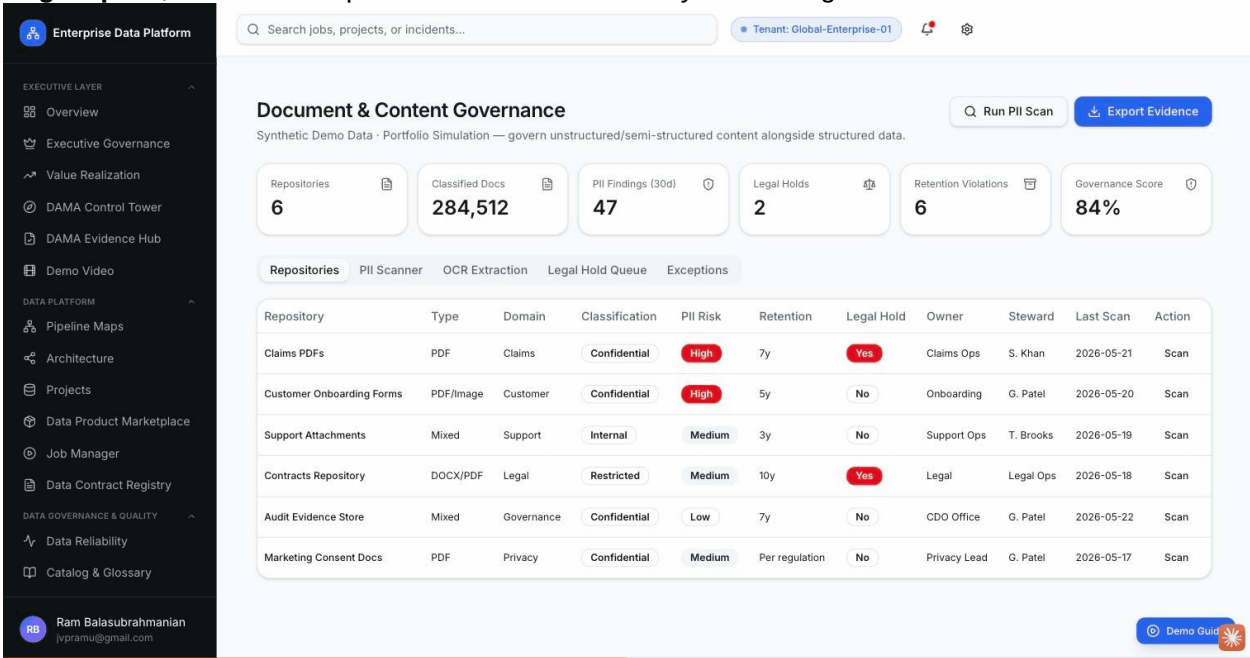


Figure 4.21 — Document & Content Governance: govern unstructured/semi-structured content alongside structured data.

Figure 4.21 — Document & Content Governance. Alt text: Content-governance KPIs and a repository table with type, domain, classification, PII risk, retention, legal-hold flag, owner, steward, last scan; tabs for Repositories, PII Scanner, OCR Extraction, Legal Hold Queue, Exceptions.

Annotations — What you’re seeing

Marker	UI element	What it does
1	Summary tiles	Repositories 6, Classified Docs 284,512, PII Findings (30d) 47, Legal Holds 2, Retention Violations 6, Governance Score 84%

2	Tabs	Repositories · PII Scanner · OCR Extraction · Legal Hold Queue · Exceptions
3	Repository table	Claims PDFs, Customer Onboarding Forms, Support Attachments, Contracts Repository, Audit Evidence Store, Marketing Consent Docs — with classification, PII risk, retention, legal hold
4	<i>Run PII Scan / Export Evidence</i>	Scan unstructured content and export findings

Purpose. Extend governance to unstructured/semi-structured content: classify documents, find PII, apply legal holds and retention.

Why it's needed. “Unstructured PII sprawl” is a named High enterprise risk. Most sensitive data hides in PDFs, forms, and attachments — this brings them under the same controls as tables.

Step-by-step usage. (1) Review repositories and classifications. (2) *Run PII Scan*; triage findings. (3) Apply *Legal Hold* where required. (4) Resolve *Exceptions* and retention violations; *Export Evidence*.

DMBOK mapping. *Document & Content Management* (primary), *Data Security* (PII), *Data Governance* (retention/legal hold). **Principle (paraphrased):** unstructured content is in scope for governance — it must be classified, secured, retained, and discoverable like structured data.

Student Insight. OCR + PII scanning turns opaque documents into governable, searchable assets. *Reflection: the Contracts Repository is Restricted with a 10-year retention and legal hold — why might legal hold override a scheduled disposal?*

Troubleshooting. Retention violations > 0 → records past policy not yet disposed; check Lifecycle. PII scan slow → large repos scan in the background.

4.22 Data Lifecycle & Retention Manager (/lifecycle)

Nav group: AI, Ethics & Compliance · **Audience:** Admin · Legal · Stewards




Figure 4.22 — Lifecycle & Retention: stage tracking, retention policies, archive eligibility, legal hold, and evidence export.

Figure 4.22 — Data Lifecycle & Retention Manager. Alt text: A horizontal lifecycle pipeline (Create → Ingest → Classify → Store → Use → Share → Retain → Archive → Dispose → Legal Hold) above a retention-policy table and a per-dataset retention grid.

Annotations — What you’re seeing

Marker	UI element	What it does
1	Lifecycle Stages	Create → Ingest → Classify → Store → Use → Share → Retain → Archive → Dispose → Legal Hold
2	Retention Policies	Customer PII 7y after closure; Billing Events 10y; Audit Logs 7y; Marketing Consent until withdrawal + audit; Claims Documents 10y; Temporary Staging 30d
3	Per-dataset grid	Domain, classification, policy, period, archive, deletion, legal hold, approval
4	Export Retention Evidence	Produce evidence of policy enforcement

Purpose. Track data through its full lifecycle and enforce retention, archival, disposal, and legal hold.

Why it’s needed. Keeping data too long is a liability; deleting too soon breaks compliance. Explicit, enforced retention policies manage both — and legal hold can pause disposal.

Step-by-step usage. (1) Read the lifecycle stages. (2) Confirm each dataset's retention policy and period. (3) Check archive/disposal eligibility. (4) Apply legal hold where litigation requires; *Export Retention Evidence*.

DMBOK mapping. *Data Storage & Operations* and *Document & Content Management* (retention), *Data Governance* (policy/legal hold). **Principle (paraphrased):** data has a managed lifecycle from creation to disposal; retention and disposition are governed, evidence-producing controls.

Student Insight. “Dispose” and “Legal Hold” are opposing forces resolved by governance.

Reflection: a dataset is due for disposal but under legal hold — which wins, and what evidence proves the right action was taken?

Troubleshooting. Dataset overdue for disposal → check for an active legal hold before deleting. Retention period blank → policy not yet assigned; assign before the asset is certified.

4.23 Stewardship Workbench (/stewardship)

Nav group: AI, Ethics & Compliance · **Audience:** Gita Patel & all Stewards




Figure 4.23 — Stewardship Workbench: operational queues for DQ, metadata, policy, MDM, SLA, RCA, and glossary.

Figure 4.23 — Stewardship Workbench. *Alt text: Stewardship KPIs and a work-queue table (STW-9001...) with queue, severity, dataset, domain, owner, steward, SLA, status, recommended action, evidence link, and row actions; queue tabs across the top.*

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Open Items 5, In Review 2, Escalated 1, Critical Severity 2
2	Queue tabs	All · DQ Exceptions · Metadata Gaps · Policy Exceptions · Access Requests · MDM Merge

		Approvals · SLA Breaches · RCA Review · Glossary Approval
3	Work items	STW-9001...STW-9005 with severity, SLA, recommended action, linked evidence (EV-22xx), approve/reject/escalate/assign

Purpose. A single operational queue where stewards triage and resolve governance work across every domain.

Why it's needed. Stewardship is where governance becomes daily action. Consolidating queues (DQ, metadata, policy, MDM, SLA, RCA, glossary) means nothing falls through the cracks, each with an SLA.

Step-by-step usage. (1) Start in *All* or filter by queue. (2) Open a Critical item and read the recommended action. (3) Approve / reject / escalate / assign. (4) Each action links to evidence for the audit trail.

DMBOK mapping. *Data Governance* (stewardship roles), *Data Quality*, *Metadata Management*, *Reference & Master Data* (the queues themselves). **Principle (paraphrased):** stewards are accountable for data in their domain; governance operates through assigned, SLA-bound stewardship work.

Student Insight. Every queue here corresponds to a control elsewhere in the app — stewardship is the human layer over automated controls. *Reflection: STW-9001 is a Critical DQ Exception with a 4-hour SLA — trace it back to the rule in the DQ Control Center that likely raised it.*

Troubleshooting. Escalated item stuck → check the owner's availability and the SLA clock. Missing evidence link → the upstream control didn't attach an artefact; raise a metadata gap.

4.24 Access Control — RBAC + ABAC (/access)

Nav group: Operations & Admin · **Audience:** Admin · T. Brooks (Auditor)

Enterprise Data Platform

Search jobs, projects, or incidents... Tenant: Global-Enterprise-01

Access Control — RBAC + ABAC

Synthetic Demo Data · Users, role-permission matrix, attribute-based policies, and a live policy decision simulator.

Active Users: 5 Distinct Roles: 5 ABAC Policies: 6 MFA Adoption: 80%

Users RBAC Matrix ABAC Policies Policy Simulator

Search name or email... All roles

User	Role	Department	Region	Clearance	Tenants	MFA	Onboarded	Last Login
Ram Balasubrahmanian jypramu@gmail.com · USR-0001	Platform Admin	CDO Office	Global	Restricted	Apex Financial, RetailCo Global, TechStart Inc, HealthSys Corp	On	Dec 19, 2025	Jun 09, 2026 09:42 UTC
Gita Patel gita.patel@apex-fin.com · USR-0002	Data Steward	Data Governance	APAC	Confidential	Apex Financial	On	Jan 04, 2026	Jun 08, 2026 16:18 UTC
Liam Mendes liam.mendes@retailco.com · USR-0003	Data Engineer	Platform Engineering	EMEA	Confidential	RetailCo Global	On	Jan 22, 2026	Jun 09, 2026 06:55 UTC
Aisha Rivera aisha.rivera@techstart.io · USR-0004	Business Analyst	Revenue Analytics	NA	Internal	TechStart Inc.	Off	Feb 11, 2026	Jun 07, 2026 21:03 UTC
Thomas Brooks t.brooks@healthsys.org · USR-0005	Auditor	Internal Audit	Global	Restricted	Apex Financial, HealthSys Corp	On	Feb 28, 2026	Jun 09, 2026 08:11 UTC

Demo Guide

Figure 4.24 — Access Control: users, role-permission matrix, attribute-based policies, and a live policy-decision simulator.

Figure 4.24 — Access Control (RBAC + ABAC). *Alt text: Access KPIs and a users table with role, department, region, clearance, tenants, MFA, onboarding, and last login; tabs for Users, RBAC Matrix, ABAC Policies, Policy Simulator.*

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Active Users 5, Distinct Roles 5, ABAC Policies 6, MFA Adoption 80%
2	Tabs	Users · RBAC Matrix · ABAC Policies · Policy Simulator
3	Users table	Ram (Platform Admin, Restricted), Gita Patel (Data Steward), Liam Mendes (Data Engineer), Aisha Rivera (Business Analyst, no MFA), Thomas Brooks (Auditor) — with clearance, tenants, MFA, last login

Purpose. Manage identity and authorisation using both role-based (RBAC) and attribute-based (ABAC) access control, with a simulator to test decisions.

Why it's needed. Least-privilege access is the foundation of data security. ABAC adds context (region, clearance, sensitivity); the simulator lets you *prove* what a given user can do before granting it.

Step-by-step usage. (1) Review users and clearances. (2) Inspect the RBAC Matrix for role → permission. (3) Define ABAC policies on attributes. (4) Use the *Policy Simulator* to test a subject/resource/action decision. (5) Flag missing MFA (Aisha Rivera).

DMBOK mapping. *Data Security* (primary), *Data Governance*. **Principle (paraphrased):** access should follow least privilege, combining roles with attribute conditions, and be testable and auditable.

DMBOK Principle Spotlight — RBAC + ABAC together. RBAC answers “what does this *role* allow?”; ABAC adds “...given these *attributes* (region, clearance, sensitivity)?” Combining them expresses real policies like “Auditors may read Restricted data only in their region.”

Student Insight. The Policy Simulator turns access policy into a testable artifact. *Reflection: Aisha Rivera lacks MFA — write an ABAC rule that blocks her from Restricted data until MFA is enrolled.*

Troubleshooting. Unexpected allow/deny → reproduce it in the Policy Simulator. MFA adoption < 100% → a flagged user (here, one) has not enrolled.

4.25 Billing & Usage (/billing)

Nav group: Operations & Admin · **Audience:** Admin · Finance




Figure 4.25 — Billing & Usage: revenue, contracted capacity, and overage forecasts per client.
Figure 4.25 — Billing & Usage. Alt text: Billing KPIs and a per-client usage panel showing contracted vs used hours, monthly rate, overage, and a 14-day actual-vs-predicted usage bar chart.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Total Monthly Revenue \$284,500, Clients in Overage 3, Predicted Overage Revenue +\$42,000, Avg Utilization 73%
2	Per-client card	e.g., Apex Financial In Overage : contracted 500 hrs/mo, used 612 (122%), \$48,000/mo, overage + \$10,752
3	Usage chart	14-day actual-vs-predicted usage
4	Actions	<i>Upgrade Contract, Send Alert</i>

Purpose. Track consumption, capacity, and overage revenue per tenant.

Why it's needed. Capacity and cost are governance concerns: overage signals under-provisioning or runaway jobs; predictive forecasts enable proactive contract management.

Step-by-step usage. (1) Read revenue and utilisation. (2) Identify clients in overage. (3) Review the usage forecast. (4) *Upgrade Contract* or *Send Alert* as appropriate.

DMBOK mapping. *Data Storage & Operations* (capacity/cost management), *Data Governance* (chargeback/accountability). **Principle (paraphrased):** operations includes capacity and cost management; usage should be measured, forecast, and attributed.

Student Insight. Overage is where operations meets finance — capacity planning has a dollar value. *Reflection: Apex is at 122% utilisation — would you upgrade the contract or investigate the workload first?*

Troubleshooting. Overage with stable workload → likely a contract sizing issue. Forecast spikes → correlate with new jobs in Job Manager.

4.26 Incidents (/incidents)

Nav group: Operations & Admin · **Audience:** Operator · On-call




Figure 4.26 — Incidents: AI-assisted root cause analysis and remediation.

Figure 4.26 — Incidents. Alt text: Incident KPIs and an incident table (INC-0037...INC-0041) with job name, error summary, severity (P1–P3), status, owner, and created time; opening an incident reveals AI root cause, impact assessment, remediation plan, and post-incident actions.

Annotations — What you're seeing

Marker	UI element	What it does
1	Summary tiles	Open Incidents 3, P1 Critical 1, Avg Resolution 42 min, Incidents This Week 5
2	Incident table	INC-0041 (NullPointerException in schema validator, P1, Open), INC-0040 (BigQuery timeout, P2), INC-0039 (Oracle unreachable, P2, Resolved), INC-0038 (memory overflow, P3), INC-0037 (schema mismatch, P3, Resolved)
3	AI RCA panels	Root Cause · Impact Assessment · Remediation Plan · Post-Incident Action Items

Purpose. Detect, triage, and resolve operational incidents with AI-assisted root cause analysis.

Why it's needed. Incidents are inevitable; structured RCA and post-incident actions turn each one into prevention. Linking incidents to lineage and reliability closes the loop.

Step-by-step usage. (1) Triage by severity (P1 first). (2) Open an incident for AI root cause and impact. (3) Execute the remediation plan. (4) Record post-incident actions; link to the affected asset's lineage.

DMBOK mapping. *Data Storage & Operations* (incident/problem management), *Data Quality* (many incidents are quality breaks). **Principle (paraphrased):** operations includes structured incident and problem management with root-cause analysis and preventive follow-through.

Student Insight. Note INC-0037's “schema mismatch” — it ties directly to the Contract Registry's job. *Reflection: which two incidents here would a working Schema Gate have prevented?*

Troubleshooting. Recurring same-cause incidents → the preventive action wasn't completed; check post-incident items. P1 with no owner → assign immediately and escalate.

4.27 Audit Trail (/audit)

Nav group: Operations & Admin · **Audience:** T. Brooks (Auditor) · Security

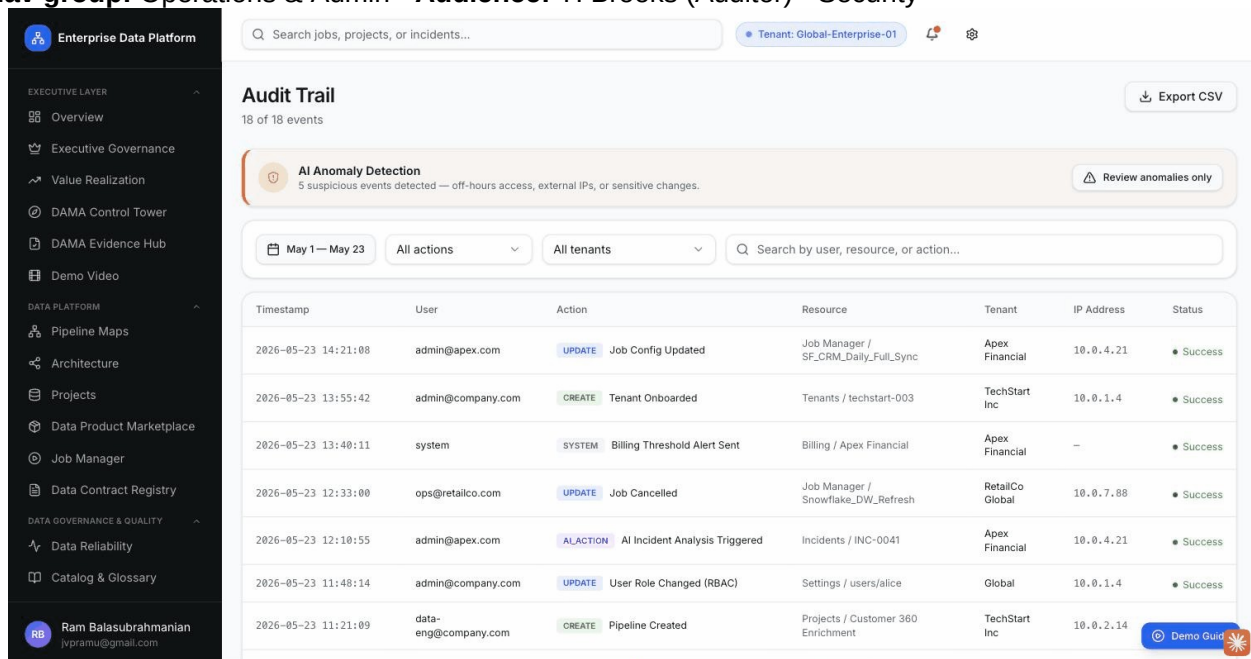


Figure 4.27 — Audit Trail: immutable activity log with AI anomaly detection.

Figure 4.27 — Audit Trail. Alt text: An “AI Anomaly Detection” banner over a filterable event log with timestamp, user, action, resource, tenant, IP address, and status.

Annotations — What you're seeing

Marker	UI element	What it does
1	Header	“18 of 18 events”; <i>Export CSV</i>
2	AI Anomaly Detection	“5 suspicious events detected — off-hours access, external IPs, or sensitive changes”; <i>Review anomalies only</i>
3	Filters	Date range, action type, tenant, free-text

4	Event log	Timestamp, user, action (UPDATE/CREATE/SYSTEM/AI_ACTION), resource, tenant, IP, status
---	-----------	--

Purpose. A complete, filterable record of who did what, where, and when — with AI flagging suspicious activity.

Why it’s needed. The audit trail is the backbone of accountability and forensic investigation. Anomaly detection surfaces risky events (off-hours, external IP) that humans would miss in a long log.

Step-by-step usage. (1) Filter by user/action/tenant/date. (2) Click *Review anomalies only* to see flagged events. (3) Investigate a sensitive change (e.g., “User Role Changed (RBAC)”). (4) *Export CSV* for the auditor’s pack.

DMBOK mapping. *Data Security* (audit logging), *Data Governance* (accountability). **Principle (paraphrased):** all access and change to sensitive data should be logged immutably and monitored for anomalies to support accountability and forensics.

Student Insight. Anomaly detection over the audit log is “governance watching governance.” *Reflection: which of the logged actions (RBAC change, off-hours access, external IP) would you alert on in real time, and to whom?*

Troubleshooting. Legitimate event flagged as anomaly → tune the detection (off-hours may be valid for some regions). Missing events → confirm logging is enabled for that resource.

4.28 Tenants (/tenants)

Nav group: Operations & Admin · Audience: Admin




Figure 4.28 — Tenants: enterprise customers with capacity and contract details.
Figure 4.28 — Tenants directory. Alt text: A card grid of four enterprise tenants, each with status, contact, start date, contracted capacity, and monthly rate, plus an Onboard New Tenant action.

Annotations — What you're seeing

Marker	UI element	What it does
1	Header	"4 enterprise customers onboarded"; <i>Onboard New Tenant</i>
2	Tenant cards	Apex Financial (500 hrs/mo, \$48K), RetailCo Global (300, \$28K), TechStart Inc (200, \$18K), HealthSys Corp (400, \$38K) — all Active
3	<i>Manage</i>	Per-tenant configuration

Purpose. Manage the multi-tenant estate: onboard customers and govern per-tenant capacity and isolation.

Why it's needed. Multi-tenancy demands data isolation, per-tenant access, and capacity governance. Onboarding is the controlled entry point (the "70% faster onboarding" value metric lives here).

Step-by-step usage. (1) Review active tenants and capacity. (2) *Manage* a tenant for its settings. (3) *Onboard New Tenant* using the governed flow. (4) Cross-check capacity against Billing.

DMBOK mapping. *Data Governance* (multi-tenant policy/isolation), *Data Security* (tenant boundaries), *Data Storage & Operations* (capacity). **Principle (paraphrased):** governance must scale across organisational boundaries with isolation, per-tenant policy, and controlled onboarding.

Student Insight. Tenant isolation is a security boundary expressed as configuration. *Reflection: what data-isolation guarantees must hold between RetailCo and HealthSys, and which screen proves them?*

Troubleshooting. Onboarding slow → it is a governed, multi-step flow by design. Capacity mismatch → reconcile Tenants vs Billing utilisation.

4.29 Settings (/settings)

Nav group: Operations & Admin · **Audience:** Admin

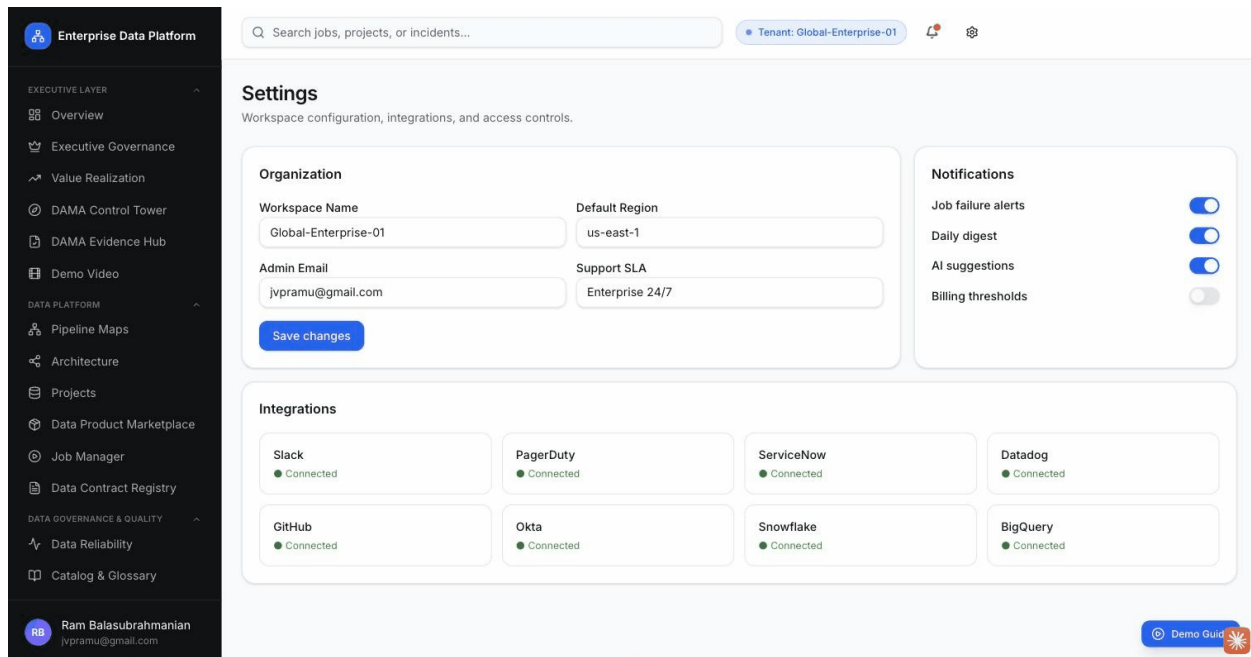


Figure 4.29 — Settings: workspace configuration, integrations, and access controls.

Figure 4.29 — Settings. *Alt text: An organisation settings panel, a notifications toggle panel, and an integrations grid showing connected tools (Slack, PagerDuty, ServiceNow, Datadog, GitHub, Okta, Snowflake, BigQuery).*

Annotations — What you're seeing

Marker	UI element	What it does
1	Organization	Workspace Name (Global-Enterprise-01), Default Region (us-east-1), Admin Email, Support SLA (Enterprise 24/7); Save changes
2	Notifications	Job-failure alerts, daily digest, AI suggestions, billing thresholds (toggles)
3	Integrations	Slack, PagerDuty, ServiceNow, Datadog, GitHub, Okta, Snowflake, BigQuery — all Connected

Purpose. Configure the workspace, notifications, and the external integrations that connect the platform to the wider tooling estate.

Why it's needed. Integrations make governance actionable beyond the app (alerts to Slack/PagerDuty, identity via Okta, warehouses via Snowflake/BigQuery). Notification policy controls signal vs noise.

Step-by-step usage. (1) Set organisation defaults and Save. (2) Tune notification toggles. (3) Review connected integrations. (4) Confirm identity (Okta) and warehouse (Snowflake/BigQuery) connections support your pipelines.

DMBOK mapping. *Data Integration & Interoperability* (integrations), *Data Governance* (workspace policy), *Data Security* (Okta/identity). **Principle (paraphrased):** a governed platform interoperates through managed, monitored integrations rather than ad-hoc connections.

Student Insight. The integration grid is the platform's "interoperability surface." *Reflection: which two integrations are most critical to the incident-to-resolution workflow, and why?*

Troubleshooting. Integration shows disconnected → re-authenticate the connector (do not store raw credentials in the app). Changes not saving → ensure *Save changes* was clicked in the Organization panel.

Section 5 — End-to-End Workflows

These walkthroughs chain several screens into a realistic task. Each is tied to the DMBOK data-lifecycle stages it exercises.

5.1 Onboarding a New Data Asset

DMBOK lifecycle: Plan → Acquire → Store → Maintain

1. **Model it** — In **Modeling Studio** (4.15), define the entity (conceptual → logical → physical) and run *Standards* validation.
2. **Catalogue it** — In **Catalog & Glossary** (4.14), register the asset, link a glossary term, assign owner + steward, and set sensitivity.
3. **Contract it** — In **Contract Registry** (4.12), publish the producer → consumer contract and enable Schema Gates.
4. **Quality-gate it** — In **DQ Control Center** (4.17), attach a metadata ruleset and run validation.
5. **Certify it** — Complete the **Catalog Approval Workflow**; the asset becomes Certified.
6. **Publish it** — Optionally list it in the **Data Product Marketplace** (4.10) with DQ + sensitivity shown. > Outcome: a findable, owned, contracted, quality-gated, certified asset — traceable end to end.

5.2 Running a Data Quality Check

DMBOK lifecycle: Maintain → Enhance

1. Open **DQ Control Center** (4.17); confirm the Incoming Feed and ruleset (DQ_RULESET_V3).
2. Review **Metadata Rules** (DQ-001...DQ-008); adjust with the *AI Rule Assistant* if needed.
3. Click **Run Data Quality Validation** and watch the 10-step pipeline.
4. Read **Validation Results** (e.g., 93 passed / 6 rejected / 1 quarantined).
5. Open **Reconciliation** to see the pipeline decision against the 10% threshold.
6. If the reject rate breaches threshold, the decision flips to *hold*; open **Analytics** to find the top failing rule, then route the exception to **Stewardship** (4.23). > Outcome: bad data rejected, risky data quarantined, an audit log produced — before anything reached the pipeline.

5.3 Tracing Lineage for an Audit

DMBOK lifecycle: Maintain → Dispose (evidence)

1. Start in **Evidence Hub** (4.5); open an *At Risk* control (e.g., a BI freshness SLA, EV-1014).
2. Jump to **Pipeline Maps** (4.7); locate the asset and click its node for AI impact analysis.
3. Identify upstream sources and downstream consumers of the failing link.
4. Cross-check **Reliability** (4.13) for the asset's SLO/error budget and root cause.
5. Capture the chain and *Generate Audit Pack / Export CSV* from the Evidence Hub. > Outcome: a defensible, end-to-end trace from control → lineage → reliability → evidence, in minutes.

5.4 Resolving an AI-Model Risk

DMBOK lifecycle: Plan → Apply → Govern

1. In **AI Governance** (4.19), open a Critical model (Credit Risk Scoring) and its Model Card.
2. Check **Bias & Fairness / Evaluation**; confirm PII handling (masked + tokenized).
3. Route the use case to the **Data Ethics Board** (4.20) Decision Register; capture consent, fairness, and oversight.
4. If approved *Conditional*, record the conditions and human-in-the-loop requirement.
5. File the model card as evidence in the **Evidence Hub** (4.5). > Outcome: a high-risk model governed, documented, ethically reviewed, and evidenced.

5.5 Onboarding a New Tenant

DMBOK lifecycle: Plan → Acquire → Operate

1. In **Tenants** (4.28), click *Onboard New Tenant* and complete the governed flow.
 2. In **Access** (4.24), assign roles + ABAC policies; require MFA; validate with the Policy Simulator.
 3. In **Settings** (4.29), confirm integrations (Okta identity, Snowflake/BigQuery) for the tenant.
 4. In **Billing** (4.25), set contracted capacity and overage alerts.
 5. Verify isolation and capacity on the **Overview** (4.1) and **Reliability** (4.13). > Outcome: a securely isolated, access-controlled, capacity-governed tenant — the “70% faster onboarding” value in practice.
-

Section 6 — DMBOK Deep-Dive Appendix (Teaching Layer)

6.1 Consolidated Mapping: Component → Knowledge Area → Principle → Practice

App Component	DMBOK Knowledge Area	Underlying Principle (paraphrased)	Real-world practice it teaches
DAMA Control Tower	All / Maturity Assessment	Govern from a central hub; assess maturity per area	DCAM/DAMA-DMM scoring & roadmap
Executive Governance	Data Governance	Governance demonstrates value & manages risk	Risk register + ROI for the data program
Value Realization	Governance / Quality / Metadata	Data is a measured asset	Benefits realisation, before/after KPIs
Evidence Hub	Data Governance	Controls require operating evidence	Audit-pack assembly, control testing
Pipeline Maps	Integration / Metadata / Architecture	Lineage is core metadata	Impact analysis, provenance, RCA
Architecture Viewer	Data Architecture	Blueprints align data to strategy	C4 diagrams, semantic layer design
Projects / Job Manager	Storage & Operations	Operations are owned & monitored	DataOps run management
Data Product Marketplace	Warehousing & BI / Governance	Data-as-a-product scales reuse	Product ownership, certification, access
Contract Registry	Integration & Interoperability	Versioned interface contracts + change control	Schema gates, breaking-change prevention
Reliability	Data Quality / Operations	Quality has measurable dimensions	SLO/SLI, error budgets
Catalog & Glossary	Metadata Management	Metadata makes data understandable	Catalogue + business glossary curation
Modeling Studio	Data Modeling & Design	Conceptual → logical → physical; standards	Model standards & change-impact
MDM & Reference	Reference & Master Data	Authoritative golden version	Match/merge/survivorship governance
DQ Control Center	Data Quality / Metadata	Prevent at the gate; rules-as-metadata	Metadata-driven validation & reconciliation
Data Governance	Data Security / Governance	Controls proportionate to sensitivity	Tokenization, masking, policy enforcement
AI Governance	AI & Big Data Governance	Govern models like assets	Model registry, risk tiers, model cards
Data Ethics Board	Data Ethics / Governance	Consent, fairness, purpose limitation	Responsible-use review & decisions

Document & Content	Document & Content Mgmt	Unstructured content is in scope	PII scan, OCR, legal hold, retention
Lifecycle & Retention	Storage / Governance	Managed lifecycle to disposal	Retention policy & legal hold
Stewardship	Governance / Quality	Stewards are accountable	SLA-bound stewardship queues
Access (RBAC/ABAC)	Data Security	Least privilege, testable	RBAC + ABAC + policy simulation
Billing & Usage	Operations / Governance	Measure & attribute usage	Capacity & chargeback
Incidents	Operations / Quality	Structured RCA & prevention	Incident/problem management
Audit Trail	Security / Governance	Immutable, monitored logging	Forensics & anomaly detection
Tenants	Governance / Security	Scale governance across boundaries	Multi-tenant isolation & onboarding
Settings	Integration / Security	Managed, monitored integrations	Connector governance, identity

6.2 The DAMA Wheel and the Environmental Factors Hexagon

The DAMA Wheel places **Data Governance at the centre**, surrounded by ten knowledge areas: Data Architecture; Data Modeling & Design; Data Storage & Operations; Data Security; Data Integration & Interoperability; Document & Content Management; Reference & Master Data; Data Warehousing & Business Intelligence; Metadata Management; and Data Quality. The platform mirrors this exactly: the **Executive Layer + DAMA Control Tower** form the governance hub, and the operational groups are the spokes. *Where this platform fits:* it is a working instantiation of the entire wheel in one application.

The Environmental Factors Hexagon explains *how* each knowledge area is delivered, through three “people” sides and three “activity/technology” sides:

Hexagon side	What it means	Where you see it in the platform
Goals & Principles	The “why”	Executive Governance, Ethics Charter
Activities	The “how” (functions)	DQ runs, validations, reviews, merges
Deliverables	The “what” (outputs)	Evidence artefacts, audit packs, model cards, catalogues
Roles & Responsibilities	The “who”	Personas: Admin, Steward, Engineer, Analyst, Auditor; RACI (7.6)
Practices & Techniques	Methods	Tokenization, survivorship, SLO/error-budgets
Tools	Technology	The platform itself + connected integrations

Teaching note. Ask students to take any one screen and label all six hexagon sides for it. The DQ Control Center, for example, has goals (trust), activities (validation), deliverables (reconciliation report), roles (DQ Lead/steward), practices (metadata rules), and tools (the rule engine).

6.3 Classroom Exercise Set

1. **Coverage audit.** Using the DMBOK Coverage Map (2.4) and the Control Tower (4.4), identify the two lowest-maturity areas and propose a one-quarter improvement plan with evidence targets.
2. **Quality at the gate.** In the DQ Control Center, change a rule's action from Reject to Quarantine and predict, then verify, the effect on the reconciliation decision.
3. **Lineage impact.** In Pipeline Maps, pick the failed sync link and list every downstream consumer at risk; propose a redundancy fix.
4. **Access policy.** Write an ABAC rule (subject/resource/action/attributes) that lets Auditors read Restricted data only in their own region, then test it in the Policy Simulator.
5. **MDM survivorship.** For the Customer golden record, design survivorship rules for email, phone, and address; justify "most-recent" vs "most-trusted" per attribute.
6. **Ethics review.** Take the Credit Risk Scoring model to the Ethics Board: capture consent, fairness, purpose limitation, and one approval condition.
7. **Retention conflict.** Find a dataset due for disposal but under legal hold; document which control wins and what evidence proves the correct action.
8. **Evidence chain.** Starting from one *At Risk* control in the Evidence Hub, assemble a complete audit trail across at least three screens and *Generate Audit Pack*.

6.4 DAMA-DMM Style Maturity Self-Assessment Rubric

Rate each area 1–5 (1 = Initial/ad-hoc, 2 = Repeatable, 3 = Defined, 4 = Managed/measured, 5 = Optimised). The "Platform reference" column shows the in-app score for calibration.

Knowledge Area	1 Initial	3 Defined	5 Optimised	Platform ref.	Your score
Data Governance	No owners	Policies defined, owners named	Continuous, value-measured	L4	—
Data Architecture	No blueprints	Standard diagrams exist	Architecture drives change	L4	—
Data Modeling & Design	Ad-hoc models	Standards enforced	Change-impact automated	L3	—
Data Storage & Operations	Manual ops	Monitored, scheduled	Self-healing, capacity-managed	L4	—
Data Security	Perimeter only	RBAC defined	RBAC+ABAC, simulated, audited	L5	—
Data Integration & Interop	Point-to-point	Contracts defined	Schema gates prevent breaks	L4	—
Document & Content Mgmt	Unmanaged files	Repos classified	PII-scanned, retained, held	L2	—
Reference & Master Data	Duplicates rife	Golden records piloted	Survivorship governed at scale	L3	—
Data Warehousing & BI	Spreadsheets	Certified marts	Data-as-a-product marketplace	L4	—
Metadata Management	Tribal knowledge	Catalogue + glossary	Auto-harvested, lineage-linked	L3	—

Data Quality	Reactive fixes	Rules defined	Metadata-driven, prevented at gate	L4	—
--------------	----------------	---------------	------------------------------------	----	---

Section 7 — Operations Runbook (Mechanics)

7.1 Routine Operational Tasks

Cadence	Task	Screen(s)
Continuous	Watch platform health, freshness SLAs	Overview (4.1)
Hourly	Triage incidents, P1 first	Incidents (4.26)
Daily	Clear steward queues within SLA	Stewardship (4.23)
Daily	Run/verify DQ validations on key feeds	DQ Control Center (4.17)
Weekly	Review reliability SLOs & error budgets	Reliability (4.13)
Weekly	Review schema drift / pending contracts	Contract Registry (4.12)
Monthly	Reconcile billing/overage; alert clients	Billing (4.25)
Quarterly	Refresh DAMA maturity + evidence pack	Control Tower (4.4), Evidence Hub (4.5)

7.2 Maintenance

Keep metadata rulesets current as new sources arrive (the Control Tower's stated DQ next action: "AI rule assistant rollout"). Re-harvest catalogue metadata after schema changes. Review ABAC policies whenever roles or tenants change. Rotate the tokenization salt per the security schedule and confirm "no raw values exported" remains enforced.

7.3 Backup, Restore & Persistence

State is held in the Postgres metadata store with a **Supabase mirror** for shared persistence; object data lives in **S3** (lakehouse tiers). Recommended posture: point-in-time recovery on Postgres, versioned S3 buckets, and periodic export of evidence/audit data via the in-app *Export CSV/JSON/Audit Pack* controls so governance artefacts survive independently of the platform. The user menu's *Reset demo session* restores the default admin/tenant posture in the training environment.

7.4 Monitoring

System Pulse (Overview) covers compute, queue depth, and I/O; Reliability covers data SLOs/error budgets; Audit Trail covers security anomalies; Settings routes alerts to Slack, PagerDuty, ServiceNow, and Datadog. Watch four signals daily: breached freshness SLAs, exhausted error budgets, open P1 incidents, and audit anomalies.

7.5 Known Errors & Resolutions

Symptom	Likely cause	Resolution
Source shows "Breached" freshness	Upstream lag (e.g., Oracle 61%)	Check Reliability root cause; open/link an Incident; consider

		redundancy
DQ reject rate > threshold	New/changed source vs old rules	Update Metadata Rules; re-run; route exception to Stewardship
Breaking schema change in prod	Schema Gate not enabled	Enable gates on the contract; run compatibility before promotion
Model stuck "Pending Review"	Awaiting governance/ethics sign-off	Complete Ethics Board decision; attach model card to Evidence
Retention violation	Disposal overdue / legal hold conflict	Resolve in Lifecycle; legal hold overrides disposal
Unexpected access allow/deny	RBAC/ABAC interaction	Reproduce in Policy Simulator; adjust policy
Integration disconnected	Expired connector auth	Re-authenticate in Settings (never store raw creds)

7.6 Escalation Path & RACI

Escalation: Operator/Steward → Domain Steward → DQ/Integration Lead → CDO Office → (Security: CISO; Ethics: Ethics Board). P1 incidents page on-call via PagerDuty immediately.
RACI for platform operations (R = Responsible, A = Accountable, C = Consulted, I = Informed)

Activity	Operator	Steward	Eng/DataOps	CISO/Security	CDO Office
Run DQ validation	R	C	C	I	A
Resolve steward queue item	I	R	C	I	A
Approve MDM merge	I	R	I	I	A
Enable schema gate / promote change	C	I	R	I	A
Tokenization / masking policy	I	C	C	R	A
RBAC/ABAC change	I	C	C	R	A
Incident P1 resolution	R	I	R	C	A
Quarterly maturity & evidence	C	C	C	C	A/R

Section 8 — Appendices

8.1 Glossary (DAMA-aligned)

Term	Definition
ABAC	Attribute-Based Access Control — authorisation from attributes (role, region, clearance, sensitivity).
Business Glossary	Agreed business definitions of terms, distinct from technical metadata.
Data Contract	A versioned agreement on the schema/meaning/SLA between a data producer and consumer.
Data Governance	The exercise of authority and control (planning, monitoring, enforcement) over data management.
Data Lineage	The record of data's origins and movement through systems and transformations.
Data Product	A reusable, owned dataset published with quality, SLA, and access controls.
Data Quality Dimension	A measurable facet of quality (completeness, accuracy, timeliness, consistency, validity, uniqueness).
Data Steward	A role accountable for the quality, definition, and governance of data in a domain.
DMBOK	DAMA's Data Management Body of Knowledge (2nd Ed.) — the reference framework this platform aligns to.
Error Budget	The allowable amount of unreliability before changes are frozen (SRE concept applied to data).
Golden Record	The authoritative, de-duplicated master version of an entity.
Knowledge Area	One of DAMBOK's data-management domains (e.g., Data Quality, Metadata Management).
Maturity Level (L1–L5)	A capability score from Initial (1) to Optimised (5).
Metadata	"Data about data" — technical, business, and operational descriptors.
MDM	Master Data Management — governing authoritative master/reference data.
Quarantine	Holding suspect records out of the pipeline pending steward review (vs outright reject).
RBAC	Role-Based Access Control — authorisation from assigned roles.
Reconciliation	The post-validation report comparing inputs, outputs, and the pass/reject/quarantine split against thresholds.
Retention Policy	Rules governing how long data is kept before

	archival or disposal.
SLO / SLI / SLA	Service-Level Objective / Indicator / Agreement — reliability targets, measures, and commitments.
Survivorship Rule	Logic deciding which source value “wins” when master records merge.
Tokenization	Replacing sensitive values with deterministic, non-reversible tokens that preserve join-ability.

8.2 Screenshot Shot List

All 29 figures were captured **live** from the running application and recorded as animated GIFs (static frames extracted for print). The DQ Control Center figure (4.17) is an 18-frame workflow tour.

Fig	Screen	Route	Asset (GIF / PNG)
2.1 / 4.8	Architecture	/architecture	08-architecture
3.1 / 4.1	Overview	/	01-overview
4.2	Executive Governance	/executive	02-executive
4.3	Value Realization	/value	03-value
4.4	DAMA Control Tower	/dama	04-dama
4.5	Evidence Hub	/evidence	05-evidence
4.6	Demo Video	/demo	06-demo
4.7	Pipeline Maps	/lineage	07-lineage
4.9	Projects	/projects	09-projects
4.10	Data Product Marketplace	/products	10-products
4.11	Job Manager	/jobs	11-jobs
4.12	Contract Registry	/contracts	12-contracts
4.13	Reliability	/reliability	13-reliability
4.14	Catalog & Glossary	/catalog	14-catalog
4.15	Modeling Studio	/modeling	15-modeling
4.16	MDM & Reference	/mdm	16-mdm
4.17	DQ Control Center (KEY SCREEN)	/dq	17-dq-control-center (18 frames)
4.18	Data Governance	/governance	18-governance
4.19	AI Governance	/ai-governance	19-ai-governance
4.20	Data Ethics Board	/ethics	20-ethics
4.21	Document & Content	/documents	21-documents
4.22	Lifecycle & Retention	/lifecycle	22-lifecycle
4.23	Stewardship	/stewardship	23-stewardship
4.24	Access (RBAC/ABAC)	/access	24-access
4.25	Billing & Usage	/billing	25-billing

4.26	Incidents	/incidents	26-incidents
4.27	Audit Trail	/audit	27-audit
4.28	Tenants	/tenants	28-tenants
4.29	Settings	/settings	29-settings

Optional deeper captures (not blocking): per-tab stills for multi-tab screens — Catalog (Glossary, Approval Workflow, Harvesting), MDM (Match/Merge, Reference Data, Steward Queue), Contracts (Schema Drift, Impact Analysis), AI Gov (Bias & Fairness), and the DQ *Rejected* and *Quarantine & Golden* tabs.

8.3 References

- DAMA International. *DAMA-DMBOK: Data Management Body of Knowledge*, 2nd Edition. Technics Publications, 2017.
- DAMA International. *The DAMA Dictionary of Data Management*, 2nd Edition.
- DAMA-DMM / DCAM maturity assessment frameworks (referenced for the maturity rubric in 6.4).
- Live source application: <https://pipeline-pulse-79.lovable.app/> (captured 14 June 2026).

Paraphrase note: DMBOK principles in this document are summarised in the author's own words for teaching; no substantial verbatim text is reproduced. Consult the DMBOK 2nd Edition directly for authoritative wording.

8.4 Change-Request Template

Field	Entry
CR ID	CR-_____
Date / Requestor	
Component / Screen	
DMBOK area impacted	
Description of change	
Reason / risk addressed	
Affected sections of this runbook	
Evidence to update (Evidence Hub)	
Approver (CDO Office)	
Status	Draft / Approved / Released

8.5 Quick-Reference Cheat Sheet

I need to...	Go to	One-line how
See platform health	Overview	Scan KPIs + Freshness Scoreboard
Check DMBOK maturity	DAMA Control Tower	Read the heatmap; export report
Validate a feed	DQ Control Center	<i>Run Data Quality Validation</i> → Reconciliation
Trace impact	Pipeline Maps	Click a node → AI impact analysis
Prove a control	Evidence Hub	Filter → <i>Generate Audit Pack</i>

Resolve duplicates	MDM & Reference	Match/Merge → survivorship → steward approve
Stop a breaking change	Contract Registry	Enable Schema Gate → Run Compatibility
Govern a model	AI Governance → Ethics Board	Model Card → Decision Register
Manage access	Access (RBAC/ABAC)	Define policy → Policy Simulator
Work my queue	Stewardship	Filter queue → approve/reject/escalate
Investigate security	Audit Trail	<i>Review anomalies only</i>
Handle an incident	Incidents	Open → AI RCA → remediation

Persona switching: sidebar avatar → *Switch persona / Switch tenant*. **Global search:** top bar.
Guided tour: *Demo Guide* (bottom-right).

End of runbook. Classification: Internal – Training Use. © 2026 — prepared for Ramachandran Balasubrahmanian. All platform data is synthetic.