TraxrGRC

User Manual

October 2026 · build 244823a

A self-hosted governance, risk and compliance platform.

About this manual

This manual describes TraxrGRC from the user's point of view: what each screen is for, what its fields mean, and how to complete each workflow. It reflects the product as shipped.

Who this manual is for

- **GRC practitioners**: risk managers, compliance and audit staff, policy owners, privacy officers and incident responders who record and maintain the program day to day.
- **Record owners and contributors**: control owners, risk owners and other colleagues who are assigned reviews, tests, findings or questionnaires.
- **Managers and auditors** who read dashboards and reports and receive scheduled reports by email.
- **Administrators** who configure users, roles, integrations, notifications and backups.
- **Staff** who read and acknowledge policies in the Policy Portal.

New users who want the fast path should start with the Quick Start Guide. Installation, operations, single sign-on and administration in depth are in the Administrator's Guide.

Developers integrating with TraxrGRC should read the REST API chapter together with the machine-readable OpenAPI document. Engineering conventions for people working on the product itself are in `docs/CONTRACTS.md`.

How this manual is organized

The chapters follow the sidebar of the application:

Chapter	What it covers
Getting Started	Signing in, the layout, the Dashboard and the conventions shared by every module
Administration & Settings	Users, groups, roles, labels, custom fields, integrations, notifications, the audit log, API tokens and backups
Risk Management	The risk register, treatments, reviews, exceptions, controls, assets, business units and risk settings
Compliance Management	Compliance packages, analyses, the gap view, the POA&M register, exceptions and the compliance report
Policy Management	Policies, versions, approval, reviews, the Policy Portal and acknowledgements
Third-Party Management	Vendors and partners, contacts, linked risks and data protection details
Incident Management	Incidents, classifications and personal data breaches
Internal Audits	Audit plans and findings
Projects	Remediation and improvement projects
Security Awareness	Training programs, tests and results
Data Privacy	Records of processing, transfers, the Article 30 register, DPIAs, reviews and data mapping questionnaires
Account Reviews	Periodic reviews of user accounts
Evidence	Internal evidence and third parties' assurance reports, linked to what they support
Online Assessments	Questionnaire templates and requests answered by outside contacts
Report Center	Cross-module reports, CSV downloads and scheduled email delivery
REST API	Programmatic access with API tokens

Each module chapter opens with what the module is for and who uses it, then describes the list page, the record detail page and its tabs, and the procedures for the module's workflows. Read Getting Started first; the other chapters can be read in any order.

Conventions

Convention	Meaning
Bold	A label you see on screen: a menu entry, page title, tab, button, field or status, written exactly as it appears.
Settings → API Tokens	A navigation path: open **Settings** in the sidebar, then **API Tokens**.
`monospace`	Text you type or read literally: addresses, file names, permission names, codes and API values.
`module:right`	A permission, for example `risks:read` (Read on the Risks module) or `reports:create`. Roles grant rights per module; see Administration & Settings.
own and *all*	The two levels at which a role can hold a right. At *own* you work only with records that you or one of your groups own; at *all* you work with every record in the module.
Numbered steps	A procedure to follow in order.

Permission notes. What you see depends on your role. A sidebar entry, button, tab or report appears only when your role grants the right it needs, so your screens may show fewer options than the screenshots here. Where a feature depends on a specific right, the text names the module and the right. Administrators have every right in every module.

Screenshots. The figures show a demonstration installation populated with fictional data for an invented company, Quillmont Systems. Names, dates and counts on your installation will differ, and the product name and logo follow your organization's branding.

Dates and times are shown in your own time zone, which you set on your profile (**My profile → Time zone**).

Contents

1 Getting Started 10

1.1 What you can manage 10
1.2 Signing in 12
1.3 The app layout 14
1.4 The Dashboard 18
1.5 My Tasks 24
1.6 My profile 27
1.7 The Policy Portal 29
1.8 Shared UI conventions 30

2 Administration & Settings 39

2.1 Finding the Settings screens 39
2.2 Users 42
2.3 Groups 45
2.4 Roles 45
2.5 Labels 48
2.6 Custom Fields 48
2.7 Integrations 50
2.8 Provisioning (SCIM) 56
2.9 Sessions 58
2.10 Branding 59
2.11 Policies 61
2.12 Privacy 61
2.13 AI assistants 62
2.14 Notifications 65
2.15 Audit Log 70
2.16 API Tokens 74
2.17 Backups 75

3 Risk Management 78

3.1 The risk register and its tabs 78
3.2 The Risks list 79
3.3 Creating a risk 81
3.4 How scoring works 83
3.5 The risk detail page 84
3.6 Treatments 86
3.7 Reviews 89
3.8 Exceptions 91

3.9 Comments, attachments and history 93
3.10 Controls 93
3.11 Assets 98
3.12 Business Units 99
3.13 Risk Settings 99
3.14 Who can do what (Risk Management) 102

4 Compliance Management 103
4.1 Packages 103
4.2 Importing a package 107
4.3 Exporting a package 108
4.4 Common controls and inheritance 108
4.5 The NIST SP 800-53 catalog 109
4.6 Crosswalks between frameworks 110
4.7 The analysis workspace 111
4.8 Analysis statuses 114
4.9 Working on an item 114
4.10 The gap view (POA&M) 119
4.11 The POA&M register 121
4.12 Exceptions on a clause 124
4.13 The compliance report (PDF) 125
4.14 Deleting packages and items 126
4.15 Who can do what (Compliance Management) 127

5 Policy Management 129
5.1 The policies list 129
5.2 Creating a policy 130
5.3 The policy detail page 131
5.4 Writing content (Content tab) 132
5.5 The draft → approval → publish workflow 134
5.6 Versions tab 139
5.7 Reviews tab 139
5.8 Exceptions 141
5.9 Archive and reactivate 143
5.10 Comments, Attachments, History and Custom Fields 143
5.11 Policy portal and acknowledgements 144
5.12 Deleting a policy 150
5.13 Who can do what 151

6 Third-Party Management 153
6.1 The register 153

6.2 Creating a third party 154
6.3 The detail page 156
6.4 Data protection 158
6.5 Deleting a third party 162
6.6 Who can do what 162

7 Incident Management 164
7.1 Who can do what 164
7.2 Managing the classification catalog 165
7.3 The incidents list 165
7.4 Creating an incident 167
7.5 The incident detail page 168
7.6 Closing and reopening an incident 170
7.7 Personal data breaches 171
7.8 Editing and deleting 178
7.9 Comments, attachments and history 178

8 Internal Audits 180
8.1 Who can do what 180
8.2 The audit plans list 180
8.3 Creating an audit plan 182
8.4 The plan detail page 183
8.5 Working with findings 185
8.6 Open findings 188
8.7 Comments, attachments and history 190

9 Projects 191
9.1 Who can do what 191
9.2 The projects list 191
9.3 Creating a project 193
9.4 The project detail page 194

10 Security Awareness 200
10.1 The programs list 200
10.2 Creating a program 201
10.3 Program status: draft, active, archived 202
10.4 The program detail page 203
10.5 Content tab: the training material 205
10.6 Test tab: building the question set 206
10.7 Tracking tab: who has completed the program 207
10.8 Taking a test 209

10.9 Comments, attachments, history and custom fields 211
10.10 Deleting a program 212
10.11 Awareness permissions 212

11 Data Privacy 213
11.1 The records list 213
11.2 Creating a record 215
11.3 The record detail page 216
11.4 Documenting the record 218
11.5 Article 30 completeness 220
11.6 Transfers 221
11.7 Linked assets, risks and controls 223
11.8 The Article 30 register 224
11.9 DPIAs 224
11.10 The Reviews tab 229
11.11 Review notifications 230
11.12 Record status: draft, active, archived 231
11.13 Comments, attachments, history and custom fields 231
11.14 Deleting a record 232
11.15 Data mapping questionnaires 232
11.16 The data map 241
11.17 Data privacy permissions 243

12 Account Reviews 245
12.1 The programs list 245
12.2 Creating a program 246
12.3 The program page 247
12.4 Adding items to a program 249
12.5 Reviewing items: keep or disable 250
12.6 Comments, attachments and history 251
12.7 Archiving, reactivating and deleting 253
12.8 Who can do what 254

13 Evidence 255
13.1 The Evidence page 255
13.2 Adding evidence 257
13.3 An evidence record 258
13.4 Evidence on other records 259
13.5 Replacing earlier evidence 261
13.6 Locked evidence and file fingerprints 261
13.7 Expiry and reminders 262

13.8 Evidence packages for auditors 263
13.9 Who can do what with evidence 264

14 Online Assessments 265
14.1 The Assessments page 265
14.2 Templates and the question editor 266
14.3 Requests 269
14.4 Reviewing answers 275
14.5 Repeating assessments 280
14.6 The public assessment page (for external contacts) 280
14.7 Notifications and audit trail 284
14.8 Who can do what 285

15 Report Center 286
15.1 The landing page 286
15.2 Anatomy of a report 287
15.3 Risk Register 288
15.4 Compliance Status 289
15.5 Exceptions 290
15.6 Reviews Due 292
15.7 Awareness Completion 293
15.8 Incidents 293
15.9 Policy Acknowledgements 294
15.10 Scheduled reports 295
15.11 Who can see reports 302

16 REST API 303
16.1 Authentication 303
16.2 Conventions 304
16.3 What is available 304
16.4 Related access paths 306

Index

CHAPTER 1

Getting Started

TraxrGRC is a self-hosted governance, risk and compliance platform. It gives your organization one place to manage risks and controls, track compliance against frameworks, maintain policies, oversee third parties, respond to incidents, run audits, deliver awareness training, and coordinate projects, data privacy records, account reviews and online assessments.

This chapter is for everyone who uses TraxrGRC. It explains how to sign in, how the application is laid out, how to read the Dashboard, where to find your own profile and the Policy Portal, and the on-screen conventions you will meet in every module. Administrators should continue with Administration & Settings afterward.

1.1 What you can manage

Each major area of GRC work has its own module, reached from the sidebar:

Module	What it is for
Risk Management	Risk register, treatments, reviews and exceptions, plus the supporting controls, assets, business units and risk settings
Third Parties	Vendors and partners, their contacts, contracts and review lifecycle
Incidents	Security and operational incidents with response tracking, including personal data breaches
Compliance	Compliance analyses against requirement packages (for example NIST SP 800-53 or the GDPR), the packages themselves, and the gap view
Policies	Policy documents with owners, versions, approvals and periodic reviews
Audits	Audit plans and their findings
Assessments	Online questionnaires sent to internal or external respondents
Account Reviews	Periodic reviews of user accounts and their items
Awareness	Training programs and who has completed them
Projects	GRC-related projects (remediation, rollout) tracked by status
Data Privacy	Records of processing activities, DPIAs, transfers and periodic reviews
Reports	The Report Center, and report emails sent on a schedule
My Tasks	Everything waiting on you across the modules you can read: decisions, acknowledgements, training, questionnaires, reviews, control tests, findings, POA&M items and projects
Policy Portal	The published policies everyone can read, and the policies you are asked to acknowledge
My Questionnaires	For Data Liaisons: the data mapping questionnaires assigned to you

User accounts, groups, roles, labels, custom fields, sign-in integrations, notifications, the audit log, API tokens and backups are administered from Settings (see Administration & Settings).

1.2 Signing in

1. Open TraxrGRC in your browser. If you are not signed in you are taken to the `/login` page, which shows the **Sign in to TraxrGRC** form.
2. Enter your **Email** and **Password**.
3. Click **Sign in**. When your credentials are accepted you land on the page you originally asked for, or on the Dashboard if you opened the sign-in page directly.

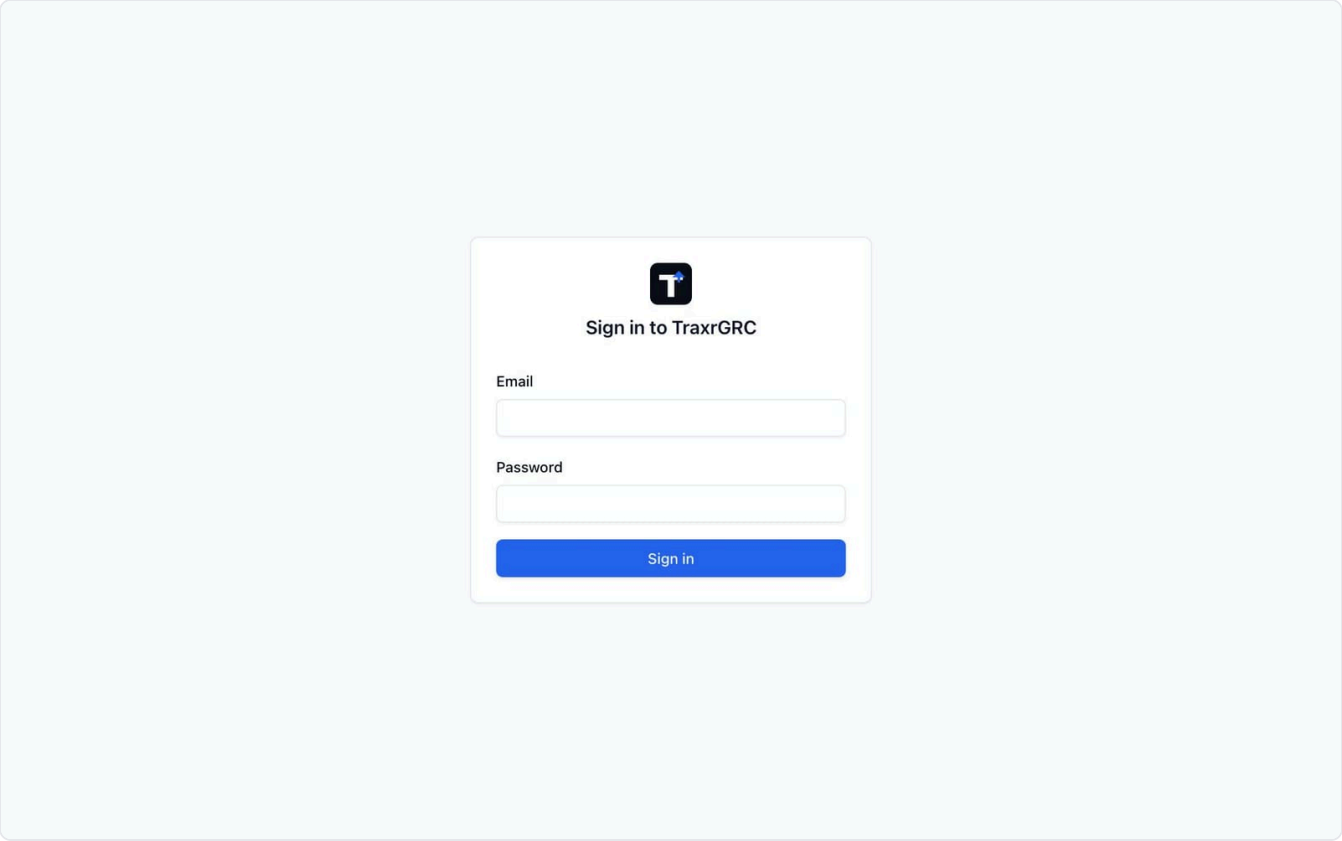


Figure 1.1. The sign-in page with the Email and Password fields and the Sign in button.

Where you land depends on your rights: someone without Dashboard access is taken to the **Policy Portal**, and a Data Liaison with no Dashboard access to **My Questionnaires**.

When an administrator has turned on **SSO only** (see SAML single sign-on), the page shows just **Sign in with SSO**, and you sign in through your organization's identity provider. The **Emergency administrator sign-in** link below it opens the password form, which works only for break-glass administrators.

If sign-in fails, the form tells you why:

Message	Meaning
Invalid email or password.	The email/password combination was not recognized.
Your account is inactive. Contact an administrator.	Your account has been deactivated; an administrator must re-enable it.
Sign-in is misconfigured. Contact an administrator.	The sign-in setup on the server is broken; report it to your administrator.
Single sign-on failed. Please try again or contact an administrator.	The identity provider's answer was refused (expired, for another service, not a reply to your sign-in, or the account cannot be linked). The administrator can see the reason in the audit log.
Sign-in failed. Please try again.	Any other failure, such as a lost connection.

Repeated wrong passwords lock **your address** out of **that account** for 15 minutes after five failures, and an address out of every account after 30. The account itself is never locked: from another network you (or a colleague) can still sign in with the right password. After ten failures on one account in 15 minutes, every address that has failed against it is refused for the rest of the window and administrators are alerted.

Two directory options can be enabled by an administrator (see Administration & Settings):

- **Single sign-on (SAML).** When SAML is enabled, a **Sign in with SSO** button appears below the password form. Click it to authenticate with your organization's identity provider. You return to a page reading "Completing single sign-on..." and then to the page you asked for. If the identity provider's answer did not arrive within 10 minutes, the page says "The SSO sign-in window expired or no assertion was received. Please try again." and offers **Back to sign in**.
- **LDAP.** When LDAP is enabled, directory authentication happens invisibly: if your email and password do not match a local account, TraxrGRC automatically tries your directory account before reporting a failure. There is no separate LDAP button. The part of your email before "@" is the directory username, and your directory entry must have an email address: without one, sign-in is refused.

Two-step verification. If you have set up two-step verification (see Sign-in security under My profile), a correct password or directory password takes you to the **Two-step verification** page, which shows the email you are signing in as:

1. Enter the **Code from your authenticator app** and click **Verify**. Each code works once; if one is refused as already used, wait for the next.
2. Without your device, click **Use a recovery code instead**, enter one of your recovery codes and click **Verify**. That recovery code is then spent.

Cancel and sign out returns you to the sign-in page. The page waits 15 minutes for a code; after that, sign in again. After 10 wrong codes in 15 minutes (app codes and recovery codes together) the code step for your account is locked from every address: the page says "Too many wrong codes. Wait 15 minutes and try again.", signing in again with your password does not lift it, and administrators are alerted. A correct code resets the count. When your administrator requires two-step verification and you have not set it up, the same page asks you to set it up (**Set up now**) before you can continue, with the steps described under Sign-in security. Single sign-on never asks for a code; your identity provider handles its own verification.

An LDAP or SSO sign-in never takes over an existing account that has its own password, administrator rights or break-glass status, even with the same email. Such an account signs in through LDAP or SSO only once it is linked: an administrator allows the link for that user in Settings → Users, and the next directory or SSO sign-in makes it. Accounts created by a first LDAP or SSO sign-in link automatically.

1.2.1 Signing out and session length

To sign out, open the user menu in the top-right corner of any page (your initials and name) and click **Sign out**. You return to the sign-in page. **Sign out everywhere** in the same menu also ends your sessions on every other browser and device, within a minute; use it after signing in on a shared computer or if you think someone else has your session.

A session ends after the session length (12 hours unless your administrator changed it) or after a period without use (60 minutes unless changed), whichever comes first; you then sign in again. An administrator resetting your password or your two-step verification, deactivating your account, or changing your administrator rights or groups also ends your sessions within a minute. Changing your own password ends every session except the one you changed it from.

1.3 The app layout

Every screen after sign-in shares the same frame.

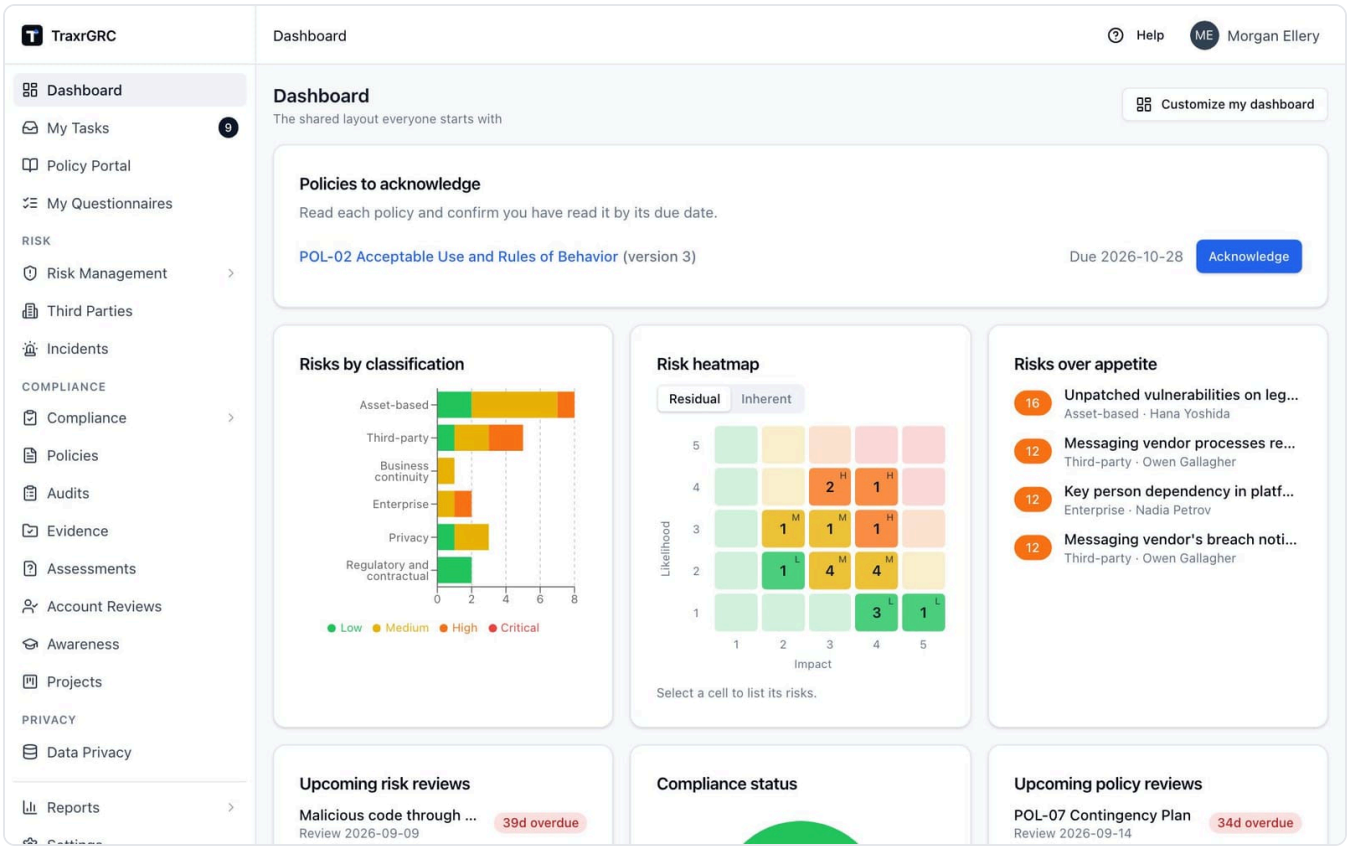


Figure 1.2. The application frame: the sidebar on the left with its Risk, Compliance and Privacy sections, the header with the user menu, and the Dashboard.

1.3.1 Sidebar

The sidebar on the left is the main navigation. Entries are grouped into sections:

Section	Entries
(top)	**Dashboard**; **My Tasks** (every signed-in user); **Policy Portal** (every signed-in user); **My Questionnaires** (Data Liaisons)
Risk	**Risk Management** (with **Risks**, **Controls**, **Assets**, **Business Units** and **Risk Settings**); **Third Parties**; **Incidents**
Compliance	**Compliance** (with **Analyses**, **Packages**, **Gaps**, **POA&M** and **Catalog**); **Policies**; **Audits**; **Assessments**; **Account Reviews**; **Awareness**; **Projects**
Privacy	**Data Privacy**
(bottom)	**Reports** (with **Report Center** and **Scheduled**); **Settings**

An entry with sub-pages opens automatically while you are on one of its pages; click its name to open or close it yourself.

The **Collapse** button at the bottom of the sidebar hides the labels and section headings and leaves only icons; hover an icon to see its name, and click the button again (now showing a right chevron) to expand.

You only see entries for modules you have access to. Access is controlled by permissions: each module grants rights such as read, create, update and delete, at the level *all* (every record), *own* (records owned by you or your groups) or *none*. Roles bundle these rights and are assigned to groups; you inherit the roles of every group you belong to. Administrators bypass all checks. If a module is missing from your sidebar, you do not have its read permission; ask your administrator.

1.3.2 Header, help and user menu

The header across the top shows where you are on the left, as a trail such as **Risk Management > Controls > AC-02 Account management**: each earlier step links back to its list. **Help** and your user menu are on the right.

Help opens a panel beside the page with the part of this manual that explains the page you are on, so the page and anything you have typed stay as they are. The panel ends with **Open in the full User Manual**, which opens the whole manual at that section in a new tab, with a contents list and links to the Quick Start Guide and, for administrators and anyone with `settings:read`, the Administrator's Guide. The help is this manual with your product name, so it matches the version you are using.

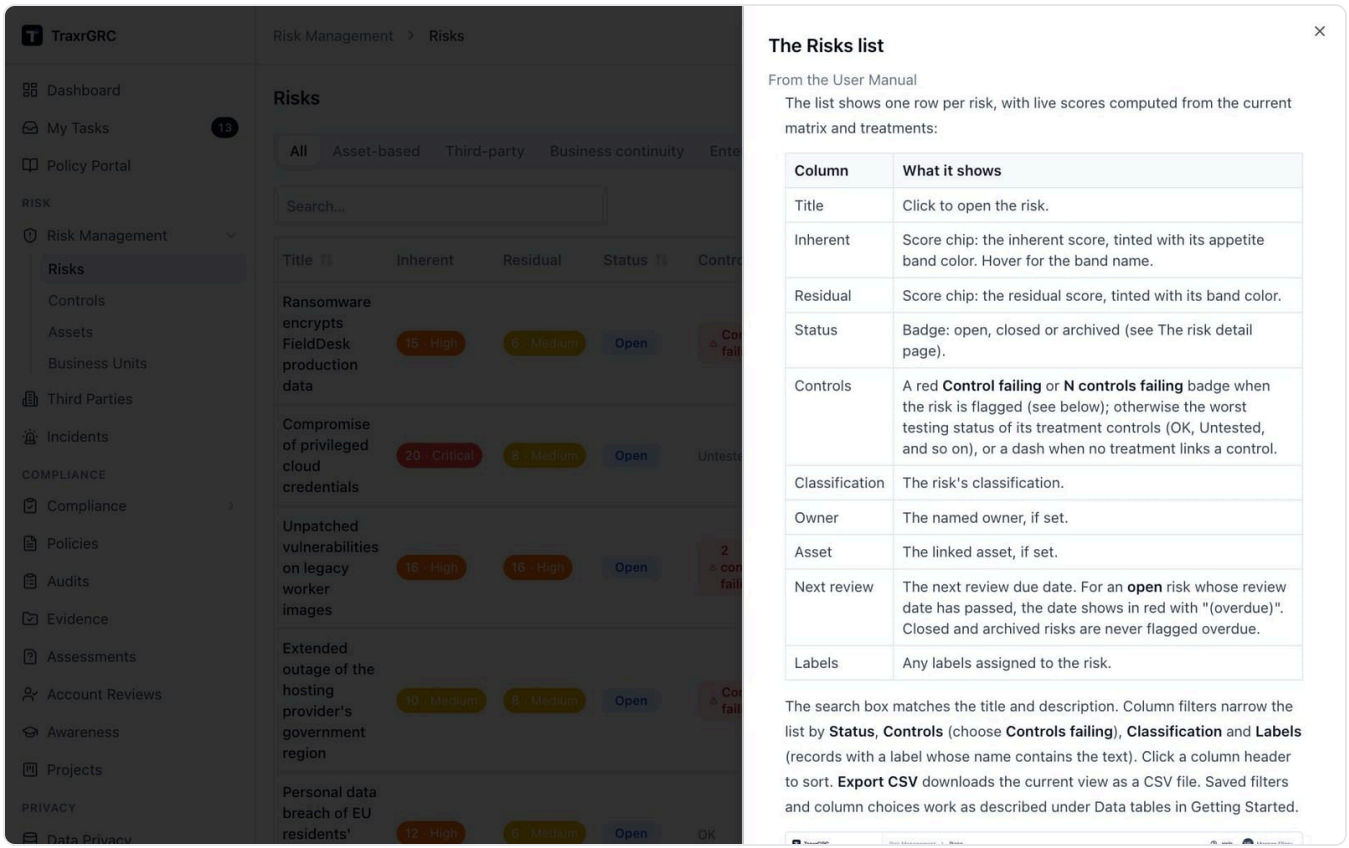


Figure 1.3. The Help panel open beside the Risks list, showing that page's section of this manual and Open in the full User Manual at the bottom.

Click your initials and name to open the user menu. It shows your full name and email and contains:

Item	What it does
My profile	Opens your profile: time zone, your policies and your notification choices (see My profile below).
Sign out	Ends this session and returns you to the sign-in page.
Sign out everywhere	Ends this session and every other session you have, on any browser or device, within a minute.

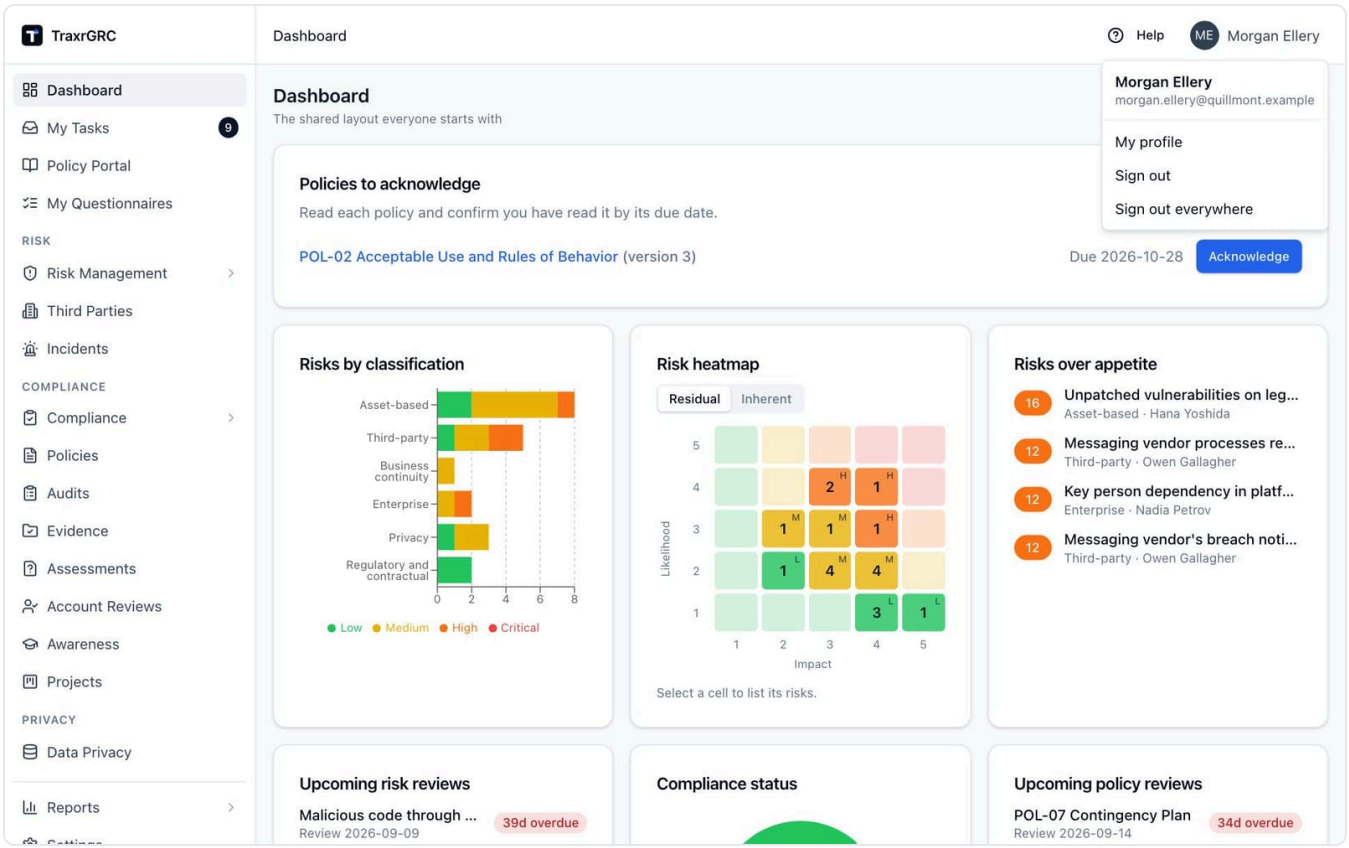


Figure 1.4. The user menu open in the header, showing My profile, Sign out and Sign out everywhere.

1.4 The Dashboard

The Dashboard is the first screen after sign-in and gives a one-glance health check across the program. When you have policies to acknowledge, a **Policies to acknowledge** card sits above the widgets (see Policy portal and acknowledgements). Someone with no Dashboard read right is taken to the **Policy Portal** instead (a Data Liaison to My Questionnaires).

Everyone starts on the shared dashboard, the same set of widgets for every user, and can make their own (see Your own dashboard). Each widget only counts records you are allowed to see (widgets honor *own*-level permissions), and a widget is hidden entirely if you lack read access to its module.

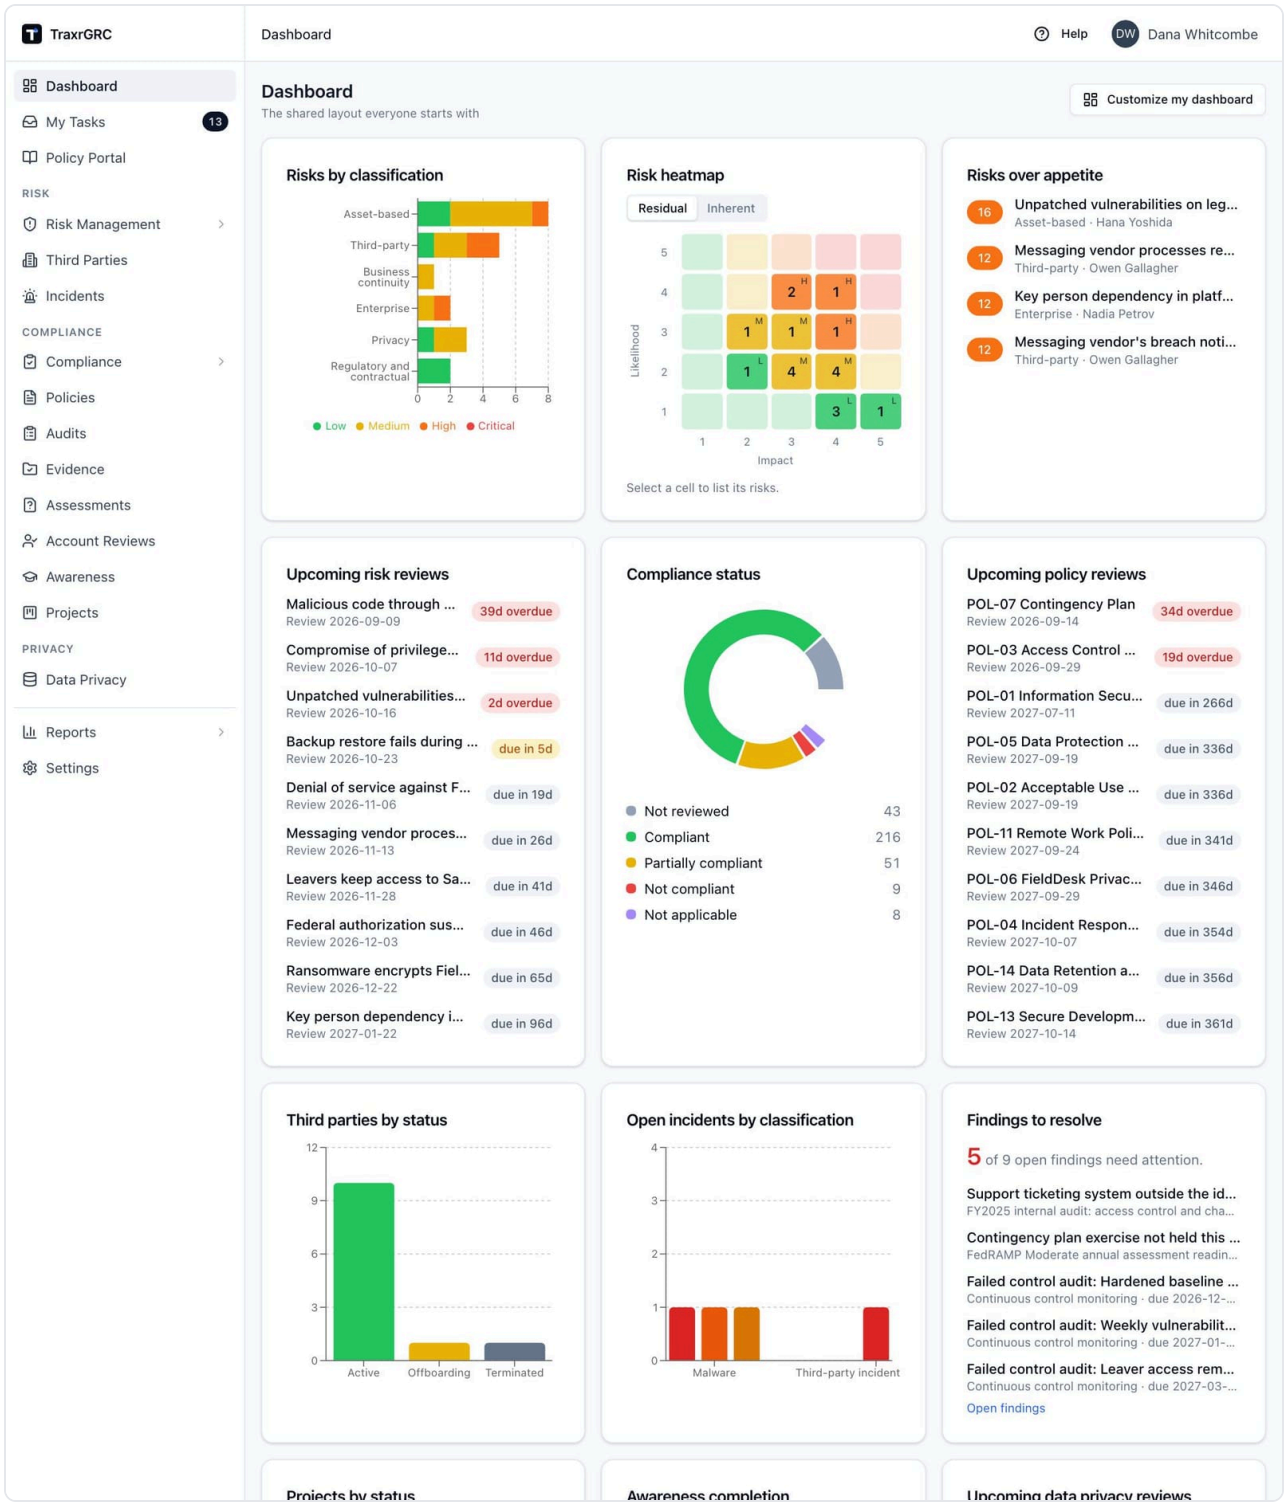


Figure 1.5. The Dashboard as a GRC manager sees it: the top of the shared widget grid.

The default dashboard contains these widgets in this order. When an update ships a new default layout, it replaces the shared dashboard on the next deploy, unless an administrator has changed the shared dashboard (added, removed, moved or retitled a widget); a changed dashboard is never overwritten, and neither is anyone's own:

Widget	What it shows	Needs
Risks by classification	Horizontal bar per risk classification counting your visible, non-archived risks, each bar split by the appetite band of the risk's residual score (lowest band first, in the band colors; "Not scored" in gray when a score falls in no band). Empty message: "No risks yet."	risks:read
Risk heatmap	The likelihood × impact matrix with the number of your visible **open** risks in each cell, colored by the cell's appetite band, with the band's first letter (for example L, M, H, C) in the corner of each occupied cell so the band does not rely on color alone. Likelihood runs up the side (highest on top), impact along the bottom; hover a number for the level's name. **Residual** (the default) and **Inherent** switch which scores place the risks. Click a cell to list its risks below the matrix, each linked to the risk; click it again to clear. Empty: "No open risks."	risks:read
Risks over appetite	Up to ten **open** risks whose residual score sits above your organization's risk appetite, highest score first. Each row shows a colored score chip, the risk title (click to open it) and, below it, its classification and owner. Empty: "No open risks currently exceed appetite."	risks:read
Upcoming risk reviews	The ten nearest incomplete risk reviews on risks that are not closed or archived, soonest first, with the scheduled date below the title and a due chip (see below). Click a title to open the risk. Empty: "No upcoming risk reviews."	risks:read
Compliance status	Donut of the compliance analyses you can read, by effective status (an inherited requirement counts with its provider's status): Not reviewed (gray), Compliant (green), Partially compliant (yellow), Not compliant (red), Not applicable (purple). Empty: "No compliance analyses yet."	compliance:read
Upcoming policy reviews	The ten nearest incomplete policy reviews for non-archived policies, soonest first. Empty: "No upcoming policy reviews."	policies:read
Third parties by status	Counts by status: Active (green), Offboarding (yellow), Terminated (gray). Empty: "No third parties yet."	thirdParties:read

Widget	What it shows	Needs
Open incidents by classification	Bar chart of incidents currently **open**, grouped by incident classification. Empty: "No open incidents."	incidents:read
Audit findings by status	Counts of findings by status: Open (red), In remediation (orange), Closed (gray). Empty: "No audit findings yet."	audits:read
Findings to resolve	How many open findings need attention out of all open findings you can read, then up to ten of them, soonest deadline first, each with its plan, deadline and reason (see Open findings under Internal Audits), and a link to the filtered Open findings view. It replaces Audit findings by status in the default dashboard for new installations. Empty: "No open audit findings." or "Every open finding is on track."	audits:read
Projects by status	Counts by status: Planned (gray), In progress (orange), Completed (green). Empty: "No projects yet."	projects:read
Awareness completion	Stacked bar per **active** awareness program (earliest due date first) showing the share of active users who passed versus everyone else. Empty: "No active awareness programs."	awareness:read
Upcoming data privacy reviews	The ten nearest incomplete reviews on data privacy records that are not archived, soonest first. Empty: "No upcoming data privacy reviews."	dataPrivacy:read
Account review items	Counts of items in **active** account review programs by status: Pending (orange), Kept (green), Disabled (gray). Empty: "No account review items in active programs."	accountReviews:read
Assessment requests	Counts of online assessment requests by status: Sent (orange), Completed (green), Expired (gray). A request counts as expired once its due date has passed. Empty: "No assessment requests yet."	assessments:read
Control testing status	Counts of your visible controls by calculated testing status: Failed (red), Overdue (amber), OK (green), Untested (gray). Empty: "No controls yet."	controls:read
Compliance at risk	How many requirements are **At risk** (see Compliance Management), then up to ten of them, each a link to the requirement with its clause code, package and the reason below: a failed, overdue or retired control, or	compliance:read

Widget	What it shows	Needs
	policy issues. Counts only analyses you can read. Empty: "No compliance claims are at risk."	
Out-of-date documents	Published policies, procedures and other controlled documents past their next scheduled review (the Next review date the policy list shows). A document with no review scheduled falls back to its review interval, counted from the effective date (or the publish date) or from a later completed review with the result Approved. The count is how many are past due for reissue; documents due within 90 days follow. Up to ten rows, oldest due first, with the effective date below the title, an overdue chip (in days, or in months once 60 days or more late) or a due chip, and an upload button that opens **Upload new version** on that policy (shown only when outside approvals are on and you hold update and review on it). Empty: "Every published document is within its review interval."	policies:read
Personal data breaches	Open incidents marked as personal data breaches, undecided ones first and soonest deadline first, each with a chip for the 72-hour authority notification: time left (red under 24 hours), overdue, notified within 72 hours, notified after 72 hours, not required, or time of awareness not set. The count names how many are past the deadline. Up to ten rows. Empty: "No open personal data breaches."	incidents:read
Data mapping questionnaires outstanding	How many questionnaires are awaiting a Data Liaison's answers (and how many of those are overdue) and how many are awaiting review, then up to eight of them, overdue first, each with the Data Liaison and answer-by date or "to review". Click one to open its answers. Counts only questionnaires you can see (see Data mapping questionnaires). Empty: "No questionnaires outstanding."	dataPrivacy:read
Risks with failing controls	How many **open** risks are treated by failing or overdue controls (**Controls failing**, see The Risks list), then up to ten of them by title, each linked to the risk with the number of failing controls. Counts only risks you can read. Empty: "No open risks have failing controls."	risks:read

Widget	What it shows	Needs
Evidence expiring	How many pieces of evidence you can read have expired and how many expire within 30 days, then up to ten of them, oldest first, each with its type, third party, valid-until date and status (see Evidence). A superseded report and a report about a terminated third party are left out. Empty: "No evidence is expiring."	evidence:read

The "upcoming reviews" widgets use the same due-chip convention everywhere in the product:

Chip	Meaning
Red **Nd overdue**	Past the scheduled date by N days.
Amber **due in Nd**	Due within 14 days.
Gray **due in Nd**	Due later than that.

If a widget cannot load, it shows "This widget failed to load." The rest of the dashboard still renders.

1.4.1 Your own dashboard

Everyone starts on the shared dashboard. Under the **Dashboard** heading a line says which one you see: "The shared layout everyone starts with" or "Your own layout".

To make the dashboard your own, click **Customize my dashboard**. Your copy of the shared layout is made, and a **My widgets** card at the bottom of the page lists its widgets. Changing it needs only Dashboard read and affects nobody else:

1. To add a widget, click **Add widget**, choose the widget type from **Choose a widget type** and click **Add**. It is appended at the end.
2. To rearrange, use the **up** and **down** arrow buttons beside each widget.
3. To remove a widget, click its **trash** button and confirm with **Remove**.

Your own dashboard is never replaced by later default layouts. When the default gains widgets you do not have, a **New widgets available on the default dashboard** note appears above your widgets with an **Add** button for each; **Dismiss** hides the note. A widget you removed, added or dismissed is not offered again. Your edits are your own preference and are not written to the audit log.

Reset to default removes your own layout after you confirm, and you see the shared dashboard again, including the default changes that come with future updates.

1.4.2 Managing the shared dashboard

The shared dashboard is the default for everyone who has not customized theirs. Changing it requires the **dashboard:update** permission at the *all* level (the built-in GRC Manager role has it). If you have it, a **Shared dashboard (the default for everyone)** card at the bottom of the Dashboard lists its widgets, with the same **Add widget**, arrow and **trash** buttons as above. Changes reach everyone on the shared dashboard immediately and are recorded in the audit log.

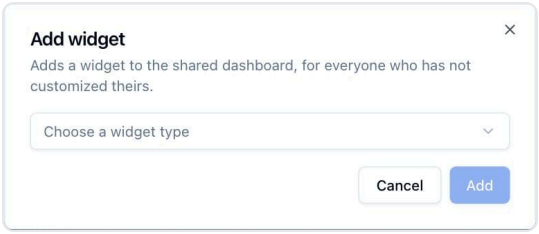


Figure 1.6. The Add widget dialog for the shared dashboard. Choose a widget type, then click Add.

Once the shared dashboard has been changed, later default layouts no longer replace it, so your organization's arrangement is kept; people who customized their own see the new defaults in their **New widgets available** note either way.

1.5 My Tasks

My Tasks in the sidebar (right under **Dashboard**) lists everything waiting on you, from every module you can read, on one page. Every signed-in user has it. When you have tasks, the sidebar entry shows a dark badge with the number (**99+** above 99; on the collapsed sidebar the badge sits on the icon).

Under the **My Tasks** heading a summary line reads either "Nothing is waiting on you." or, for example, "5 tasks waiting on you, 2 overdue. Reviews and control tests appear here 14 days before they are due."

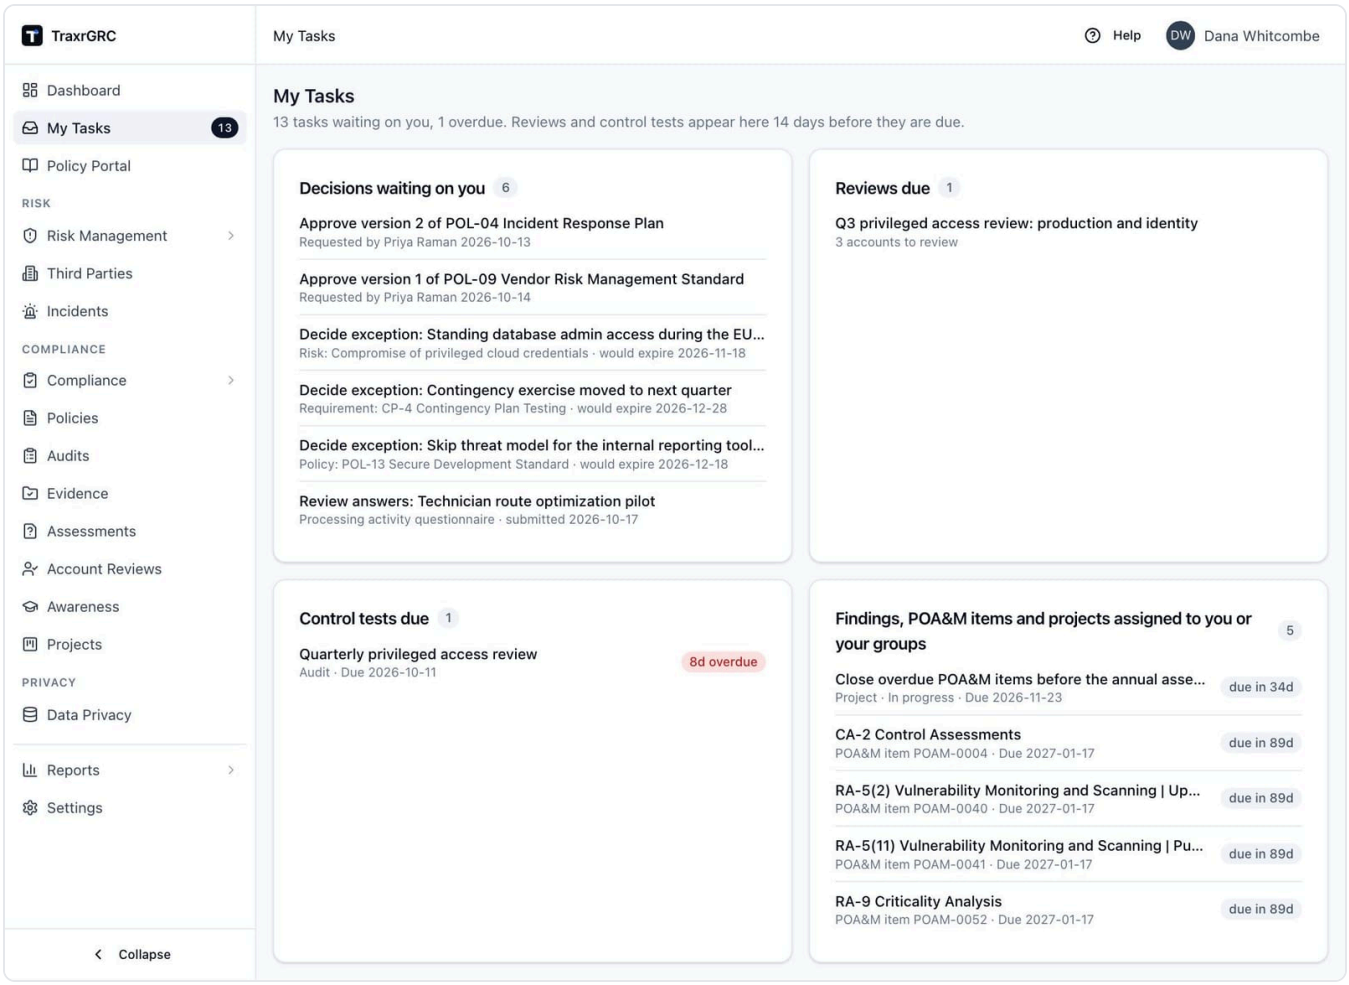


Figure 1.7. My Tasks: everything waiting on you, grouped, overdue first.

Tasks are shown in cards, one per group, each titled with the group name and a count. A group with nothing in it is not shown. Each row's title is a link to the page where you act on the task, opened at the tab for it (a risk's or policy's **Reviews** or **Exceptions**, a policy's **Approvals**, a control's **Maintenance**, a record of processing's **Reviews**); the line below says what it is and, when it has one, "Due *date*". The chip on the right follows the due-chip convention (see The Dashboard): red **Nd overdue**, amber **due in Nd** or **due today** within 14 days, gray beyond that. Within each group, overdue tasks come first, then the soonest due, then tasks without a date.

Group	What it lists
Decisions waiting on you	Policy approvals you can decide ("Approve version N of *policy*"; needs policies:update on the policy, and never your own request unless separation of duties is relaxed); risk, compliance and policy exceptions awaiting your decision (you are the named approver, or none is named and you hold the review right; never your own request under the default separation of duties); DPIAs awaiting your approval (needs dataPrivacy:approve, respects the second-person setting, and waits for the DPO's advice when it is required) and, for the named DPO, DPIAs awaiting your advice; submitted Data Liaison answers you can accept or send back ("Review answers: *activity*"; needs dataPrivacy:review and update on the record, never answers you gave yourself); and personal data breaches you own whose authority decision is still pending, with the 72-hour deadline (overdue once it passes).
Policies to acknowledge	The same items as the **Policies to acknowledge** card on the Policy Portal and the Dashboard.
Training to complete	Active awareness programs with a test that you can open and have not passed (needs awareness:read). The link opens the test.
Questionnaires to answer	Data mapping questionnaires sent to you as a Data Liaison and not yet submitted.
Reviews due	Risk reviews on open risks you own, directly or through an owner group, policy reviews on published policies you own, and periodic reviews of records of processing you own that are not archived, when the review date is within 14 days or past; transfer impact assessment reviews (the assessment's owner, or the vendor's owner when it has none) and signed DPAs of vendors you own, from when their reminders start (30 and 60 days ahead) and once overdue or expired; and account review programs you own that still have pending items ("N accounts to review").
Control tests due	Audits and maintenance scheduled within 14 days or overdue on controls you own that are not retired.
Findings, POA&M items and projects assigned to you or your groups	Open or in-remediation audit findings, open POA&M items (with their due date) and projects not yet completed that you own, directly or through an owner group, with their deadline.

Every group includes only records you can read.

1.6 My profile

Open the user menu (your name, top right) and choose **My profile**. Every signed-in user has a profile page; it holds your personal settings and nothing anyone else can see.

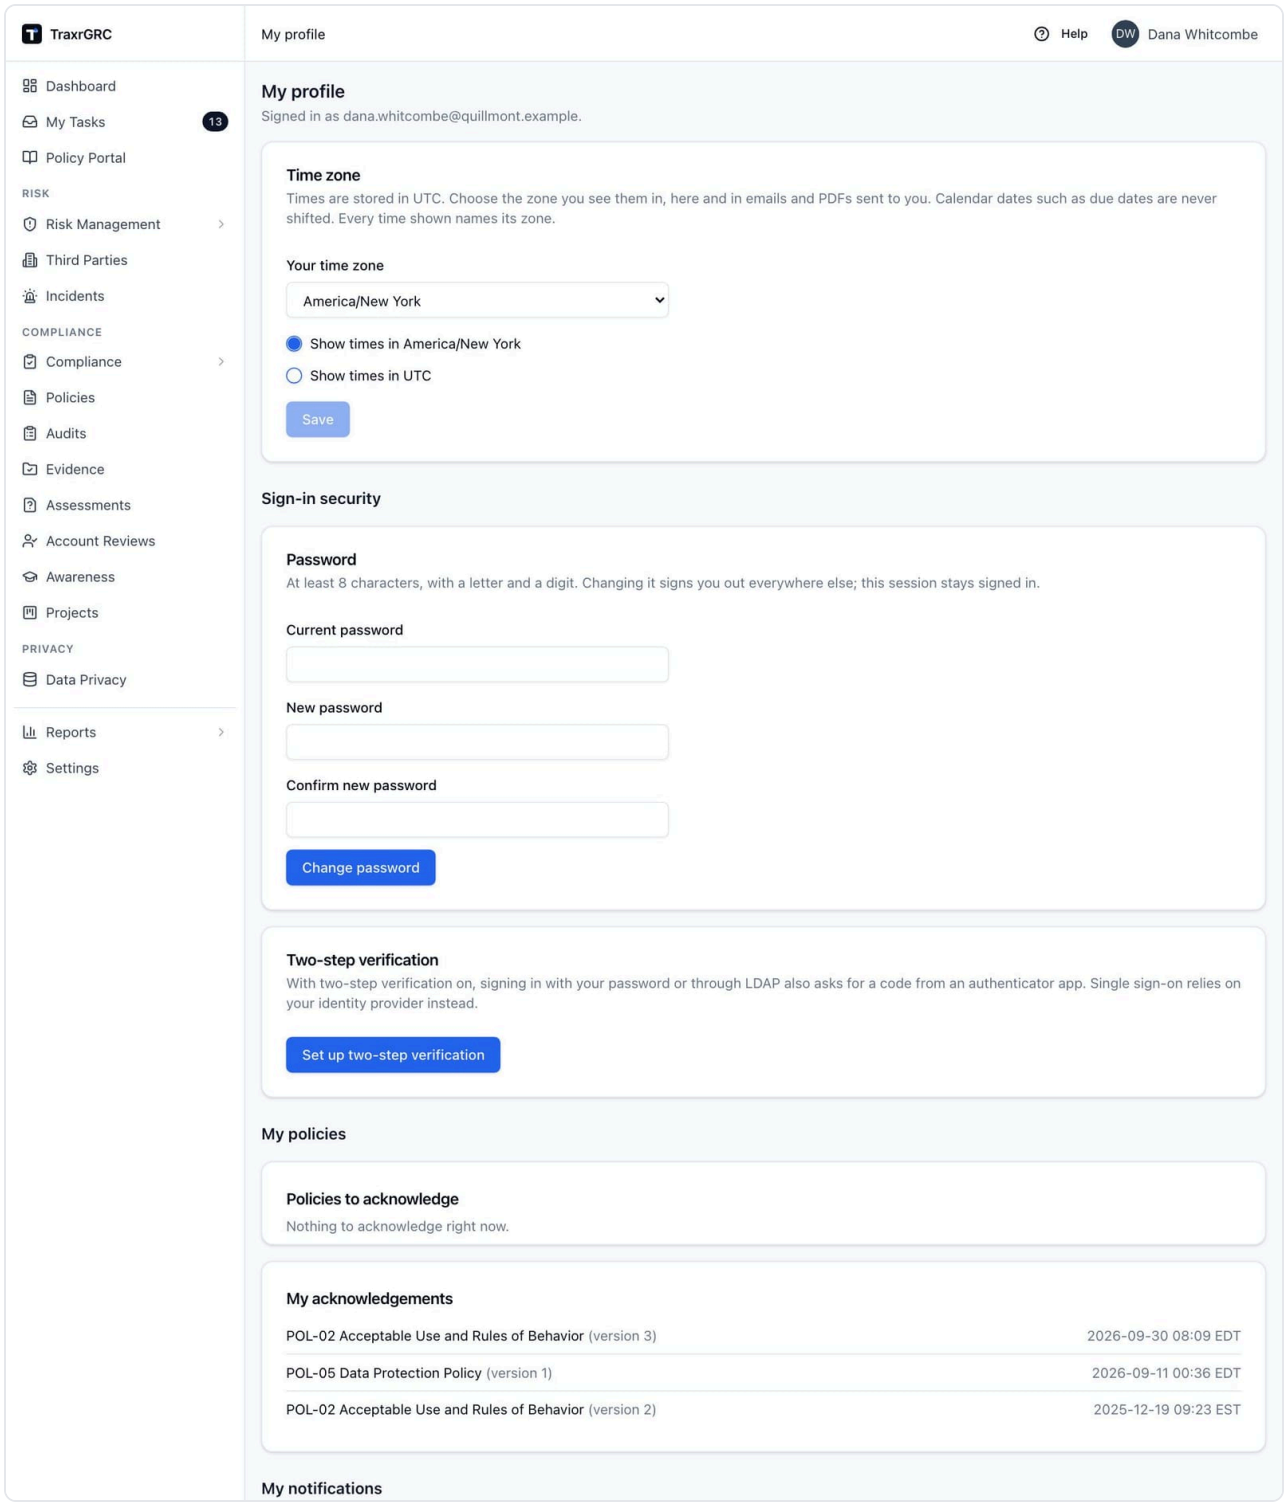


Figure 1.8. My profile, with the Time zone card, the Sign-in security cards and My policies.

The page has up to four parts:

Part	What it holds
Time zone	The zone you see times in, and whether you see them in that zone or in UTC (see Times and time zones).
Sign-in security	The **Password** card, if you sign in with a TraxrGRC password, and the **Two-step verification** card, if you sign in with a password or through LDAP (see Sign-in security). Accounts that sign in only through single sign-on do not see this part.
My policies	**Policies to acknowledge** (the same list as on the Policy Portal) and **My acknowledgements**: every policy version you have acknowledged, newest first, with the time in your zone.
My notifications	**Notification rules** and **Scheduled reports** you can switch off for yourself (see My notifications).

1.6.1 Sign-in security

Changing your password. The **Password** card is shown when your account has a TraxrGRC password:

1. Enter your **Current password**.
2. Enter the new one in **New password** and again in **Confirm new password**. It needs at least 8 characters, with a letter and a digit, and must differ from the current one.
3. Click **Change password**.

A "Password changed" message confirms it. Your sessions on every other browser and device are signed out; this one stays signed in. You receive an email saying when the password was changed; if you did not change it, tell your administrator straight away. After too many wrong current passwords the card says "Too many wrong passwords. Wait 15 minutes and try again." If you sign in through your directory or single sign-on, change your password there. There is no "forgot password" link: if you have forgotten your password, an administrator sets a new one.

Setting up two-step verification. Two-step verification adds a code from an authenticator app (Microsoft Authenticator, Google Authenticator, 1Password and others) to password and directory sign-in. Your administrator may require it; the card then says "Your administrator requires it."

1. In the **Two-step verification** card, click **Set up two-step verification**.
2. In your authenticator app, add an account and scan the QR code shown. If you cannot scan it, type the setup key shown under it, or click **Open in authenticator app** on the device that has the app.
3. Enter the six-digit code the app shows and click **Confirm**. The key expires if you do not confirm within 15 minutes; **Get a new key** starts again.

4. Ten recovery codes appear. Save them somewhere safe, such as a password manager: each works once in place of a code from your app, and they are not shown again. Click **I have saved my recovery codes**.

From then on the card shows since when two-step verification has been on and how many of the 10 recovery codes are left, and every password or directory sign-in asks for a code. You cannot turn it off or get new recovery codes yourself: if you lose your device or run out of codes, ask an administrator to reset two-step verification, then set it up again here. You receive an email when an administrator resets it.

1.6.2 My notifications

Under **My notifications**, the **Notification rules** card lists every enabled notification rule your administrators have set up, with what triggers it. Each has a switch: turn it off and that rule stops emailing you; turn it on again to resume. Reminders about your own reviews and tasks (the fixed notifications) and the security alerts sent to administrators are not rules and cannot be switched off here.

Below the rules, **Scheduled reports** lists every scheduled report that names you (directly, through one of your groups, or through a role one of your groups carries). Turn one off to unsubscribe; turn it on to receive it again. If none names you, the card says "No scheduled reports are sent to you."

1.7 The Policy Portal

Policy Portal in the sidebar is open to every signed-in user, whatever their role. It is where staff read the policies in force and acknowledge the ones they are asked to. A person whose only access is the portal sees just that entry and lands on the portal after signing in.

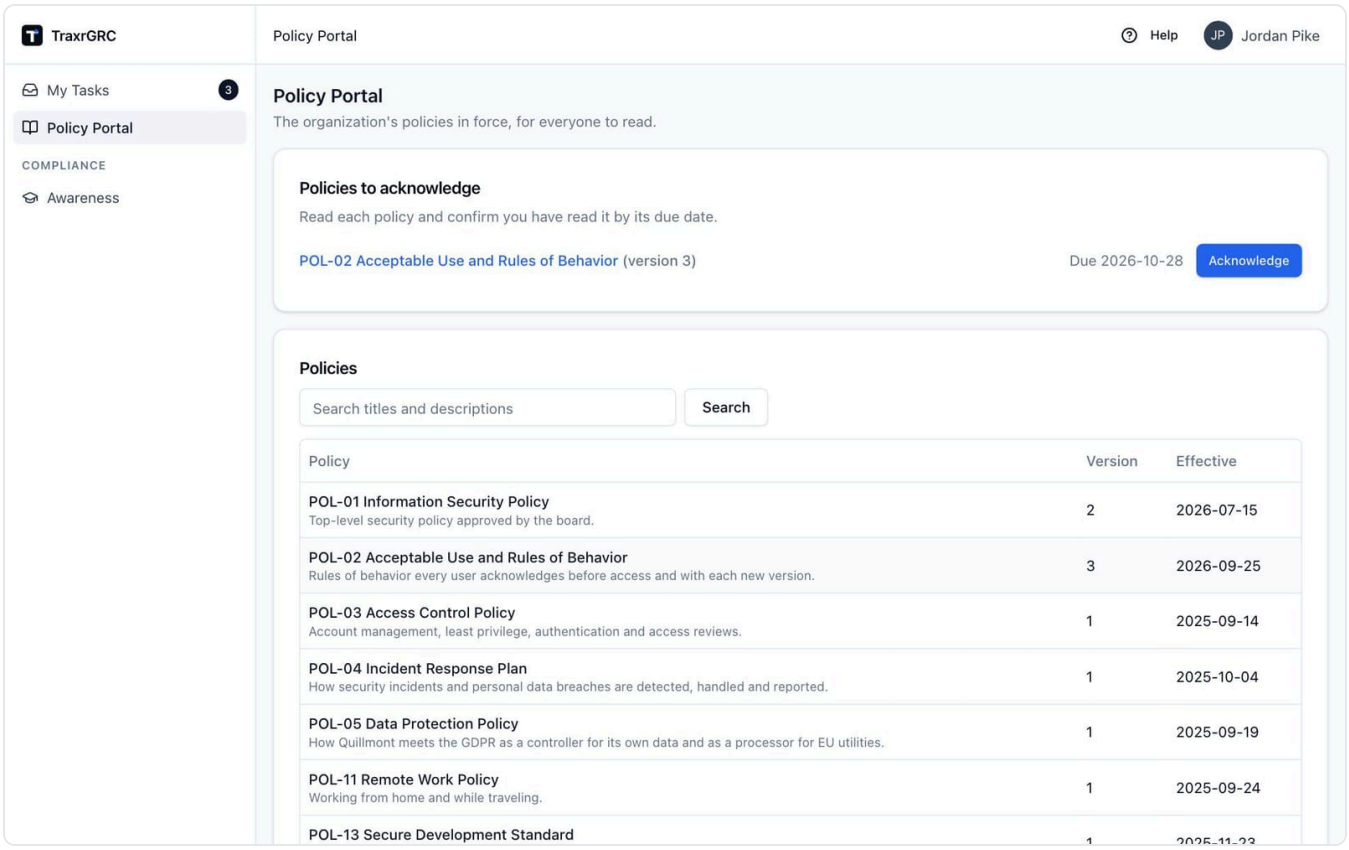


Figure 1.9. The Policy Portal as a staff member sees it, with Policies to acknowledge above the list of published policies.

The portal page shows:

- **Policies to acknowledge**: each policy version you have been asked to acknowledge, with its due date (red "Overdue, was due ..." once late). Click a title to read it, or **Acknowledge** (red once overdue) to go straight to the acknowledgment below the policy text. When there is nothing to do, the card says "Nothing to acknowledge right now."
- **Policies**: every policy published to the portal, with its **Version** and **Effective** date. Type in **Search titles and descriptions** and click **Search** to narrow the list. Click a title to read the policy.
- **My acknowledgements**: your most recent acknowledgements, once you have made any.

Reading and acknowledging a policy, and how policies reach the portal, are described in Policy portal and acknowledgements (Policy Management).

1.8 Shared UI conventions

The same building blocks appear on every list page and record page. Learning them once covers all modules.

1.8.1 Data tables

Module list pages use a consistent table:

- **Global search.** A **Search…** box above the table matches text across the searchable columns at once.
- **Column filters. Filters** in the toolbar opens a panel with one filter per column, each labeled with its column: free-text boxes ("Contains…"), dropdown lists (starting with **All**) for fixed-choice columns, and date pickers for date columns. Filters apply as you enter them and the list jumps back to page 1. The button shows how many filters are active, and each one appears as a chip under the toolbar: click a chip's cross to remove that filter, or **Clear all** to remove them all (the search text stays).
- **Sorting.** Click a column heading to sort ascending; click again for descending. The arrow beside the heading shows the current direction. Sorting also returns to page 1.
- **Pagination.** The footer shows "N rows · page X of Y". Choose **10 / page**, **25 / page** (the default), **50 / page** or **100 / page**, and step with the chevron buttons.
- **Column visibility.** The **Columns** button opens a checklist to show or hide individual columns (the actions column always stays).
- **Saved filters.** The **Saved** button (bookmark icon) lists the views you have stored for that table, or "No saved filters". Click **Save current view…**, give it a name, and the current search text, column filters, sorting, page size and column visibility are stored on your account. Click a saved name to reapply it, or its trash icon to delete it. Saved filters are personal: other users do not see them.
- **CSV export. Export CSV** downloads a comma-separated file of **all** rows matching your current filters and sort, up to the first 5,000 rows, in place of just the visible page; when there are more, a notice says how many were left out, and narrowing the filters exports the rest. Any text cell that starts with `=`, `+`, `–`, `@` or a tab is written with a leading apostrophe (`'`), so a spreadsheet shows it as text instead of running it as a formula; numbers are written as they are.
- Your page, page size, sort and filters are kept in the browser address bar, so you can bookmark a filtered view or share the link with a colleague (they still need permission to see the records).

While data loads, placeholder rows shimmer; a failure shows "Failed to load data" with a **Retry** link. When a search or filter matches nothing, the table shows "No matching records" with a **Clear all** button that clears the search and every column filter; a list with no records at all says so: the main lists (risks, controls, policies, incidents, third parties, evidence, audit plans and projects) explain what the list is for and, if you may create one, offer the same **New** button as the page header; other tables show "Nothing here yet."

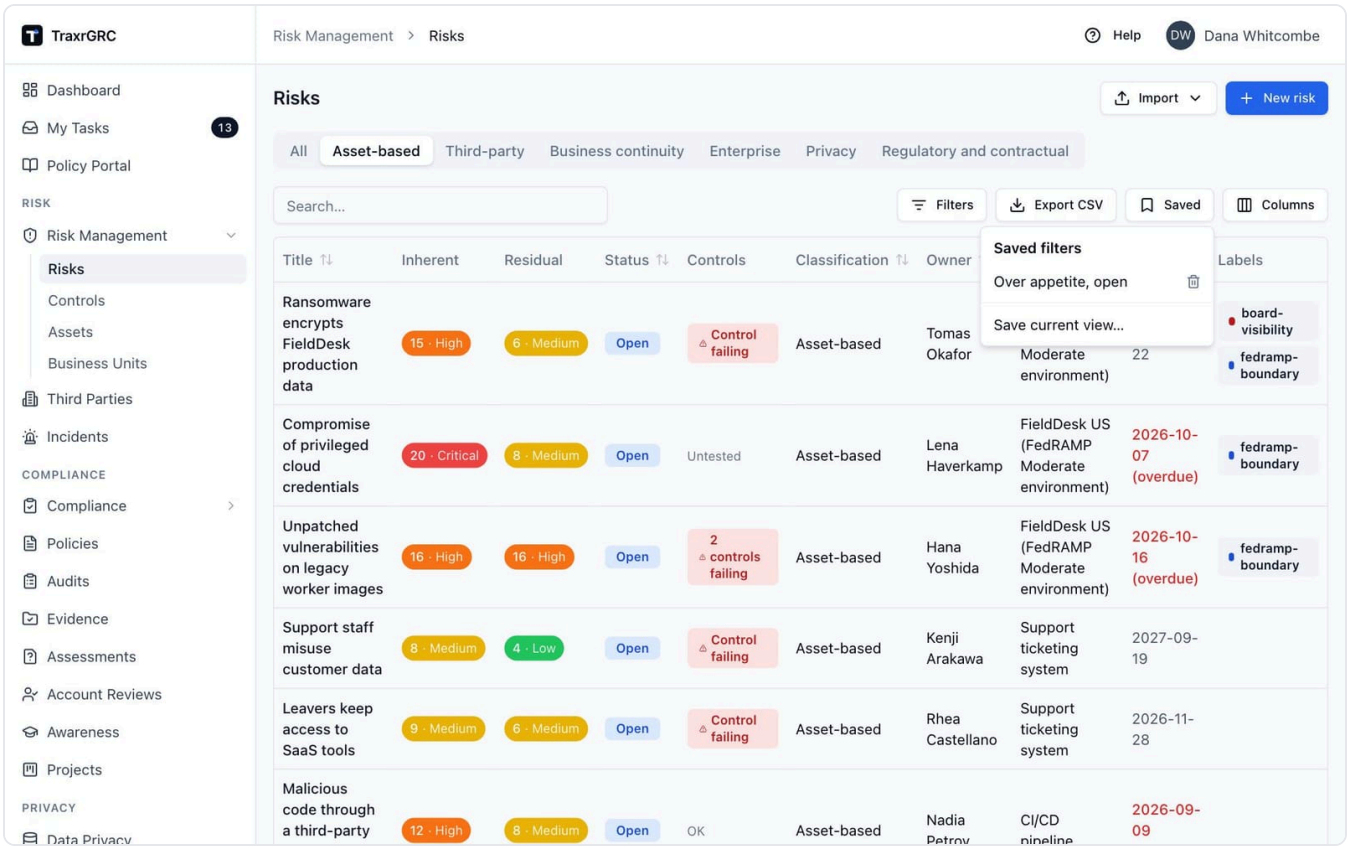


Figure 1.10. The risk register table with the Saved menu open, showing a saved view named Over appetite, open.

1.8.2 Importing from CSV

Risks, controls, assets, policies, third parties, incidents, projects and audit findings can be created and updated in bulk from a CSV file, and risk, policy and compliance exceptions created. Click **Import** beside the list's **New** button (on Risks, Policies and Compliance a menu offers the record type; audit findings are on the Audits page). The button appears only when you may create or update those records, or, for exceptions, update their parent records.

1. Click **Import** (and choose the record type if a menu opens). The **Import** dialog for that record type opens.
2. **Download template** gives the columns for that record type, including its active custom fields (headed `Custom: <field name>`), with one example row. **Download current records** gives the records you can see in the same columns with their ids filled in (at most 5,000): edit it and import it back to change many records at once. The **Columns** panel in the dialog explains each column.
3. Choose your file and click **Check file**. Every row is run through the same checks as the forms, and the table shows what each row would do: **Create**, **Update** (with the columns that change), **Unchanged**, or **Error** with every reason, by line number. Nothing is saved by a check.
4. When no row has an error, click **Import** (the button shows how many rows it will write). All rows are applied together: if anything fails at that moment, nothing is imported and the table shows why.

5. Click **Done** to reload the list.

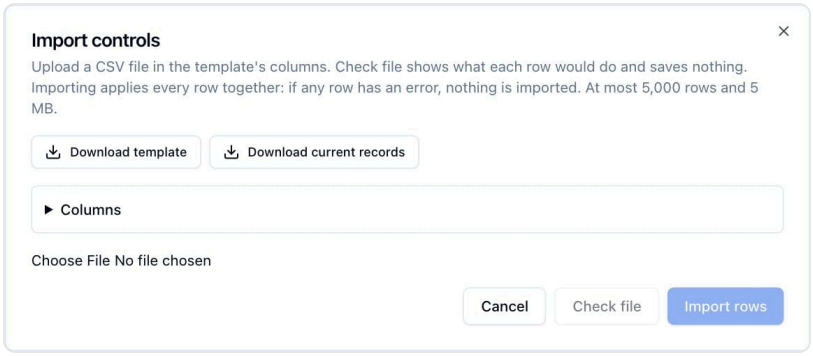


Figure 1.11. The Import controls dialog with the Download template and Download current records buttons.

Rules:

- A row with an `id` updates that record; a blank `id` creates one. Third parties without an id are matched by name, which is unique. Records are never deleted. Exceptions are only created; decide them in the app as usual.
- A blank cell leaves an existing record's value unchanged, so a file can carry only the columns you want to change. Values cannot be cleared by import.
- Owners and approvers are active users, named by email address or full name (use the email when two people share a name). Owner groups, classifications, business units and the risk scale levels are named as they appear in the app. Linked records (a risk's asset, an incident's asset or third party, a finding's control, an exception's risk or policy) are given by name or id and must be records you can read; a name that matches several records asks for the id. Compliance exceptions name the package and the requirement code.
- Dates are written `YYYY-MM-DD`; checkbox custom fields take yes or no; multi-select custom fields separate options with a semicolon.
- You can update only records you could edit in the app: with "own" access, only your own. An id you cannot read is reported as not found.
- Policies import their title, description, owners and review frequency only; content and approvals stay in the app. Personal data breach details on incidents and the data map fields on assets are not imported. Risks' and controls' other details (treatments, testing schedules) are edited in the app.
- Files are limited to 5,000 rows and 5 MB. Text that a spreadsheet would treat as a formula is stored as plain text, and exports still neutralize it; a value exported with a leading apostrophe imports without it.
- Each created or updated record's History tab shows the change as usual, marked as coming from the import, and one summary entry records the import itself. Imports send no notification rule emails unless an administrator has turned on **Notify on API changes and CSV imports**.

1.8.3 Status badges and overdue highlighting

Status values are shown as small colored pills so you can scan a list quickly:

Color	Used for	Examples
Green	Live or complete states	Active, Compliant, Completed
Red	Negative states	Not compliant, Open findings
Amber	In-between states	In remediation, Pending, Sent, Partially compliant
Gray	Inactive end states	Terminated, Closed, Expired, Not reviewed

Exact colors per status are listed in each module chapter.

Dates that matter are emphasized: in the risk register, the **Next review** date of an open risk turns red and gains "(overdue)" once it has passed, and the Dashboard due chips described above flag overdue and soon-due reviews the same way.

1.8.4 Labels

Labels are colored tags for grouping records across modules. Eight record types carry them: risks, controls, assets, third parties, incidents, policies, audit findings and projects. The record pages of risks, controls, third parties, incidents, policies and projects offer a **Labels** control (a tag icon) below the description:

1. Click **Labels**. A panel opens with every label and a **Search or create...** box; labels already on the record show a checkmark.
2. Click a label to attach or detach it. The change saves immediately.
3. To create a label on the fly, type a name that does not exist yet and choose **Create "name"**; it is created with a default gray color and attached at once.

The button shows up to three attached labels, then "+N" for the rest. If you only have read access to the record, labels appear as read-only colored badges instead of the picker. Assets and audit findings have no record page, so the asset list and an audit plan's **Findings** tab show the same picker in each row's **Labels** column, where you may update that row.

Every list of these eight record types has a **Labels** column showing each record's labels as colored badges. Type in its filter to list only records carrying a label whose name contains that text. **Export CSV** writes the label names, separated by "; ". Label definitions (names and colors) are maintained under Settings → Labels (see Administration & Settings).

1.8.5 Attachments

Record pages carry an **Attachments** section listing supporting files, with an **Upload file** button when you have permission to add files to that record:

1. Click **Upload file** and pick a file. The limit is 25 MB.

2. Allowed file types: PDF, PNG, JPG, JPEG, GIF, WEBP, BMP, TXT, CSV, LOG, DOC, DOCX, XLS, XLSX, PPT, PPTX and ZIP.
3. If the upload is rejected you are told why: the file is empty, exceeds the 25 MB limit, is not an allowed type, its type does not match its extension, its contents do not match its extension (a PDF must start like a PDF, a Word or Excel file like one, and a text or CSV file must be text), or it would take your attachments past your storage limit (2 GB per person unless an administrator changed it under Settings → Sessions).

The list shows each file's name (click to download), size, who uploaded it and when. A file can be deleted by anyone with delete permission on the record, and by whoever uploaded it for as long as they can still edit the record; deletion asks for confirmation and permanently removes the file. A document a published policy version was made from (Upload new version) is kept with that version and cannot be deleted, by administrators either. Deleting a record removes its files with it. Uploads and deletions are both recorded in the audit log (see Administration & Settings).

Files can also arrive through an AI assistant's upload link (see AI assistants under Administration & Settings); they appear in the same list.

1.8.6 Comments

Record pages include a **Comments** thread:

1. Type in the **Add a comment...** box (up to 4,000 characters; a live counter shows your usage).
2. Click **Post comment**.

Comments appear oldest first with the author's name and how long ago they were posted. You can delete your own comments; administrators can delete any. If a former user's account has been removed, their old comments show "Deleted user" as the author.

To bring a colleague into the conversation, write `@` followed by their email address, for example `@jane@example.com`. When an administrator has set up a notification rule for comments that includes mentioned people (see Notification rules under Administration & Settings), the colleague gets an email saying they were mentioned, with a link to the record, provided they are allowed to see that record. The email never contains the comment itself.

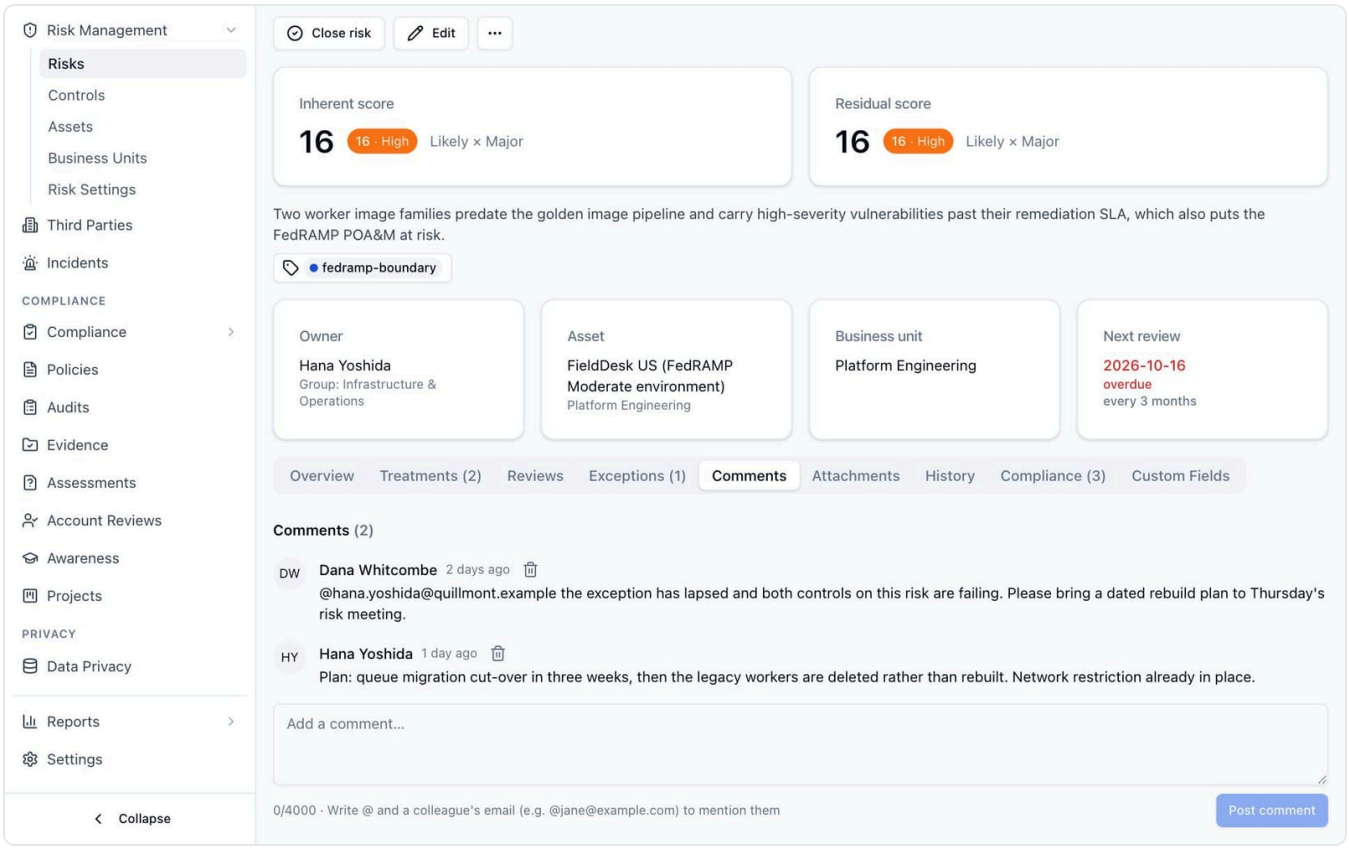


Figure 1.12. The Comments tab of a risk, with a comment that mentions a colleague by email address and the reply below it.

1.8.7 Times and time zones

Every time is stored in UTC. Each person chooses whether to see times in UTC or in their own time zone, and every time on screen, in emails sent to them and in PDFs they download names its zone, for example `2026-10-14 09:05 EDT` or `2026-10-14 13:05 UTC`. Until you choose a zone you see UTC.

Calendar dates are never shifted: a due date, review date, deadline, expiry or effective date is the same day for everyone, whatever their zone. Only moments (created, updated, sent, signed off, completed) move with the zone.

Which day it is today comes from the organization's time zone (set by an administrator under Settings, Notifications; US Eastern unless changed), the same for everyone. Something due today is overdue only once that day is over there, an exception or certificate is valid through its expiry or valid-until date, a new review date counts from today there, and a POA&M item's opened date is the day it opened there. Form defaults such as "today" and the dates in exported file names use the same day. CSV exports, the CSV import templates, the REST API and the AI assistant connection keep UTC (ISO timestamps, `YYYY-MM-DD` dates).

To choose your zone:

1. Open the user menu (your name, top right) and choose **My profile**.

2. In the **Time zone** card, pick **Your time zone** from the list (IANA names such as `Europe/London`).
3. Choose **Show times in** your zone, or **Show times in UTC**.
4. Click **Save**.

If your organization's directory, single sign-on or SCIM provisioning supplies your zone, the list is grayed out with "From the LDAP directory", "From the SSO identity provider" or "From SCIM provisioning", updated at each sign-in: the zone is changed there, not here. You can still choose to see times in UTC. An administrator can set the zone of an account that signs in with a password (Settings → Users). Every change of zone is recorded in the audit log.

1.8.8 History tabs

Every record page has a **History** tab: the audit trail for that single record. Each entry shows what happened (for example created, updated, deleted), who did it, when, and which fields changed. It answers "who touched this record and what did they change" without leaving the page. The organization-wide audit log for all modules lives under Settings → Audit Log (see Administration & Settings).

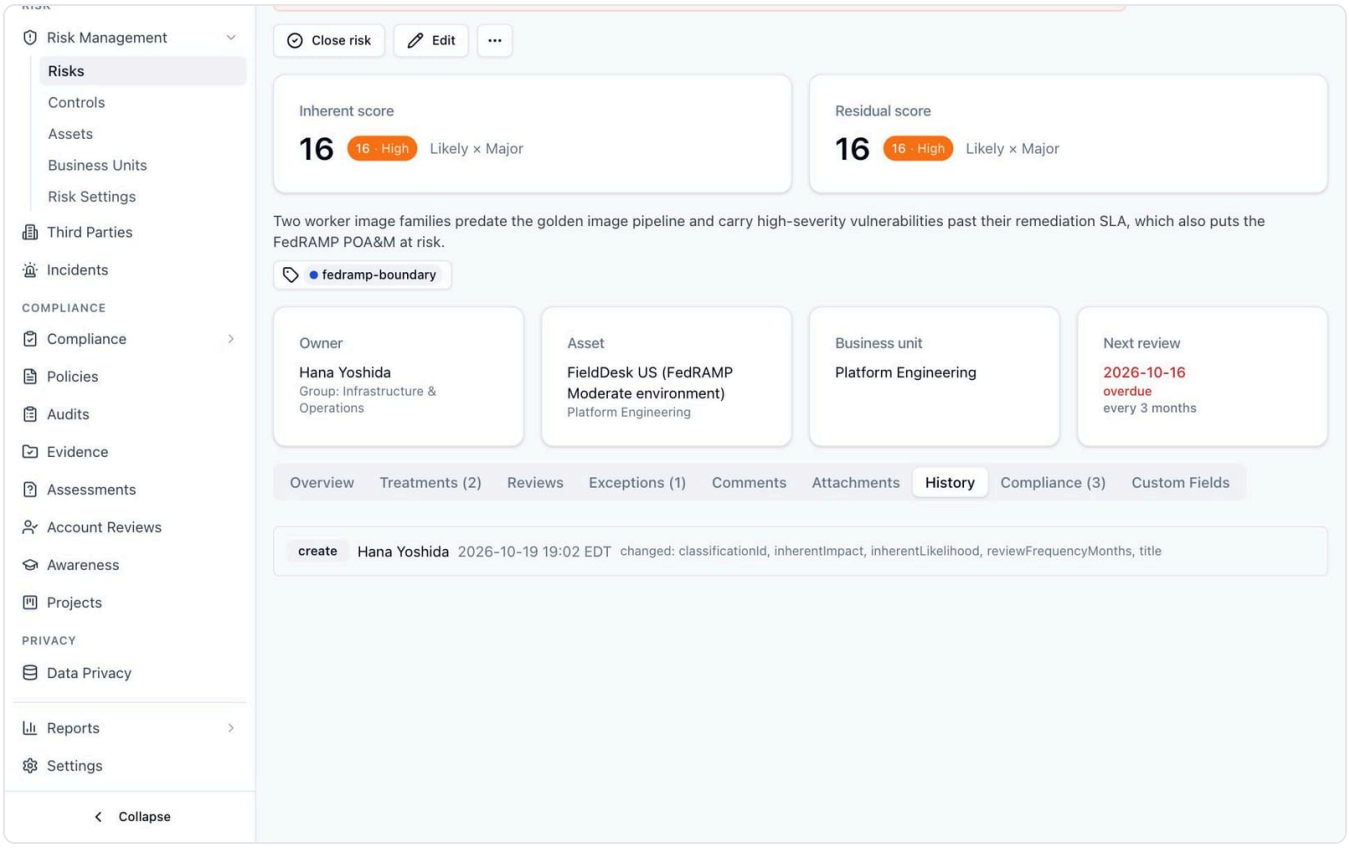


Figure 1.13. The History tab of a risk, listing each change with who made it, when and which fields changed.

1.8.9 Deleting a record

On a record's page, **Delete** is not among the action buttons. It sits in the **More actions** menu (the **...** button at the end of the button row), in red, so it cannot be clicked by mistake next to **Edit**. The menu appears only when you hold the delete right on the record. Every record delete then asks you to confirm, and says what goes with the record. This applies to risks, policies, third parties, incidents, audit plans, projects, awareness programs, DPIAs, records of processing, account review programs, evidence and assessment templates. Rows inside a record (contacts, treatments, attachments and the like) keep their own trash icons.

1.8.10 When a page is off limits

If you open a page you do not have permission for, you see a page headed **403: Forbidden** stating "You do not have permission to view this page. Ask your administrator to grant you access." **Go to My Tasks** and **Policy Portal** take you back to pages everyone can open; you can also use the sidebar, or ask your administrator for the missing permission. If you are signed out entirely you are sent to the sign-in page instead, and returned to the page after signing in.

A link to something that does not exist shows **Page not found** with a **Back to dashboard** button.

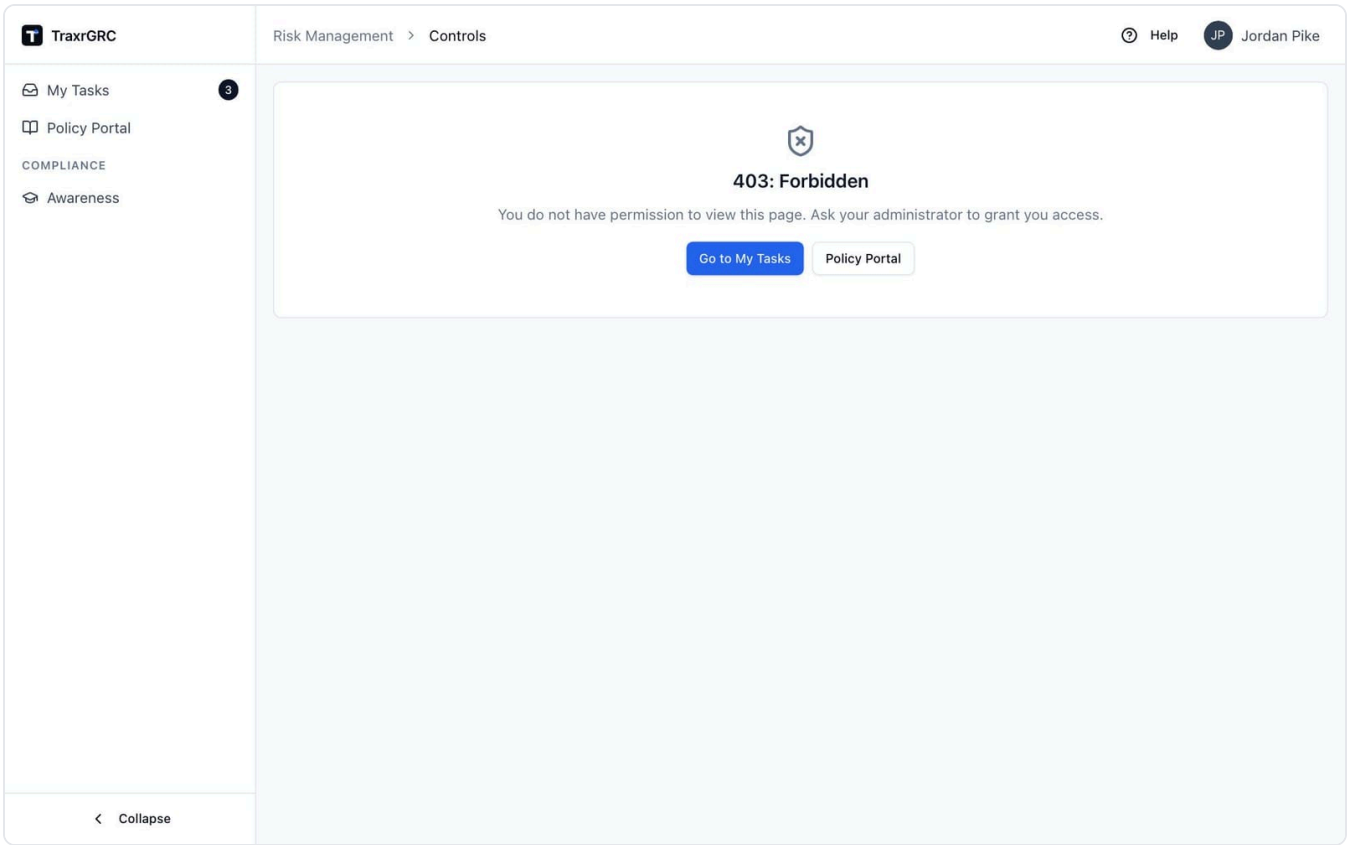


Figure 1.14. The 403: Forbidden page shown to a staff member who opened the controls list, with the Go to My Tasks and Policy Portal buttons.

CHAPTER 2

Administration & Settings

Settings is where you run the platform itself: accounts and sign-in, permissions, the shared vocabulary your records use (labels and custom fields), email and single sign-on connections, provisioning, sessions, branding, AI assistant access, notification scheduling, the audit trail, API access and backups.

This chapter is for administrators. Most Settings screens check the **settings** module right (settings:read to open them, settings:create or settings:update to change them), and several changes are reserved for administrator accounts whatever a role grants. None of the built-in roles holds settings rights. A few screens check their own module instead, and the built-in roles do grant some of those, so a GRC Manager, Risk Owner or Auditor also sees the **Settings** entry in the sidebar, with only the cards for those screens:

Screen	Right checked to open it	Built-in roles that hold it
Custom Fields	customFields:read (create, update and delete at *all* to change)	GRC Manager (all rights); Risk Owner and Auditor (read)
Notifications	notifications:read (notifications:update to change)	GRC Manager
Audit Log	auditLog:read	GRC Manager, Auditor
API Tokens	apiTokens:read (apiTokens:create to issue tokens)	None
Backups	Administrator account	None

Administrators always pass every permission check, so the procedures below assume you are signed in as an administrator.

Almost everything you change here is recorded automatically: every create, update and delete writes an entry to the Audit Log (see Audit Log below), and saving integration, scheduler or template changes is likewise audited.

2.1 Finding the Settings screens

Click **Settings** at the bottom of the sidebar to open the Settings page ("Platform configuration and identity management"). It shows a card for each Settings screen you may open, grouped under **Identity & Access**, **Governance & Customization**, **Integrations** and **System**; click a card to open it. A group with no card you may open is hidden. The sidebar entry appears for anyone who can open at least one screen. An administrator sees every card:

Card	What it covers
Users	Local accounts, administrators, group membership.
Groups	User groups that carry roles and own records.
Roles	Permission matrices assigned to groups.
Integrations	SMTP email delivery and LDAP / SAML single sign-on.
Provisioning	SCIM 2.0 provisioning from an identity provider such as Microsoft Entra ID.
Sessions	Session length, idle timeout, the longest API token lifetime and separation of duties.
Branding	Product name, logo and color for a white-labeled deployment.
Policies	Whether approvals made outside the app may be recorded.
Privacy	Controller, representative and DPO details for the Article 30 register.
AI assistants	Model Context Protocol access for Copilot, VS Code and Claude, with Entra ID sign-in.
Labels	Colored tags for grouping and filtering records.
Custom Fields	Extra fields on the forms of each module.
Notifications	Email notification rules and the templates they send.
Audit Log	Who changed what and when, with filters and CSV export.
API Tokens	Scoped tokens for the REST API.
Backups	Export and restore database snapshots, and the scheduled weekly backups.

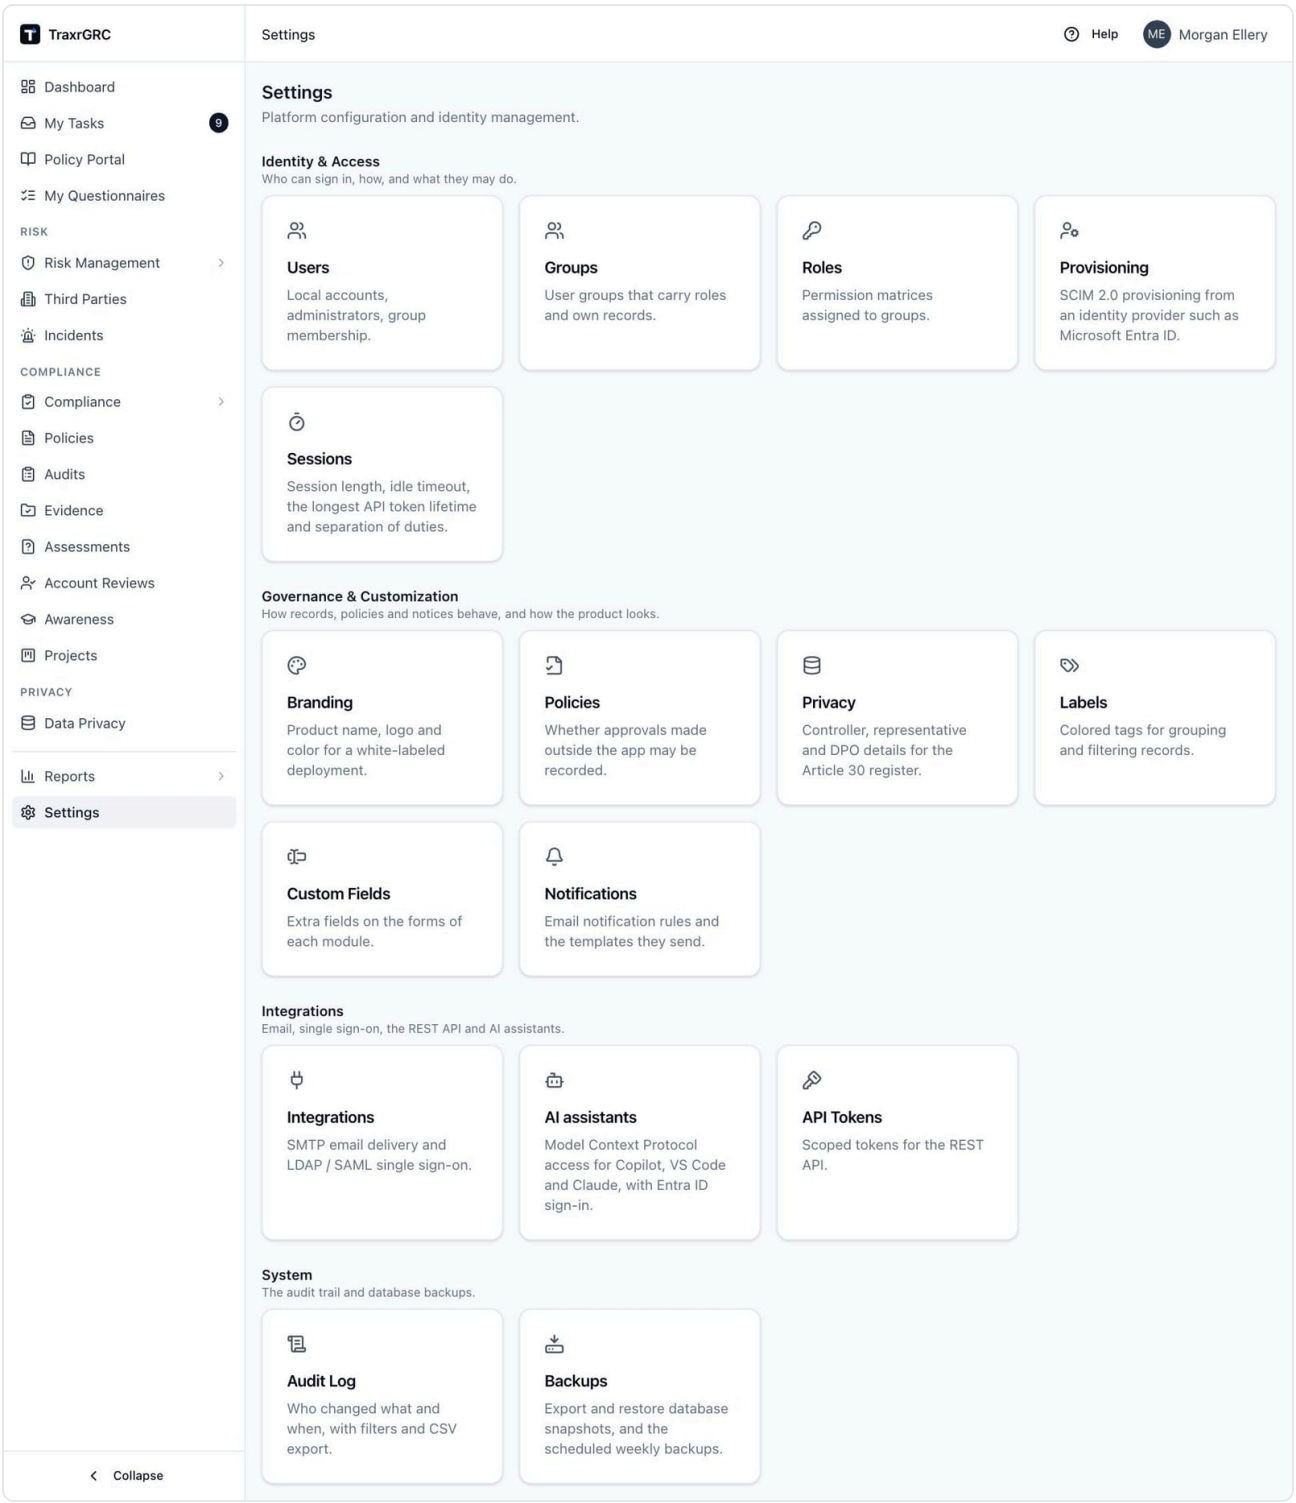


Figure 2.1. The Settings page with its cards grouped under Identity & Access, Governance & Customization, Integrations and System.

Each card appears only to people its screen admits. A GRC Manager, for example, sees **Custom Fields**, **Notifications** and **Audit Log**; an Auditor sees **Custom Fields** and **Audit Log**; and only administrators see **Backups**. Security alert emails link straight to the Audit Log.

2.2 Users

Settings → Users lists every account that can sign in: local accounts and directory-linked identities alike. The list is paginated (10 per page) and shows, for each user:

Column	What it shows
Name	The display name, with a "SCIM" badge for accounts provisioned by SCIM.
Email	The sign-in address, with the user's time zone under it ("(directory)" when it comes from the directory or identity provider).
Groups	The user's groups, as badges.
Admin	A solid "Admin" badge (plus "Break-glass" for break-glass administrators), or a dash.
Active	An outlined "Active" badge, or a red "Inactive" badge, plus an "MFA" badge when the user has two-step verification on.
Actions	Edit (pencil), Reset password (rotate) and Delete (trash).

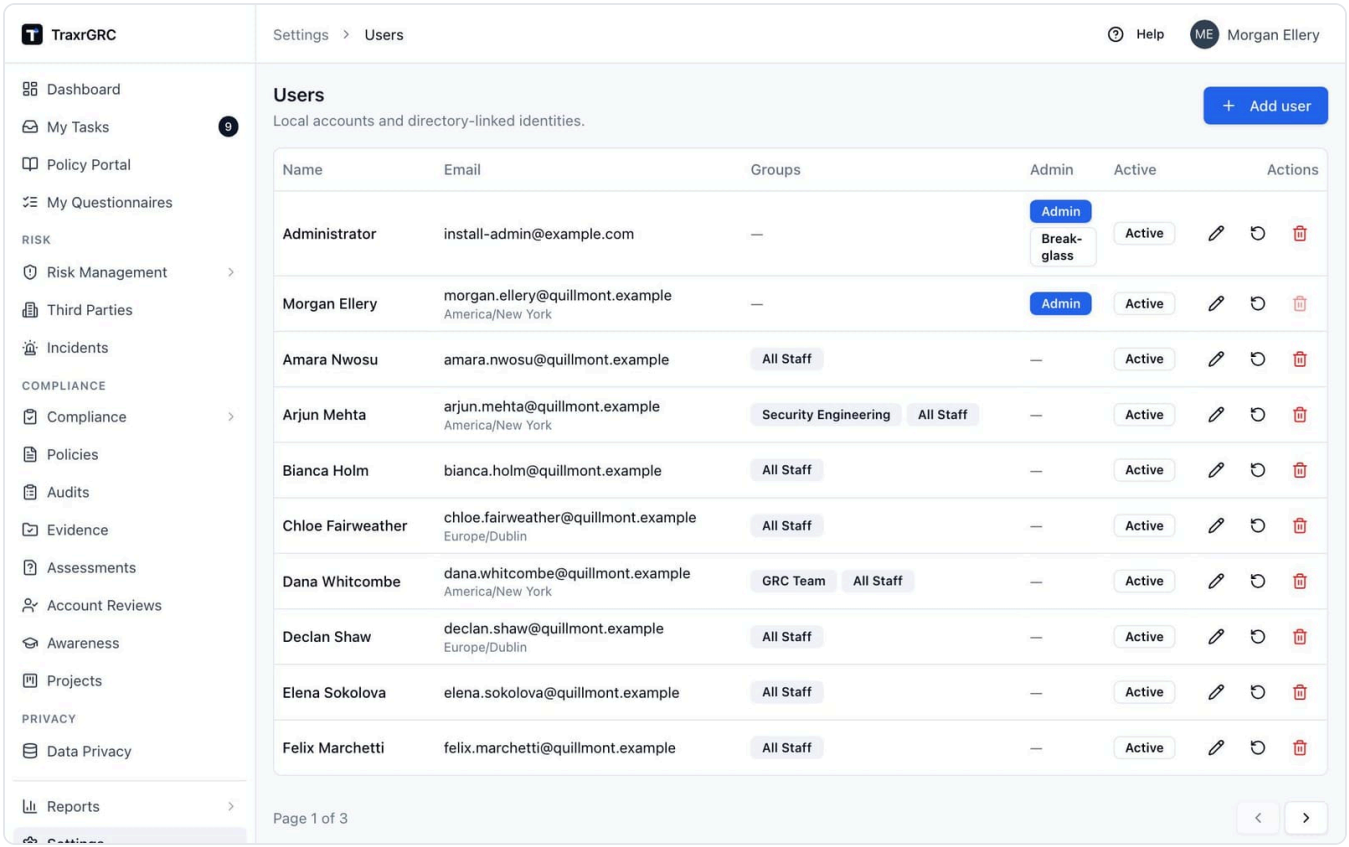


Figure 2.2. The Users list, showing each account's groups, administrator and active badges and the row actions.

To create an account:

1. Click **Add user**. The dialog explains: "Create a local account. The user signs in with email and password."
2. Complete the dialog:

Field	What it means
Email	Sign-in address. Must be unique; it cannot be changed afterward.
Name	The display name shown across TraxrGRC.
Password	Initial password for a local account: at least 8 characters, with at least one letter and one digit.
Administrator	Grants full access to every module and setting; bypasses all permission checks.
Break-glass (emergency password sign-in)	Appears once **Administrator** is on. When SSO-only sign-in is on, only break-glass administrators may sign in with a password, for use when the identity provider is unavailable. The account needs a password (use **Reset password** if it was created by SSO). Shown as a "Break-glass" badge in the list.
Groups	Tick the groups the user belongs to. Group membership is how permissions and record ownership reach the user.

3. Click **Save**. The user can sign in immediately with email and password. An audit entry is written.

To manage an existing user, click the pencil icon on their row. The **Edit user** dialog opens:

- **Name**, **Administrator** and **Active** can be changed, and **Groups** re-ticked.
- The **Active** switch only appears when editing. Turning it off deactivates the account: the user can no longer sign in.
- The email address is shown in the dialog header for reference but is not editable.
- **Allow the next LDAP or SSO sign-in, or SCIM, to link this account** appears for accounts with a password or administrator rights. LDAP, SSO and SCIM provisioning never take such an account over by email on their own; turning this on (administrators only) lets the next directory or SSO sign-in, or SCIM request, with that email link to it, after which the switch turns itself off. The help text shows whether the account is already linked to LDAP, SSO or SCIM provisioning.
- **Time zone** sets the zone the user sees times in ("Not set (times in UTC)" until chosen). It is grayed out when the user's directory or identity provider supplies the zone. Setting it for an administrator's account needs an administrator. The user can still choose to see UTC on their profile. See "Times and time zones" in Getting Started.

- **Sign out everywhere** (in the edit dialog, for other users) ends all of that user's sessions within a minute. Signing out an administrator needs an administrator.
- Turning the account off, changing its administrator rights or changing its groups ends that user's open sessions within a minute (not when you edit your own account). Adding or removing members of a group in Settings → Groups does the same for those members.
- Reactivating an account whose role rights go beyond your own needs an administrator, as does a password reset of such an account.
- **Reset two-step verification** (in the edit dialog, for other users with two-step verification on) is for a lost device or used-up recovery codes. Confirm with **Reset** in the dialog that follows: the user's authenticator app and recovery codes stop working, their sessions are signed out, they are emailed (template `user.mfa.reset`) and every administrator is alerted. They set it up again from My profile, or at their next sign-in when your administrator requires it. It needs the same rights as a password reset. Nobody can reset their own; another administrator must do it.

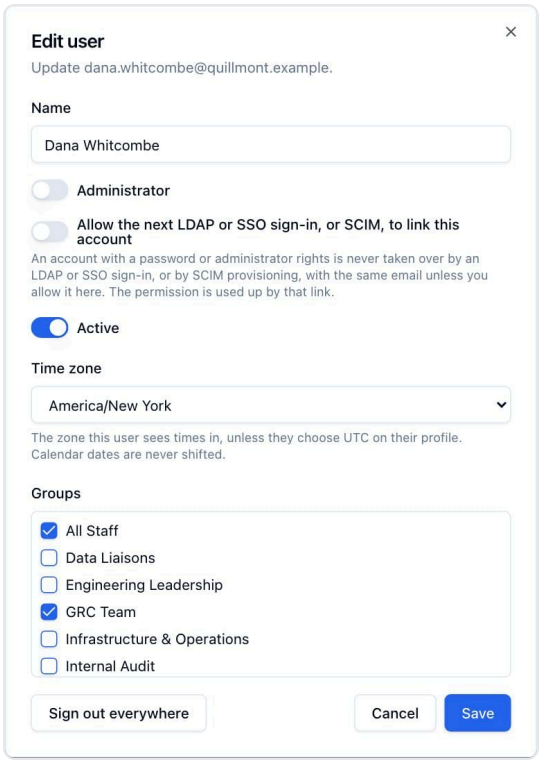


Figure 2.3. The Edit user dialog with the Administrator, Active and Time zone settings and the group list.

To set a new password, click the rotate (reset) icon, enter the **New password** in the **Reset password** dialog (the same rules as above) and click **Save**. The user's open sessions end within a minute and they receive an email (template `user.password.reset`) saying who reset it and when.

To remove an account, click the trash icon and confirm in the **Delete user** dialog. Deleting a user keeps the records they owned, but those records become unassigned.

Two protections keep you from locking yourself out:

- You cannot deactivate or demote your own account, and you cannot delete your own account (the delete button is disabled on your own row, and the switches are disabled when you edit yourself).
- TraxrGRC always requires at least one active administrator: if your change would leave the organization with zero active admins, it is rejected with "At least one active administrator account is required."

2.3 Groups

Settings → Groups bundles users together and attaches roles to them; members inherit the permissions of the roles on their groups (see Roles below). The list shows each group's **Name**, **Description**, **Members** (a count), **Roles** (as badges) and actions.

To create or edit a group, click **Add group** or the pencil icon, complete the fields and click **Save**:

Field	What it means
Name	The group name (must be unique).
Description	Optional free-text purpose, e.g. "Governance, risk and compliance function."
Members	Tick every user who belongs to the group. Users are listed as "Name (email)".
Roles	Tick the roles this group carries. Every member inherits the union of these roles' permissions.

The dialog reminds you of the rule: members inherit the permissions of the roles attached here. Group membership also matters for "own"-level access: records owned by a user's groups count as theirs.

Groups created by SCIM provisioning carry a "SCIM" badge: their name and members come from the identity provider, so in the edit dialog the name and member list are grayed out and only the description and roles can be changed. Likewise their membership cannot be ticked or unticked from Settings → Users.

To delete a group, click the trash icon and confirm in the **Delete group** dialog. Members keep their accounts but lose this group's roles; records are unaffected.

2.4 Roles

Settings → Roles defines permission matrices. A role is a table of access levels, one cell per module and right, and groups carry roles to their members. The list shows each role's **Name** (with a "System" badge for the built-in ones), **Description**, a **Permissions** summary (e.g. "Risks (Read: all, Create: all…)" or "No permissions"), the **Groups** it is attached to, and actions.

TraxrGRC ships with four system roles:

Role	What it grants
GRC Manager	Full access to the GRC modules, including Review on Risks, Compliance, Policies and Data privacy, the Dashboard, Custom fields, Reports, Notifications and read access to the Audit log.
Risk Owner	Manage own risks; read shared catalogs.
Auditor	Read-only across modules, including the Audit log.
Data Liaison	Answers the data mapping questionnaires assigned to them, and nothing else (see Data mapping questionnaires). Give a Data Liaison other roles as well where they need more.

System roles cannot be deleted; the delete button is disabled for them. They can be renamed: TraxrGRC recognizes its built-in roles by a fixed internal key, so a renamed role stays the built-in one.

To create or edit a role:

1. Click **Add role** or the pencil icon.
2. Enter a **Name** and optional **Description**.
3. Set the **permission matrix**. Rows are the 23 modules: Risks, Compliance, Policies, Third parties, Incidents, Audits, Awareness, Projects, Data privacy, Data Liaison questionnaires, Account reviews, Assessments, Custom fields, Reports, Controls, Assets, Business units, Dashboard, Settings, Users, Notifications, Audit log and API tokens. Columns are the rights **Read**, **Create**, **Update**, **Delete** and **Review**.
4. Set each cell's level from its dropdown:

Level	What it means
None	No access to that module/right.
Own	Only records owned by the user or by one of their groups appear in lists and can be opened or changed. For Data Liaison questionnaires, the questionnaires assigned to the user.
All	Every record of that module.

   Own is not offered on modules whose records have no owner (Custom fields, Reports, Dashboard, Settings, Users, Notifications, Audit log and API tokens), since there it would mean the same as All. Set at least one cell away from "None": a role with an empty matrix grants nothing. The matrix scrolls inside the dialog; **Save** stays in view at the bottom.

5. Click **Save**.

The **Review** column applies to Risks (completing risk reviews and deciding risk exceptions with no designated approver), Compliance and Policies (deciding their exceptions with no designated approver; on Policies also recording outside approvals) and Data privacy (the decision on Data Liaison answers); every other row shows a dash in that column. The built-in GRC Manager role has Review on all four; an upgrade adds Compliance and Policies review to it where it still has update at "all" on them. Add Review to any custom role whose holders should decide exceptions, and Data privacy review to the role of the people who review Data Liaison answers.

The **Approve** column applies only to Data privacy: it is the Approve DPIAs right, to approve or reject a DPIA. Every role that held Data privacy review before this right existed was given it at the same level, so nobody lost the ability to decide DPIAs on upgrade; remove it from roles that should not. Approving is separate from being the DPO, who is named in Settings → Privacy.

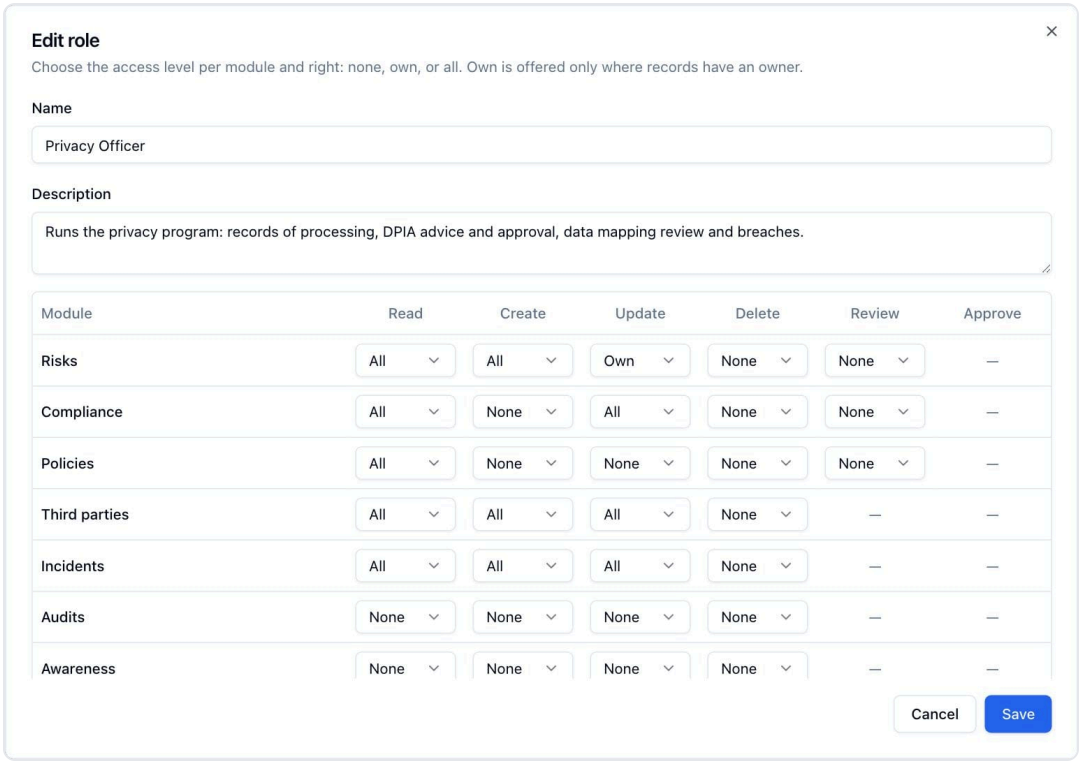


Figure 2.4. The Edit role dialog showing the permission matrix, with a level dropdown for each module and right.

When a user belongs to several groups, their effective permissions are the union of all their roles' matrices, with "all" winning over "own", which wins over "none". Administrators bypass the matrix entirely.

Deleting a role (click the trash icon, then confirm in the **Delete role** dialog) detaches it from its groups; members lose the granted permissions. Groups themselves are kept.

2.5 Labels

Settings → Labels (address `/settings/labels`) manages the shared tag vocabulary. In the application, labels are attached and detached with the **Labels** picker on the record pages of risks, controls, third parties, incidents, policies and projects, and in the **Labels** column of the asset list and of an audit plan's findings (see Labels in Getting Started). You need update rights on the record in its own module; without them the labels are read-only. The screen shows a count of labels and a table with **Name** (colored dot plus name), **Color** (hex value), **Assignments** (how many records carry it), **Created** (date) and actions. Viewing needs settings:read; creating, editing and deleting need settings:create, settings:update and settings:delete respectively, each at *all*.

To create a label:

1. Click **Create label**.
2. Enter a **Name**. Names are lowercase slugs of letters, digits and single dashes, e.g. `critical-vendor`; duplicates are rejected.
3. Pick a **Color**: either one of the ten preset swatches or any custom hex color typed into the field below (for example `#0ea5e9`). The button stays disabled until the name is filled in and the hex is valid.
4. Click **Create label**.

To rename or recolor a label, click its pencil icon. To remove one, click the trash icon and confirm in the **Delete label?** dialog: the label is unassigned from every record it was on. This cannot be undone.

2.6 Custom Fields

Settings → Custom Fields (address `/settings/custom-fields`) extends GRC Platform with your own attributes. Fields are defined per entity type and then appear on every record of that type, where people fill in their values (see below). The screen is organized into tabs, one per entity: **Risks**, **Controls**, **Policies**, **Third parties**, **Incidents**, **Audit plans**, **Projects**, **Awareness programs**, **Records of processing**, **Account review programs**, **Compliance analyses** and **Assessment requests**.

Each tab lists that entity's fields with columns **Name**, **Type**, **Required**, **Options** (for select fields), **Sort**, **Active** and actions. Fields are ordered by their sort value. Opening the screen needs customFields:read; creating, changing and deleting fields need the matching customFields right at the *all* level.

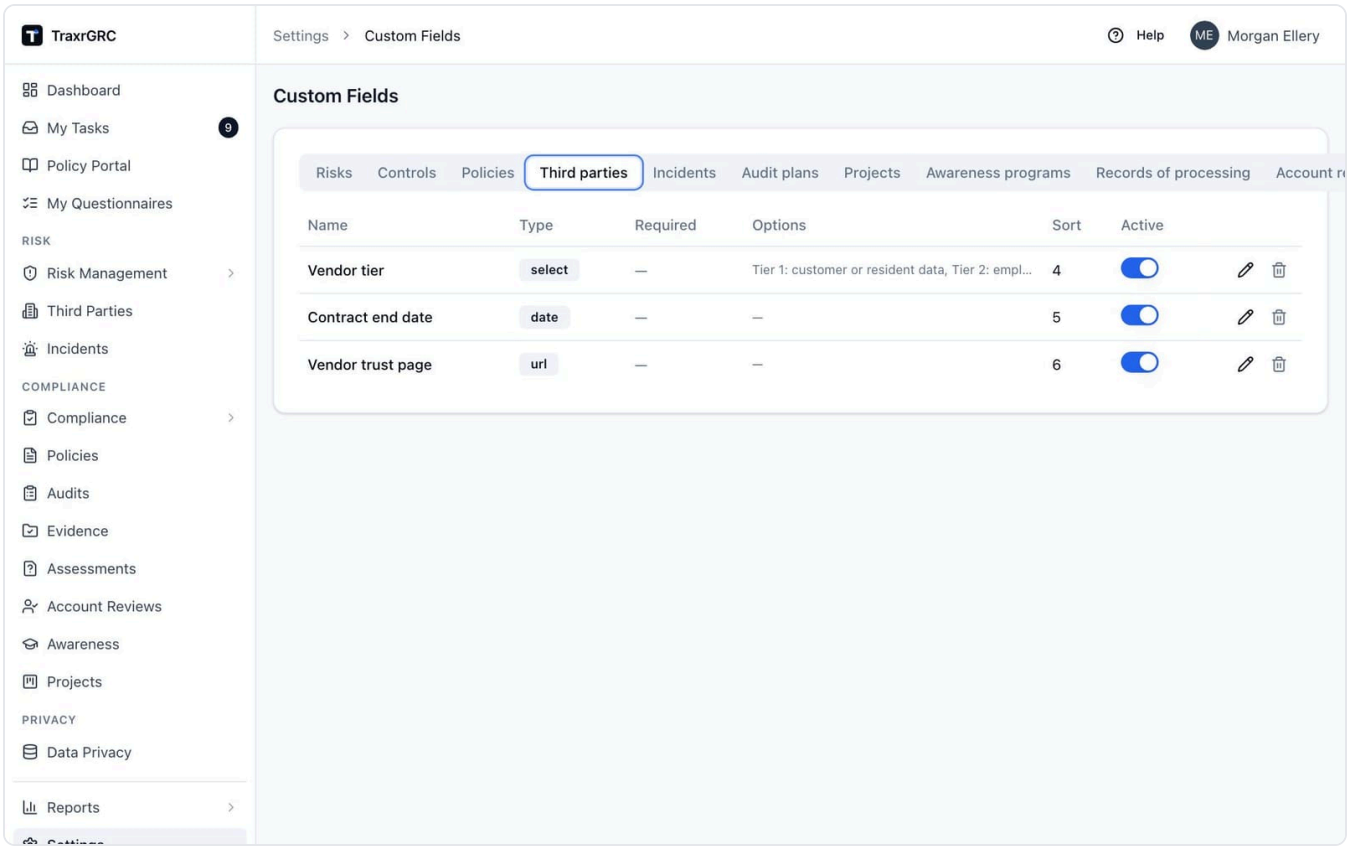


Figure 2.5. The Custom Fields screen on the Third parties tab, listing fields of several types.

To create a field:

1. Open the tab for the entity type and click **Create field**.
2. Complete the **Create custom field** dialog:

Field	What it means
Name	The field label shown on records, e.g. "Mitigation notes". Names must be unique per entity.
Type	One of: **Text**, **Textarea**, **Number**, **Date**, **Select**, **Multi-select**, **Checkbox** or **URL**.
Sort order	A number controlling where the field appears; lower values first. New fields default to the end of the list.
Options	Shown only for Select and Multi-select. Use **Add option** to build the list and the × button to remove one. Options must be unique and at least one is required.
Required	When ticked, records must have a value for this field before they can be saved.

3. Click **Create field**.

In the table, the **Active** switch turns a field off without deleting it (inactive fields no longer appear on records); the pencil icon edits the definition (save with **Save changes**); the trash icon deletes the field. Deleting also destroys every value already stored against it: the **Delete custom field?** confirmation warns about this, and it cannot be undone.

Where values are filled in: open any record page (a risk, policy, audit plan, record of processing and so on) and switch to its **Custom Fields** tab. The card lists each active field (required fields are marked with a red asterisk); click **Edit** to open the value dialog, save, and the values are stored with the record. Whoever edits values needs update rights on that record's module.

2.7 Integrations

Settings → Integrations connects TraxrGRC to your mail server and identity providers. It holds three cards: **SMTP**, **LDAP** and **SAML SSO**, each with its own **Save** button. Saving on any card writes an audit entry and sends every administrator a security alert. The SMTP password, LDAP bind password and SAML private key are stored encrypted and never sent back to the browser: a stored one shows "Stored. Type to replace." and is kept unless you type a new value.

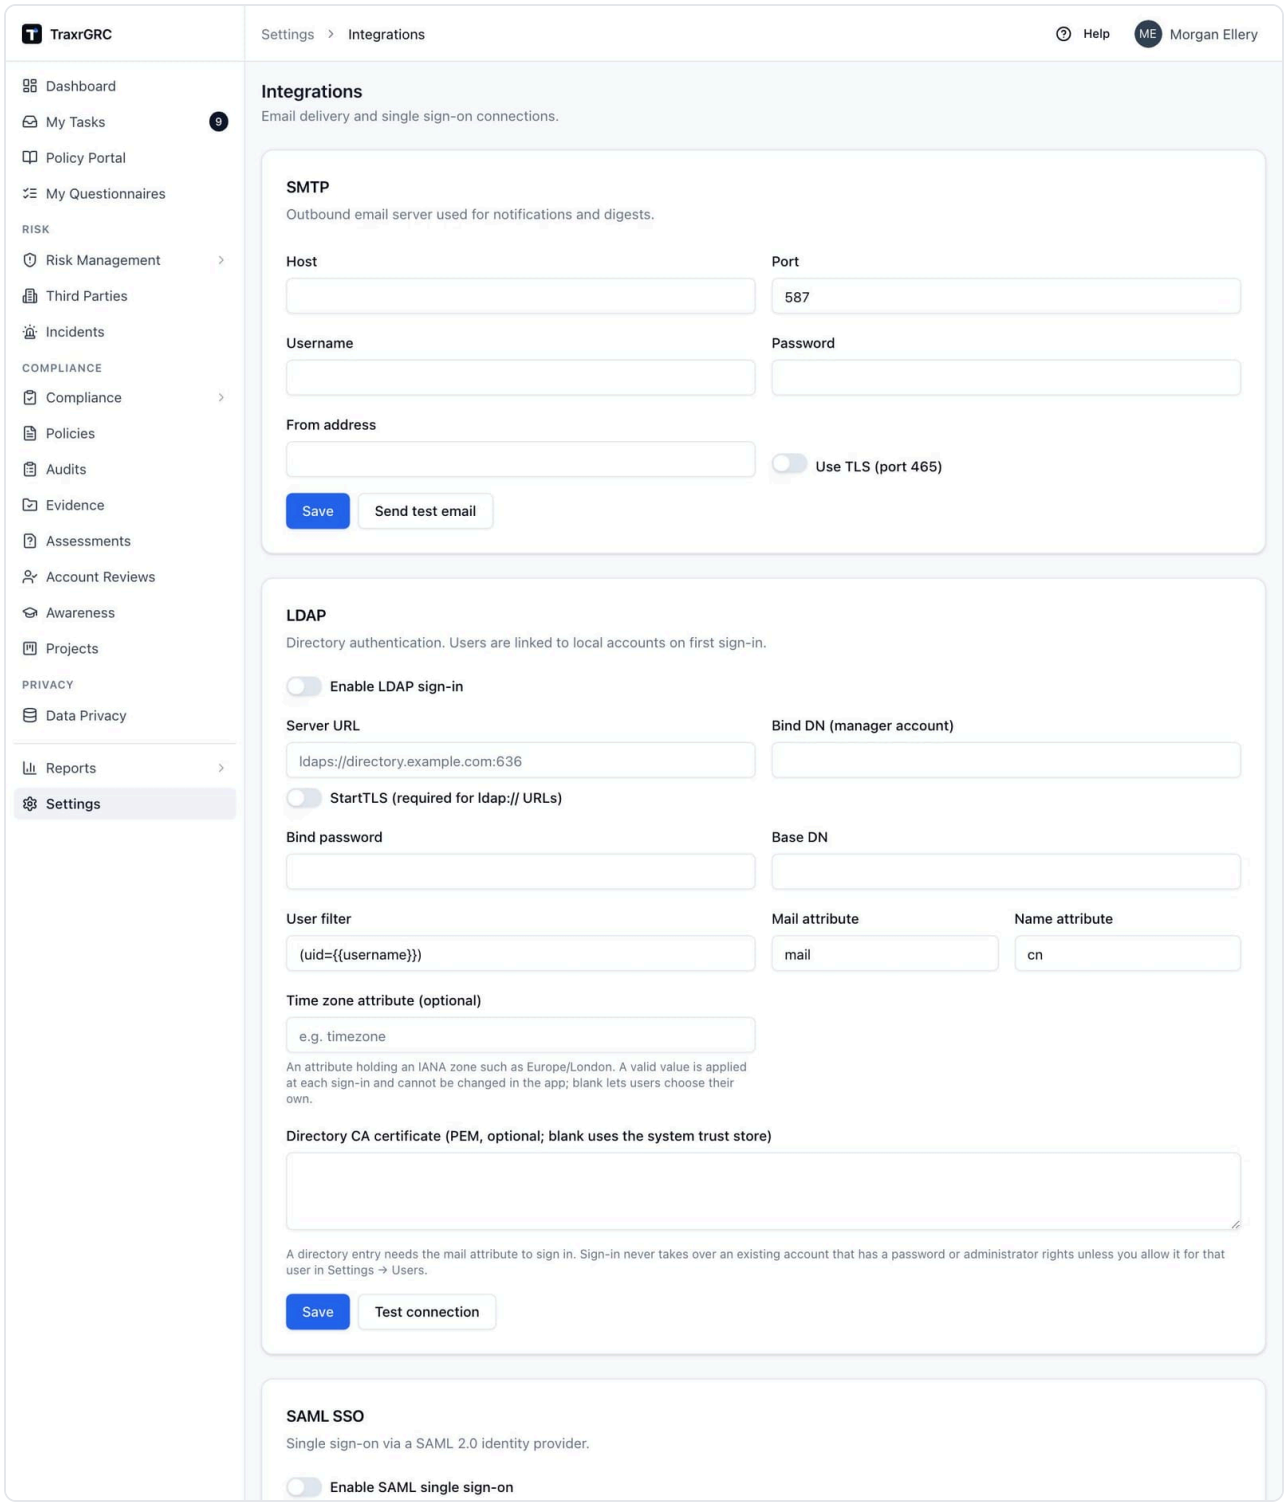


Figure 2.6. The Integrations screen with the SMTP and LDAP cards; the SAML SSO card follows below.

2.7.1 SMTP

The outbound email server used for notifications and digests (see Notifications below). Fields:

Field	What it means
Host	SMTP server hostname.
Port	SMTP server port (587 by default).
Username	Account used to authenticate, if your server requires it.
Password	Password for that account.
From address	The sender address on outgoing TraxrGRC mail.
Use TLS (port 465)	Toggle for an implicit-TLS connection.

With a username set and **Use TLS** off, the server must offer STARTTLS: GRC Platform refuses to send the password over an unencrypted connection, the same rule LDAP follows. A relay without TLS that takes a password fails with an error (Send test email shows it); use port 465 with **Use TLS**, enable STARTTLS on the relay, or use an unauthenticated relay, which sends no password and is unaffected.

To configure and check the mail server:

1. Fill in the fields and click **Save**.
2. Click **Send test email**. TraxrGRC sends a test message, with the product name and "SMTP test" in its subject, to your own address.
3. Read the toast: "Test email sent to your address." on success, or "Test email failed:" with the reason. This is the fastest way to verify a new configuration before notifications depend on it.

2.7.2 LDAP

Directory authentication: with LDAP enabled, users sign in with their directory credentials. A first sign-in creates an account, or links to an existing one with that email when it has no password, administrator rights or break-glass status (or an administrator has allowed the link in Settings → Users).

Field	What it means
Enable LDAP sign-in	Master switch. When on, a failed local password check falls back to the directory.
Server URL	The directory address: `ldaps://directory.example.com:636`, or `ldap://directory.example.com:389` with StartTLS on. Plain `ldap://` without StartTLS is refused, so passwords never cross the network in clear text.
StartTLS (required for ldap:// URLs)	Upgrades an `ldap://` connection to TLS before binding.
Bind DN (manager account)	The service account TraxrGRC binds with.
Bind password	Password for the bind account.
Base DN	Where in the directory user search starts.
User filter	The search filter used to find the signing-in user; must contain `{{username}}`, e.g. `(uid={{username}})`.
Mail attribute	Directory attribute mapped to the TraxrGRC email (defaults to `mail`). An entry without it cannot sign in.
Name attribute	Directory attribute mapped to the display name (defaults to `cn`).
Time zone attribute (optional)	Directory attribute holding the user's IANA time zone, for example `Europe/London`. Blank (the default) reads none and users choose their own. A valid value is applied at every sign-in and cannot be changed in the app; an invalid one is ignored and logged. When the directory stops sending one, the user's last zone stays and becomes theirs to change.
Directory CA certificate (PEM, optional; blank uses the system trust store)	The certificate authority that signed the directory's TLS certificate, when it is not publicly trusted (an internal CA). Changing it clears a stored bind password, as changing the URL does.

To check the connection:

1. Click **Save**.

2. Click **Test connection**. A result panel appears under the buttons: on success, "Connected. Found N directory user(s)." with the directory's user count; on failure, "Connection failed:" with the reason.

Connections give up after 5 seconds and each directory operation after 10.

2.7.3 SAML single sign-on

The **SAML SSO** card sets up single sign-on via a SAML 2.0 identity provider.

Field	What it means
Enable SAML single sign-on	Master switch.
SSO only: turn off password and LDAP sign-in	Everyone signs in through the identity provider, so its MFA covers every account (TraxrGRC's own two-step verification, under Sessions, covers password and LDAP sign-in instead). Password sign-in stays open only to break-glass administrators (Settings → Users); LDAP sign-in is refused. Password sessions already open end within a minute. API tokens are unaffected.
Allow sign-in started at the identity provider	Off by default: TraxrGRC accepts only the identity provider's answer to a sign-in started from its own **Sign in with SSO** button, in the same browser, within 10 minutes. That stops someone from signing a visitor in as themselves with a response they captured. Turn it on only if people launch TraxrGRC from the identity provider's app portal.
IdP sign-in URL (entry point)	Where TraxrGRC sends users to authenticate.
Issuer (SP entity ID)	Your TraxrGRC service's entity ID, e.g. `grc`.
IdP entity ID (recommended: only this issuer may sign in)	The identity provider's own entity ID, e.g. `https://idp.example.com/metadata`. When set, responses from any other issuer are refused.
Name ID format	The identifier format, defaulting to the email-address format.
Time zone attribute (optional)	The assertion attribute holding the user's IANA time zone (its exact name, which may be a URI). Blank (the default) reads none. Treated like the LDAP time zone attribute above.
IdP signing certificate (PEM)	The identity provider's public signing certificate.
SP private key (PEM, optional; signs AuthnRequests)	Your service's key; when present, authentication requests are signed.

Click **Save** to store the configuration. The **Service provider metadata** button opens your service's metadata document in a new tab, and the metadata URL is printed next to it: give that URL (or the file) to your identity provider when setting up the trust. Once SAML is enabled, the sign-in page grows a **Sign in with SSO** button that starts the SSO flow.

TraxrGRC accepts a response only when it is signed by the pinned certificate, addressed to this service (audience, and the bearer confirmation's recipient and the response's destination are this instance's `/api/auth/saml/acs` address), still valid (Conditions NotOnOrAfter is required) and issued less than 10 minutes ago. Each response is accepted once. Single sign-on needs an `https://` APP_URL in practice: the browser only returns the sign-in request cookie on the identity provider's post when it is a secure cookie.

An SSO sign-in creates an account the first time an email is seen. For an email that already has an account, it links only when that account has no password, administrator rights or break-glass status, or is already linked to SSO, or an administrator has allowed the link in Settings → Users. Users you created in Settings → Users have a password, so allow the link for each of them (or let SSO create their accounts instead) before they first sign in with SSO. Break-glass administrators keep signing in with their password.

SSO only can be saved only when SAML is enabled with an entry point and certificate, and at least one active administrator with a password is marked break-glass. While it is on, the last break-glass administrator cannot be demoted, deactivated, deleted or have the flag removed, and SAML cannot be switched off unless SSO only is turned off in the same save. Sign in once through SSO before turning it on, so you know the identity provider works.

2.8 Provisioning (SCIM)

Settings → Provisioning connects an identity provider such as Microsoft Entra ID over SCIM 2.0, so that joiners get an account before they first sign in, names, emails and time zones stay current, leavers are deactivated, and the provider's groups decide who holds the roles you attach to them. It is off by default and changes nothing until an administrator turns it on and creates a token. Built to RFC 7643 and 7644 and Entra's documented SCIM behavior and tested with a scripted client that replays Entra-shaped requests; it has not been tested against a live Entra tenant.

The screen has two cards, **SCIM 2.0 provisioning** and **SCIM tokens**:

Setting	What it means
Accept SCIM provisioning	Master switch, administrators only; click **Save** after changing it. While off, the SCIM address answers 404 and no token is accepted. Turning it on or off is audited and emails every administrator.
Tenant URL	The address to give the identity provider: `https://<your host>/scim/v2`.
SCIM tokens	Bearer tokens for the identity provider, listed with **Name**, **Created**, **Last used**, **Expires** and **Status**. **New token** (administrators only) asks for a **Name** and an optional **Expires** date and shows the token once; only its hash is stored. Tokens last at most the API token lifetime set in Settings → Sessions (365 days by default). **Rotate** revokes a token and shows its replacement at once (provisioning fails until you paste the new one into the identity provider); **Revoke** stops it. Both ask for confirmation. Creating and rotating email every administrator. A SCIM token works only at the SCIM address: it grants nothing on the REST API or to AI assistants, and API tokens do not work at the SCIM address.

What provisioning may do:

- **Accounts are keyed by email**: the work or primary email the provider sends, else the userName (which must then be an email address). A new email creates an account with no password and no administrator rights; the person signs in through SSO.
- **An existing account** with the same email is linked only when it has no password, no administrator rights and no break-glass status, or is already linked to SSO, or an administrator has ticked **Allow the next LDAP or SSO sign-in, or SCIM, to link this account** in Settings → Users (the same rule as SSO sign-in). Otherwise the provider gets a conflict and the refusal is in the audit log. Accounts that were not provisioned are invisible to the provider.
- **Leavers**: setting the user inactive in the provider deactivates the account and ends its sessions within a minute; reactivating it there reactivates the account. Removing the user from the app's assignment (DELETE) also deactivates the account, removes it from the SCIM groups and unlinks it from SCIM. Nothing is ever deleted, since acknowledgements, evidence and the audit trail refer to people.
- **Never**: grant administrator rights, set a password, change or deactivate a break-glass administrator, reactivate an administrator or change an administrator's email, or deactivate the last active administrator. Deactivating any other administrator is allowed and emails every administrator.
- **Groups**: the provider creates, renames and deletes only the groups it created, and manages their members (provisioned users only). A group name already used by a group made here is a conflict. Roles are never set by the provider: attach roles to a SCIM group in Settings → Groups, and its members receive them. Attaching a powerful role to a SCIM group lets the identity provider decide who holds it, so treat that group's membership in the provider as part of your access control.

- **Time zone**: a valid IANA name in the user's `timezone` attribute becomes their zone with source "SCIM provisioning", like a directory zone.
- Every change is audited with the actor "SCIM" and the token that made it. Each token may make 1,200 requests and 600 changes a minute; an address that sends 20 bad tokens in 10 minutes is refused (429) for the rest of that window.

Setting up Microsoft Entra ID (Enterprise applications → your app → Provisioning):

1. Here: turn on **Accept SCIM provisioning**, **Save**, then **New token** and copy it.
2. In Entra: Provisioning mode **Automatic**; Tenant URL the address shown here (`https://<host>/scim/v2`); Secret Token the token. **Test connection**.
3. Mappings, users: keep `userPrincipalName` → `userName` (or `mail` if UPNs are not email addresses), `Switch([IsSoftDeleted], , "False", "True", "True", "False")` → `active`, `displayName` → `displayName`, `mail` → `emails[type eq "work"].value`, `objectId` → `externalId`. Optional: a time zone source attribute → `timezone`. Other default mappings (name parts, job title, department, manager) are accepted and ignored. Use the same email the SSO assertion sends, so SSO sign-in finds the provisioned account.
4. Mappings, groups: `displayName` → `displayName`, `objectId` → `externalId`, `members` → `members`.
5. Scope: "Sync only assigned users and groups", assign the users and groups, and start provisioning. Then attach roles to the new groups in Settings → Groups.

2.9 Sessions

Settings → Sessions (administrators change it; anyone with settings:read can view it) holds the **Sign-in sessions, API tokens and storage** card:

Field	What it means
Session length (hours)	A sign-in ends this long after it started, whatever happens. 1 to 168, default 12.
Idle timeout (minutes)	A session nobody has used for this long ends. 5 to 1440, default 60. The browser tells the server about activity (clicks, typing, moving between pages) at most once a minute.
API token lifetime (days)	The longest expiry a new API token may have; 0 for no maximum. Default 365. Existing tokens keep their expiry.
Attachment storage per user (MB)	The total size of the files one person may upload; 0 for no limit. Default 2048 (2 GB). Lowering it keeps existing files and refuses new uploads past the limit.

Saving writes an audit entry and emails every administrator.

Two-step verification, the next card, has one switch with its own **Save**: **Require two-step verification for password and LDAP sign-in**, off by default and changed only by an administrator. When on, anyone who signs in with a password or through LDAP without two-step verification must set it up straight after signing in, before they can use the app; this includes break-glass administrators. Sessions already open are held at the same step within a minute. Single sign-on is never asked for a code. Set it up on your own account first (see Sign-in security under My profile); the card reminds you when you switch it on. Saving writes an audit entry and, when the value changes, emails every administrator. When someone loses their authenticator device, reset their two-step verification in Settings → Users.

Separation of duties, on the same page, is a separate switch with its own **Save**, off by default and changed only by an administrator. By default the person who requests an exception or a policy approval cannot decide it, and a Data Liaison cannot review their own questionnaire answers. **Relax separation of duties** is meant for teams too small to have a second person decide: when on, the person who requested a risk, compliance or policy exception or an in-app policy approval may decide it, a Data Liaison who also holds dataPrivacy review may review their own answers, and, where DPIA approval needs a second person (see Privacy), a DPIA's author may decide it. They still need the right the decision always needs. Every such decision is recorded as self-approved in the audit log and shows a **Self-approved** badge (**Self-decided** for a rejection) on the record. Turning the switch off stops new self-approvals; past ones keep their badge. Naming yourself as the approver of your own exception stays refused, and the other safeguards (approved content checks, the DPIA lock while awaiting approval, outside policy approvals) are unchanged. Saving writes an audit entry and, when the value changes, emails every administrator.

2.10 Branding

Settings → Branding white-labels the application (settings:read to view, settings:update to change; every save is audit-logged). It has two cards: **Name and color**, with its own **Save**, and **Logo**.

Setting	Effect
Product name	Shown in the sidebar, browser tab titles, the sign-in screen, public assessment pages, PDF exports, the SMTP test email and the API document title. Email templates show it through the `{{app.name}}` variable. Leave it empty to use the deployment default (`APP_NAME`, otherwise `GRC Platform`). Up to 60 characters.
Brand color	A `#rrggbb` color applied to primary buttons, focus rings and links. Button text switches between light and dark to stay readable; a color too light to read as link text on white is darkened, keeping its hue, until it meets WCAG AA contrast. **Use default** restores the standard theme.
Show "Powered by" credit	A small credit naming the platform vendor at the foot of public assessment pages. Off by default.
Logo	**Upload logo** takes a PNG, JPEG, WebP or SVG up to 256 KB, checked by its content rather than its file name. It replaces the shield icon in the sidebar and on the sign-in and assessment pages, and becomes the browser tab icon. **Remove** goes back to the shield.

Changes reach every page within about 30 seconds.

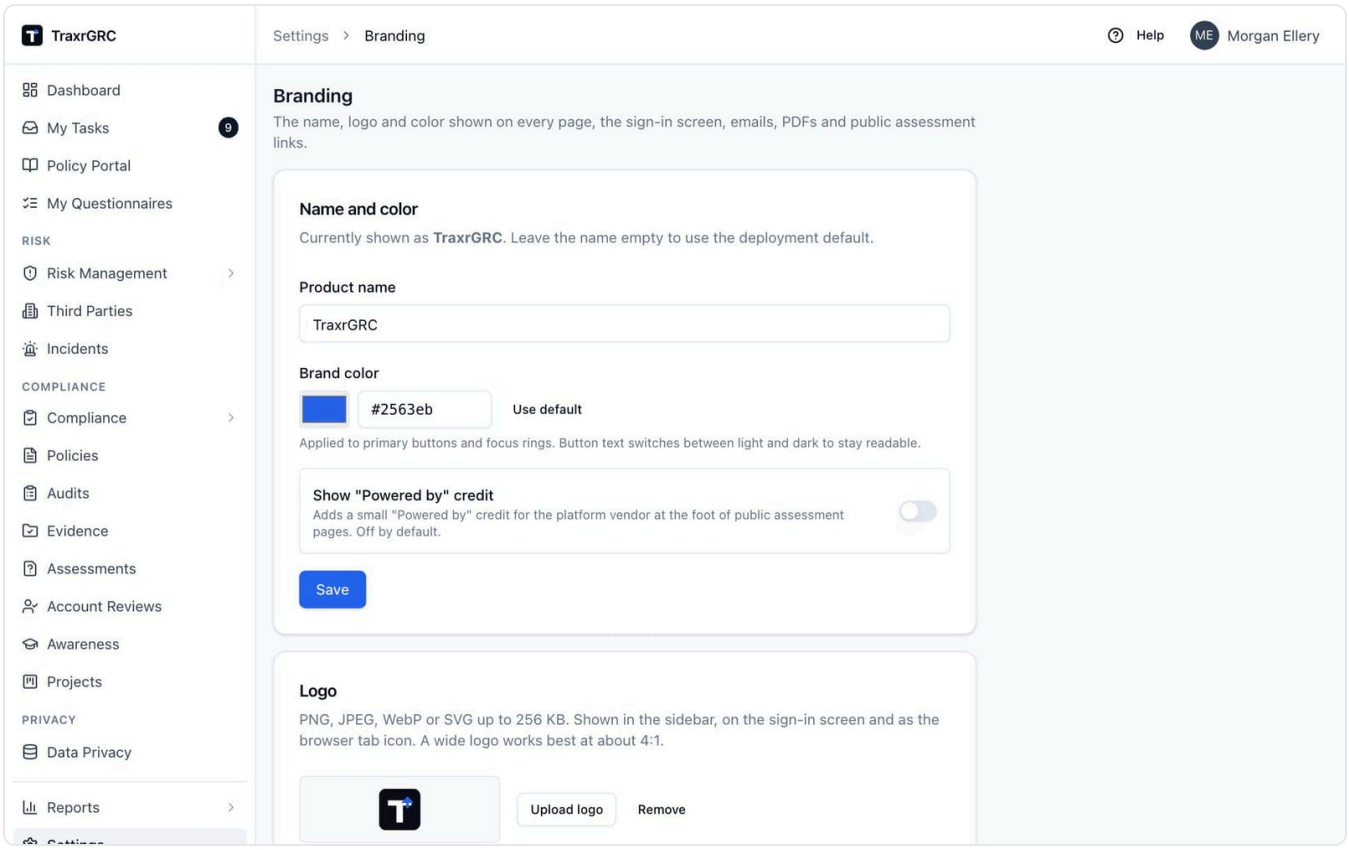


Figure 2.7. The Branding screen with the product name, brand color and logo of the demo instance.

2.11 Policies

Settings → Policies has one card, **Approvals made outside the app**, with one switch, **Allow recording outside approvals** (off by default), and a **Save** button. On, people with update and Review on Policies can record an approval made outside the app and upload reissued documents with one (see Policy Management). Anyone with settings:read can see it; only an administrator can change it, and the change is audit-logged. Turn it on where documents are approved outside the tool, such as by a board or a controlled-document process.

2.12 Privacy

Settings → Privacy holds the **Records of processing header**, printed at the top of the Article 30 register export (settings:read to view, settings:update to change; saves are audit-logged): the **Controller** (legal name), **Controller contact details**, **EU or UK representative** where Article 27 requires one, and the **Data protection officer**'s name and contact details. Empty fields print empty.

DPIA approval, on the same page, has its own **Save** and is changed only by an administrator. It holds three settings:

Setting	Effect
Data Protection Officer	Optional. The one user who gives DPO advice on DPIAs. Only active users can be chosen; deactivating the DPO leaves none named (the card says so) until you choose another. Separate from the DPO contact details printed in the register header above.
DPO advice before a DPIA is decided	**Required when a DPO is named** (the default), **Required** (needs a DPO named) or **Optional**. While required, a submitted DPIA waits for the DPO's advice before anyone can approve or reject it, and approving against the advice needs a reason, which the DPO is emailed. With no DPO named, nothing waits.
Require a second person to approve DPIAs	Off by default, so the person who wrote a DPIA may decide it, as a one-person privacy office needs. When on, whoever created the DPIA or submitted it for approval cannot approve or reject it; another holder of the Approve DPIAs right must. If separation of duties is relaxed (see Sessions), an author may still decide, recorded as self-approved with a **Self-approved** badge on the DPIA.

Saving writes an audit entry for each changed setting and, when any value changes, emails every administrator.

2.13 AI assistants

Settings → AI assistants (the screen is headed **AI assistants (MCP)**) lets AI assistants such as Microsoft Copilot Studio agents, GitHub Copilot in VS Code and Claude connect to TraxrGRC through the Model Context Protocol (MCP) at `https://<your address>/api/mcp`. Anyone with settings:read can see this page; only an administrator can change it, and every change is audit-logged and emailed to all administrators.

This section summarizes what an administrator decides here. The full reference, with the tools an assistant can call, the guardrails, and step-by-step setup for VS Code, Claude Code and Copilot Studio, is the MCP guide (`docs/MCP.md`) shipped with TraxrGRC; keep it at hand when connecting an assistant.

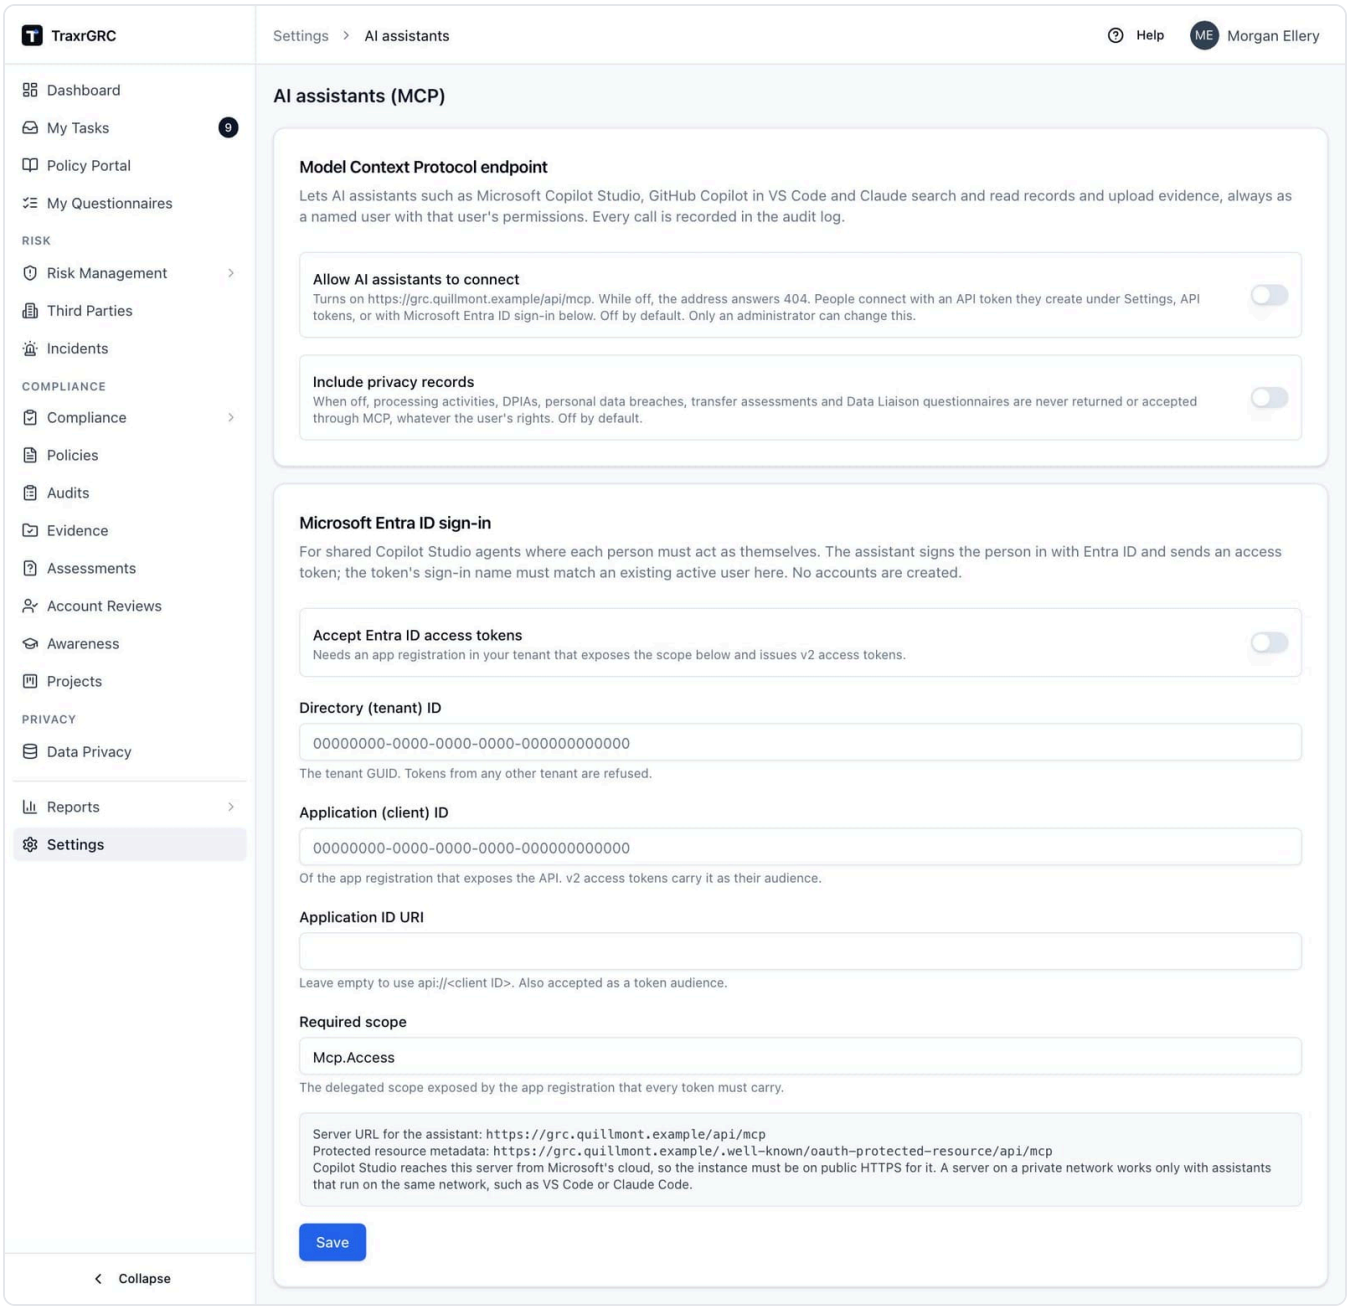


Figure 2.8. The AI assistants (MCP) screen with the endpoint switches and the Microsoft Entra ID sign-in card.

What an assistant can do. Search records across the modules; read one record in full; list risks, controls, compliance requirements and policies with filters; list what is waiting on its user; and attach evidence to a record its user may update. It cannot delete, edit or decide anything.

The **Model Context Protocol endpoint** card holds two switches:

Switch	What it does
Allow AI assistants to connect	Off by default. Turns on the `/api/mcp` address; while off, the address answers 404 ("not found").
Include privacy records	Off by default. While off, processing activities, DPIAs, personal data breaches, transfer assessments and Data Liaison questionnaires are never returned or accepted through MCP, whatever the user may see in the app.

The **Microsoft Entra ID sign-in** card is for shared Copilot Studio agents where each person signs in as themselves:

Field	What it means
Accept Entra ID access tokens	Turns Entra sign-in on. Needs an app registration in your tenant that exposes the scope below and issues v2 access tokens.
Directory (tenant) ID	The tenant GUID. Tokens from any other tenant are refused.
Application (client) ID	Of the app registration that exposes the API.
Application ID URI	Leave empty to use `api://<client ID>`.
Required scope	The delegated scope every token must carry.

Below the fields, a panel shows the server URL to give the assistant, the protected resource metadata address and (once configured) the authorization server. Click **Save** to store both cards. The person signing in must already have an active account here with the same email as their Entra sign-in name; no accounts are created.

An assistant always acts as a named user and sees exactly what that user can see. With an API token (see API Tokens) it is the token's creator, limited to the token's scopes; with Entra sign-in it is the person who signed in. Every tool call is recorded in the Audit Log under the module "mcp". Evidence uploaded by an assistant goes through the same file checks and storage limit as uploads in the app, with a 10 MB limit per file, and appears on the record's Attachments tab. For a control or a compliance requirement it becomes new evidence linked to the record instead (see Files on controls, requirements and account reviews). That needs read access to the record and `evidence:create`, the same as linking evidence in the app; update access to the record is not needed, since the record itself does not change.

Uploading files through an assistant. An assistant cannot carry a real file (a PDF, spreadsheet or screenshot) to TraxrGRC itself. Instead it gives you an upload link:

1. Open the link in your browser, signed in as yourself. If you are signed out, you sign in first and return to the link.
2. Check the record type and name the files will be attached to, and any note from the assistant.
3. Choose up to 10 files (25 MB each, the same file types and storage limit as on a record's Attachments tab) and select **Upload**. They appear on the record's Attachments tab, or for a control or requirement on new evidence linked to it.

A link expires 15 minutes after it is made and works only for the person whose assistant asked for it: anyone else who opens it is refused and the attempt is logged. Your rights are checked again when you open the link and for every file, so a link stops working if you lose access to the record or an administrator turns AI assistants off. The Audit Log records the link being issued and each file uploaded through it (module "mcp").

Copilot Studio calls TraxrGRC from Microsoft's cloud, so it works only when this instance is reachable over public HTTPS. A server on a private network works only with assistants that run on the same network, such as VS Code or Claude Code.

2.14 Notifications

Settings → Notifications (address `/settings/notifications`) controls the email TraxrGRC sends. Opening it needs notifications:read; changing templates and the scheduler needs notifications:update. It has four tabs: **Templates**, **Rules** (shown to administrators only), **Delivery log** and **Scheduler**.

2.14.1 Templates

Every email TraxrGRC can send has a template, keyed by a stable identifier. The **Templates** table lists each template's **Key** (click it to edit), **Name**, **Module** (as a badge), **Subject** and an **Enabled** switch, with the usual search and column filters; the **Module** filter offers every area in the table below. The templates are:

Area	Template keys
Risks	`risk.review.due`, `risk.review.overdue`, `risk.appetite.breach`, `risk.exception.expiring`
Controls	`control.audit.due`, `control.audit.overdue`, `control.audit.failed`, `control.maintenance.due`, `control.maintenance.overdue`
Compliance	`compliance.exception.expiring`
Policies	`policy.review.due`, `policy.review.overdue`, `policy.exception.expiring`, `policy.acknowledgement.requested`, `policy.acknowledgement.due`, `policy.acknowledgement.overdue` (see Policy portal and acknowledgements)
Awareness	`awareness.due`
Data privacy	`dataprivacy.review.due`, `dataprivacy.review.overdue`, `dataprivacy.dpa.expiring`, `dataprivacy.dpa.expired`, `dataprivacy.tia.due`, `dataprivacy.tia.overdue`, `dataprivacy.questionnaire.sent`, `dataprivacy.questionnaire.submitted`, `dataprivacy.dpia.approvedAgainstAdvice`, `dataprivacy.attestation.sent`, `dataprivacy.questionnaire.due`, `dataprivacy.questionnaire.overdue`
Evidence	`evidence.expiring`, `evidence.expired`
Assessments	`assessment.sent`, `assessment.reminder`, `assessment.overdue`, `assessment.returned`, `assessment.completed`
Reports	`report.scheduled` (see Scheduled reports)
Platform	`security.alert` (see Security alerts under Audit Log), `user.password.reset`, `user.password.changed` (a user changed their own password), `user.mfa.reset` (an administrator reset a user's two-step verification), `notification.rule` (every email a notification rule sends) and the summary `digest.daily`

To tailor a template:

1. Click its key. The **Edit template** dialog opens.
2. Edit the **Name**, **Subject** and **Body**. Both subject and body accept `{{dotted.vars}}` placeholders, e.g. `{{risk.title}}`.
3. Check the **Available vars** panel on the right of the dialog. It lists, for that template, every placeholder with its meaning. For example, `risk.review.due` supports `{{risk.title}}` and `{{review.dueDate}}`; `risk.appetite.breach` adds `{{risk.score}}` and `{{risk.band}}`; `digest.daily` supports `{{digest.date}}` and `{{digest.body}}`. Every template also lists `{{app.name}}`. The Administrator's Guide has the full table.
4. Click **Save template**. The change applies to the next email queued through that template.

Use the **Enabled** switch to turn a template off without editing it; disabled templates send nothing. Template edits and toggles are audited.

2.14.2 Notification rules

The fixed notifications above go to record owners on fixed events. Notification rules let an administrator email other people about other events. The **Rules** tab is shown to administrators only, and every rule created, changed or deleted is recorded in the Audit Log. A new installation, or one upgraded from an earlier version, has no rules ("No rules yet, so nothing beyond the fixed notifications is sent."), so nothing beyond the fixed notifications is sent until an administrator creates one.

The rules list shows each rule's **Name** (with a "Disabled" badge when paused), **Trigger**, **Records** (the record types, "All", and any status condition), **Recipients**, and **Opted out** (how many people switched the rule off on their profile), with edit and delete buttons.

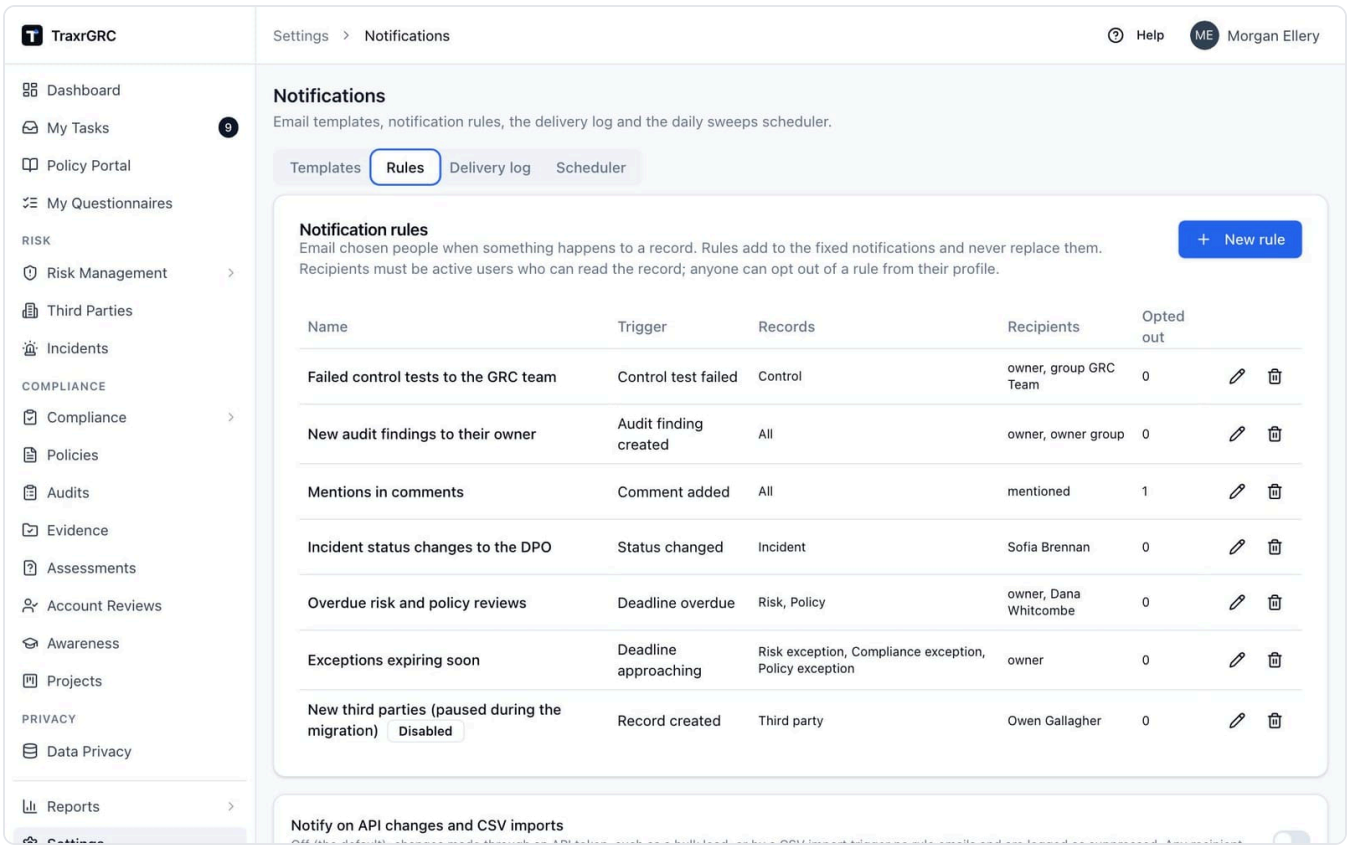


Figure 2.9. The Rules tab listing the demo's notification rules with their triggers, records, recipients and opt-out counts.

To create a rule:

1. Click **New rule**. The **New notification rule** dialog opens.
2. Fill in the fields:

Field	What it means
Name	Shown in the delivery log, in the email ("You receive this because of the notification rule ...") and on each person's My profile page.
Trigger	What the rule reacts to: **Comment added**, **Record created**, **Status changed**, **Audit finding created**, **Control test failed**, **Deadline approaching** or **Deadline overdue**.
Records	The kinds of record the rule covers: Risk, Control, Compliance requirement, Policy, Risk exception, Compliance exception, Policy exception, Incident, Project, Third party and Audit finding, as far as the trigger applies to them. Tick none to cover every kind the trigger supports.
Only when the record's status is	Optional. For a status change, the new status; for other triggers, the record's status at the time. Tick none for any status.
Recipients	Any mix of the record's owner, the members of the record's owner group, people @mentioned in the comment (Comment added only), named users, the members of chosen groups, and the members of every group that carries a chosen role. Only active users of this installation can be chosen; no outside address can be entered.
Enabled	Untick to pause the rule without deleting it.

3. Click **Create rule** (**Save rule** when editing).

To change a rule, click its pencil icon (the dialog is headed **Edit notification rule**). To remove one, click its trash icon and confirm in the **Delete this rule?** dialog.

Deadlines are the ones the daily sweep already tracks: risk and policy reviews, policy acknowledgement requests with people still to acknowledge (on the policy), control audits and maintenance, project and audit finding deadlines (approaching or overdue), and approved exceptions coming up to expiry (approaching only). The lead times are the scheduler's **Review lead days** and **Exception lead days**. A deadline rule emails once when a deadline enters the lead window and once when it passes; moving the date starts afresh. Status and comment rules work within seconds of the change; deadline rules run with the daily sweep (and **Run sweeps now**).

What every rule email contains, and who gets one:

- The email names the record (its title or reference), what happened (for example "Status: Open to Closed" or "Risk review was due 2026-10-09") and a link to it. It never contains a comment's text, a description or any other field. An incident recorded as a personal data breach is named only by a short reference ("Incident 1a2b3c4d (title withheld)"), since breach titles can name the people affected.

- A recipient gets the email only if they may open the record at that moment, by the same rule the app uses on screen (for an exception, its risk, requirement or policy; for an audit finding, its audit plan and the finding). This is checked when the rule fires and again just before sending, so someone who lost access in between gets nothing.
- Nobody is emailed about their own change.
- Each person gets at most 20 rule emails an hour; the rest are suppressed.
- Changes made through an API token, such as a bulk load with the loader script, and changes made by a CSV import send no rule emails by default; the delivery log lists them as suppressed. Turn on **Notify on API changes and CSV imports** below the rules list to include them (the hourly cap still applies).
- Anyone can switch a rule off for themselves on My profile.

Fixed notifications and security alerts are unaffected by rules: they keep going to the same people as before.

2.14.3 Scheduler

The **Scheduler** tab holds the **Daily sweeps scheduler** card, which configures when and what the background worker emails. The worker scans reviews and exceptions on a cron schedule and queues notification emails. Sweeps skip items already notified in the last 20 hours, so overlapping runs never double-send. A reminder goes out only while its recipient can still read the record it names (an owner whose role no longer includes that module gets nothing), and awareness and questionnaire reminders only to people who still hold Awareness or Data Liaison read.

Setting	What it means
Daily digest	When on, each recipient gets one summary email (the `digest.daily` template) instead of individual notifications.
Cron schedule (UTC)	Five fields (minute hour day month weekday) in UTC. Default `0 7 * * *` (07:00 daily). An invalid expression is rejected when you save.
Review lead days	Reviews due within this many days (0 to 365; default 7) are notified: risk, policy and data-privacy reviews due or overdue, plus awareness programs due.
Exception lead days	Approved risk, compliance and policy exceptions expiring within this many days (0 to 365; default 14) are flagged to their owners.

Click **Save settings** to store the schedule. A changed schedule is active at once: the confirmation reads "Scheduler settings saved. The new schedule is active." If instead a warning says the new schedule takes effect when the background worker next starts, the worker could not be reached; ask whoever runs the server to start or restart it. **Run sweeps now** performs an immediate sweep outside the cron schedule and reports how many items it found, across how many recipients, and how many digests were

queued; use it after changing settings or before relying on the schedule. The run is recorded in the Audit Log.

2.14.4 Delivery log

The **Delivery log** tab lists every queued email with columns **Created**, **Template**, **To**, **Rule or schedule** (the notification rule or scheduled report that sent it, with the report attachment's name; "Deleted rule" or "Deleted schedule" once that is gone), **Subject**, **Status**, **Error** and **Sent at**. Statuses:

Status	Badge	Meaning
Queued	Gray	Waiting to be sent.
Sending	Gray	A rule email being sent. One left in this state by a crash is marked failed after 15 minutes and is not resent, so nobody gets it twice.
Sent	Solid	Delivered successfully; **Sent at** shows when.
Failed	Red	Sending failed; the **Error** column holds the reason. A rule email is tried four times, a minute or more apart, before it fails.
Skipped	Outline	Intentionally not sent.
Suppressed	Outline	A rule email deliberately not sent to this person. The **Error** column says why: `no-access` (they may not open the record), `opted-out`, `api-write` (the change came through an API token), `rate-limit` (over 20 rule emails in an hour), `recipient-inactive`, `record-deleted`, `rule-disabled` or `template-disabled`. A suppressed row has no subject or body.

You can filter the log by status and search across recipient, subject, template key and error text. Rule emails, with the **Rule or schedule** column naming the rule, are listed to administrators only. The demo instance sends no email, so its delivery log is empty.

2.15 Audit Log

Settings → Audit Log (address `/settings/audit-log`; needs auditLog:read, and its card appears on the Settings page for anyone who holds it) is the organization's trail: every mutating action across TraxrGRC, with before/after diffs where applicable. Each create, update and delete you performed from the screens above lands here automatically.

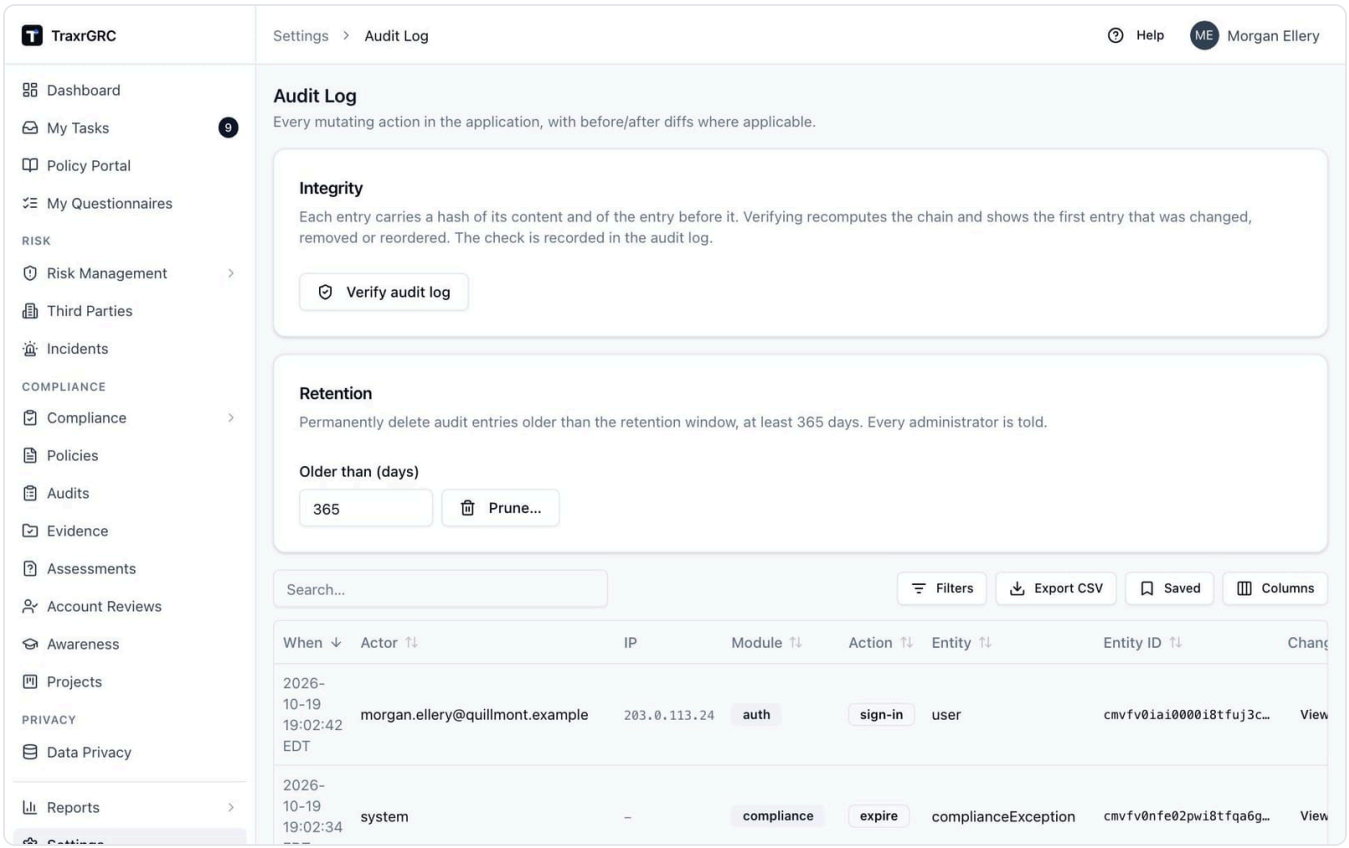


Figure 2.10. The Audit Log with the Integrity and Retention cards above the table of entries.

The table shows:

Column	What it shows
When	The time of the action, to the second.
Actor	Who acted, and the API token when the change came through the REST API.
IP	The client address the reverse proxy reported.
Module	The module, as a badge.
Action	The action, as a badge, e.g. create, update, delete, export, restore, run.
Entity	The record type.
Entity ID	The record's identifier.
Changes	A **View diff** button when the entry recorded a change.

The log opens newest first; click a column heading to sort another way. A free-text search matches the actor, entity type, entity ID and action. Column filters narrow it further: dropdowns for **Module**, **Action** and **Entity** (the options are drawn from the entries themselves), text for **Actor**, **IP** and **Entity ID** (any

part of the value matches), and a date picker for **When** that shows the entries of that calendar day in your time zone.

Export CSV downloads the entries matching the current search and filters, up to 5,000 rows, with times in UTC. Its **Changes** column holds each entry's recorded change (before, after and, where recorded, the source) as compact JSON, and is empty for entries without one. To export a long period, export it a day at a time with the **When** filter.

Click **View diff** to open **Change details**: the Before and After values side by side, exactly as recorded.

Besides changes to records, the log holds:

- **Sign-ins** (module `auth`): every successful sign-in (`sign_in`) and sign-out, with the method (password, LDAP or SSO) and whether it was a break-glass administrator signing in with a password while SSO-only is on. Every refused sign-in (`sign_in.failed`) records the email as typed, the address and the reason: unknown user, wrong password, no password set, inactive account, locked out (and which lock), refused by SSO-only, method switched off, rejected by the directory, a directory entry with no email, a refused account link (LDAP or SSO asserting the email of an account with a password or administrator rights), an invalid or replayed SSO response (with the precise reason, such as not a reply to this browser's request), or malformed input. Nothing is created for an email that has no account, and the sign-in page answers every failure the same way. A sign-in lock (`lockout`) is recorded once per locked email-and-address pair or address, and `account.guarded` when one account reaches ten failures in 15 minutes. Accounts created by a first LDAP or SSO sign-in appear as a `users` create entry, a first link to an existing account as `users` `sso.link`, and "sign out everywhere" as `users` `sessions.revoke`.
- **Two-step verification and password changes**: a passed or failed code at sign-in (`auth` `mfa.verified`, `mfa.failed`, and `mfa.locked` when an account's code step locks after 10 wrong codes; a `sign_in` entry shows `mfaPending` when the session still waited for the code), setting it up (`users` `mfa.enroll.started`, `mfa.enabled`, and `auth` `mfa.enroll.failed` for a wrong code), an administrator's reset (`users` `mfa.reset`), and a user changing their own password (`users` `password.change`, or `password.change.failed` for a wrong current password).
- **Refused requests** (module `security`, action `permission.denied`): a screen action, API call or API token refused for lack of rights, with the module and right, or the reason the token was refused (including `rate_limited` when a token goes over its request budget). At most ten a minute are kept per person (or per address for requests with no valid token); the next entry says how many were dropped.
- **Integration settings**: saving SMTP, LDAP or SAML records every field before and after. Passwords and the SAML private key are never recorded, only whether they changed; the SAML identity provider certificate is recorded by its SHA-256 fingerprint, so a swapped certificate stands out.

Entries cannot be edited or deleted, by the application or by anything else connected to its database, except by the Retention card below.

2.15.1 Integrity

The **Integrity** card (administrators only) proves the log has not been tampered with. Each entry carries a hash of its content and of the entry before it. Click **Verify audit log** to recompute the chain: the result reads "Intact: N entries verified (chain positions ...)", or "Broken at entry ..." with the reason and the number of intact entries before the first entry that was changed, removed or reordered. The check itself is recorded in the audit log.

2.15.2 Retention

The **Retention** card (administrators only) prunes history:

1. Enter an age in **Older than (days)**: 365 to 3650; defaults to 365.
2. Click **Prune...**.
3. Confirm in the **Prune audit entries?** dialog by clicking **Prune**.

> **Warning:** All entries older than that are permanently deleted, and this cannot be undone. Export the audit log first if you must keep them.

The pruning itself is written to the log, and every administrator receives a security alert. Entries less than a year old can never be pruned.

2.15.3 Security alerts

Every active administrator receives an email (template `security.alert`, editable under Settings → Notifications) when:

- a sign-in is locked after repeated failures, or one account has ten failures in 15 minutes (at most once an hour per lock, and at most ten an hour in all; a break-glass account is named as such);
- the session or API token limits are changed;
- administrator rights are granted or removed (including deleting an administrator);
- a break-glass administrator signs in with a password while SSO-only is on, or SSO-only is turned off;
- the audit log is pruned, or the database is restored from a backup;
- SMTP, LDAP or SAML settings are saved;
- an API token is created;
- two-step verification is required or no longer required, a user's two-step verification is reset, or an account's code step is locked after 10 wrong codes in 15 minutes.

Alerts need SMTP to be configured; each one is listed in the Settings → Notifications delivery log.

2.16 API Tokens

Settings → API Tokens (address `/settings/api-tokens`) issues tokens for scripts and integrations that call TraxrGRC over the REST API (see REST API). Tokens authenticate requests with an `Authorization: Bearer aat_…` header; TraxrGRC stores only the SHA-256 hash of each token, never the token itself. Viewing the list needs apiTokens:read; creating tokens needs apiTokens:create.

The table shows each token's **Name**, **Scopes** (a compact summary such as "risks: read, update"), **Created by**, **Created**, **Last used**, **Expires** and **Status**:

Status	Badge	Meaning
Active	Solid	The token works.
Expired	Red	Its expiry date has passed.
Revoked	Outline	Revoked by an administrator.

To create a token:

1. Click **Create token**. The **Create API token** dialog opens.
2. Enter a **Name** describing its consumer, e.g. "CI compliance exporter", and optionally an **Expiry date**. Tokens last at most 365 days unless an administrator changed the limit under Settings → Sessions: a blank expiry gets the maximum, and a later date is refused. Tokens created before the limit keep their expiry.
3. Grant scopes in the matrix: tick the rights (Read, Create, Update, Delete, Review) per module, the same vocabulary as role permissions. Review is offered on the same four modules as in roles (Risks, Compliance, Policies and Data privacy) and shows a dash elsewhere. At least one scope is required. The scopes bound everything a request does, beyond the endpoint it calls: creating a risk from a DPIA also needs Create on Risks, linking a vendor or system to a record needs Read on Third parties or Assets, and a token without Read on Controls, Policies or Risks sees a compliance requirement's mappings in those modules as hidden. Any right on a module lets a request read that module's records along the way; listing them still needs Read. A token never exceeds its creator's own rights.
4. Click **Create token**.

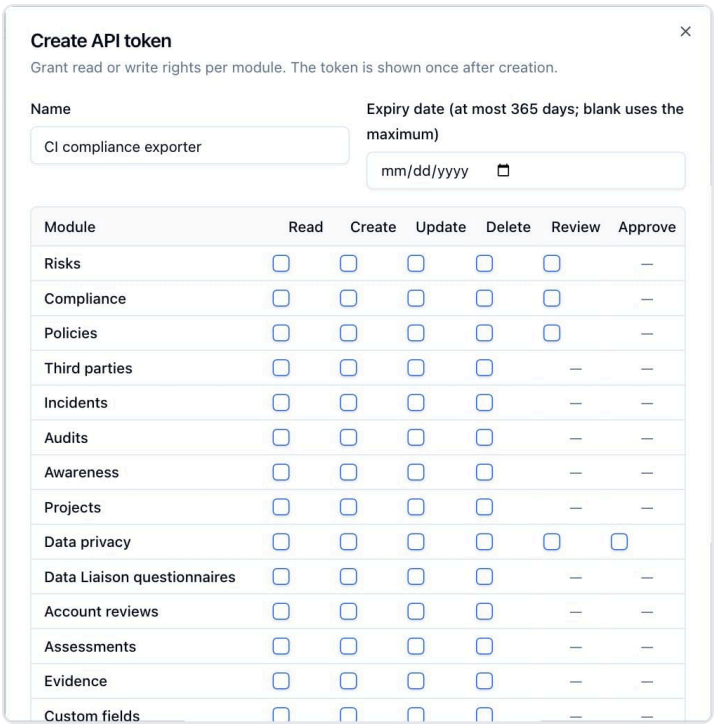


Figure 2.11. The Create API token dialog with the name, expiry date and scope matrix.

A **Token created** dialog appears with the full token value. This is the only time it is ever shown: copy it immediately (the copy button puts it on the clipboard) and store it somewhere safe, then click **Done**. If you lose it, revoke the token and create a new one.

The same tokens connect AI assistants to the MCP endpoint (see AI assistants above): give such a token only the modules the assistant needs, Read alone for a read-only assistant.

To retire a token, click its trash icon and confirm in the **Revoke token?** dialog. Requests using that token fail immediately, and the revocation cannot be undone.

2.17 Backups

Settings → Backups (address `/settings/backups`; administrators only) exports and restores full database snapshots and manages the automatic weekly backups. The screen states the stakes: exporting downloads a complete, signed snapshot of every table, restoring a backup replaces the entire database except the audit log and revokes every API token and SCIM token, and exporting, restoring and running backups are administrator operations.

Include files in backups, on by default, decides whether backups carry the attached files of every record (evidence, policies, risks, vendors and the rest). The screen shows how many files there are and how much they add to each backup. With it on, a backup is a gzipped tar (`.tar.gz`) of the signed snapshot and every file, and the snapshot lists each file with its SHA-256. There is no size limit beyond the server's free disk space: backups are assembled and unpacked on disk, and the screen says **Not enough free disk space** before starting one the disk cannot hold. With it off, a backup is the signed

JSON snapshot alone: it holds attachment records but not the files, so back up the upload directory separately. Changing the setting is recorded in the audit log.

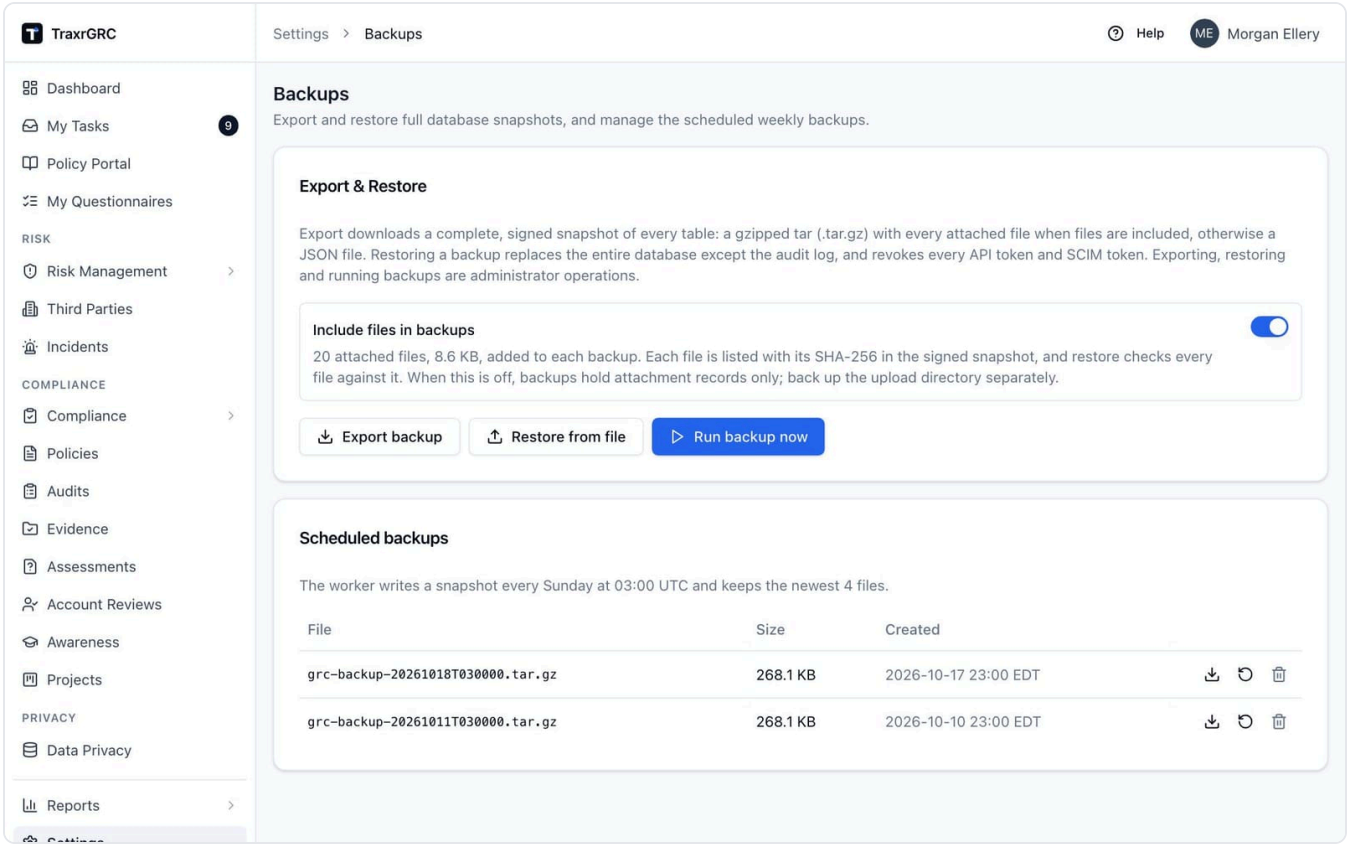


Figure 2.12. The Backups screen with the Export & Restore card, the Include files in backups switch and two scheduled backups, each with Download, Restore and Delete buttons.

The **Export & Restore** card has three buttons:

1. **Export backup** downloads a backup of the whole database to your computer: `grc-backup-YYYY-MM-DD.tar.gz` with files, or `grc-backup-YYYY-MM-DD.json` without. Keep it somewhere safe: it contains everything.

2. **Restore from file** opens a file picker; choose a previously exported backup (`.tar.gz`, `.json`, or a `.zip` made by an earlier version). To restore a scheduled backup that is already on the server, use the restore action on its row in **Scheduled backups** instead; it opens the same dialog. The **Restore backup?** dialog warns that restoring replaces the entire database with the file's contents (every current record is wiped first, including sign-in sessions, and every API token and SCIM token is revoked) and asks you to **Type RESTORE to confirm**. The **Restore** button stays disabled until the text matches exactly. When the restore finishes, a notice tells you it completed, that all sessions were wiped and that every API token and SCIM token was revoked: **sign back in** afterward.

   Every snapshot carries a signature made with a key only the server holds, and restore checks it. A file that was edited after export, or was exported by another installation, is refused. A file with no

signature at all (exported before snapshots were signed) is not refused outright: the dialog warns that it cannot be shown to be unmodified and offers **I trust this unsigned file**; tick it only if you know where the file came from, then restore again. The audit log records whether each restore was signed or trusted.

Snapshots hold the SMTP password, LDAP bind password and SAML private key encrypted with this instance's key; restored onto an instance with a different key (see the deployment guide), those three must be re-entered under Settings → Integrations. A restore never removes or rewrites the audit log: the current entries stay, and a signed snapshot's entries that are not already present are added back. It also keeps the current single sign-on replay records and sign-in lockouts, and it **revokes every API token and every SCIM token**, so integrations need new API tokens and your identity provider needs a new SCIM token afterward. A backup with files also writes its files back, each checked against the SHA-256 in the signed snapshot: a file that matches is restored, one that does not is left out and named in the completion notice, and so is a listed file the backup lacks. The files are moved into place only after the database restore succeeds. A backup with files must carry this installation's signature; an unsigned one is refused. A JSON snapshot brings back only the attachment records, and the files come from the upload directory as it stands, so a file missing there can no longer be downloaded. Snapshots taken before 14 October 2026 hold no data-privacy records (DPIAs, transfers, transfer impact assessments, sub-processors, breach notifications, data mapping questionnaires); restoring one empties those tables.

3. **Run backup now** takes an immediate on-demand snapshot (the same kind the weekly job writes) and adds it to the scheduled-backups list below.

The **Scheduled backups** card: the background worker writes a snapshot every Sunday at 03:00 UTC and keeps the newest four files. The table lists each **File** (its name), **Size** and **Created** date, with a download, a restore and a delete action per row. Downloading streams the file; restoring opens the **Restore backup?** dialog for that file; deleting asks for confirmation in the **Delete backup file?** dialog and removes the file from the list and from disk, which cannot be undone. The demo instance has no backups until one is run.

A restore rewinds the whole platform to the moment of the snapshot, well beyond moving data, so coordinate with your team before restoring over live work.

CHAPTER 3

Risk Management

Risk Management is where you record what could go wrong, score it, decide what to do about it, and prove that you review it on a schedule. Risk owners use it to maintain their risks and treatments, control owners to run control tests, reviewers to sign off scheduled reviews and decide exceptions, and administrators to configure the scoring model.

A risk in TraxrGRC carries an inherent score (before any action), a residual score (after implemented mitigations) and a status. It can hold treatments, scheduled reviews, time-boxed exceptions, comments, attachments and custom fields. This chapter covers the risk register, the scoring model, the shared catalogs (controls, assets and business units) and the Risk Settings screen. Access is controlled by the `risks` module permissions and, for the catalogs, by their own modules (see Who can do what, at the end of this chapter).

The **Risk Management** entry in the sidebar expands to five pages:

Page	What it holds
Risks	The risk register.
Controls	The control library and control testing.
Assets	The things you protect, including systems that hold personal data.
Business Units	Your organizational hierarchy.
Risk Settings	Classifications, magnitudes, impact axes, the calculation matrix and appetite bands.

3.1 The risk register and its tabs

Open **Risk Management → Risks**. The register opens on the **All** tab, which lists every risk you can see. Alongside it is one tab per *risk classification*. A new installation has four:

Classification	Purpose (seeded description)
Asset-based	Risks tied to specific assets.
Third-party	Risks introduced by vendors and partners.
Business continuity	Risks to business operations and recovery.
Enterprise	Organization-wide strategic risks.

Classifications only group the register: they drive the tabs and the **Classification** filter. Your administrator can rename them, add more, or remove unused ones on the Risk Settings screen.

Two buttons sit above the tabs:

- **New risk** creates a risk (requires risks:create; see Creating a risk).
- **Import** creates or updates risks, or creates risk exceptions, in bulk from a CSV file. A menu offers the record type. See Importing from CSV in Getting Started.

3.2 The Risks list

The list shows one row per risk, with live scores computed from the current matrix and treatments:

Column	What it shows
Title	Click to open the risk.
Inherent	Score chip: the inherent score, tinted with its appetite band color. Hover for the band name.
Residual	Score chip: the residual score, tinted with its band color.
Status	Badge: open, closed or archived (see The risk detail page).
Controls	A red **Control failing** or **N controls failing** badge when the risk is flagged (see below); otherwise the worst testing status of its treatment controls (OK, Untested, and so on), or a dash when no treatment links a control.
Classification	The risk's classification.
Owner	The named owner, if set.
Asset	The linked asset, if set.
Next review	The next review due date. For an **open** risk whose review date has passed, the date shows in red with "(overdue)". Closed and archived risks are never flagged overdue.
Labels	Any labels assigned to the risk.

The search box matches the title and description. Column filters narrow the list by **Status**, **Controls** (choose **Controls failing**), **Classification** and **Labels** (records with a label whose name contains the text). Click a column header to sort. **Export CSV** downloads the current view as a CSV file. Saved filters and column choices work as described under Data tables in Getting Started.

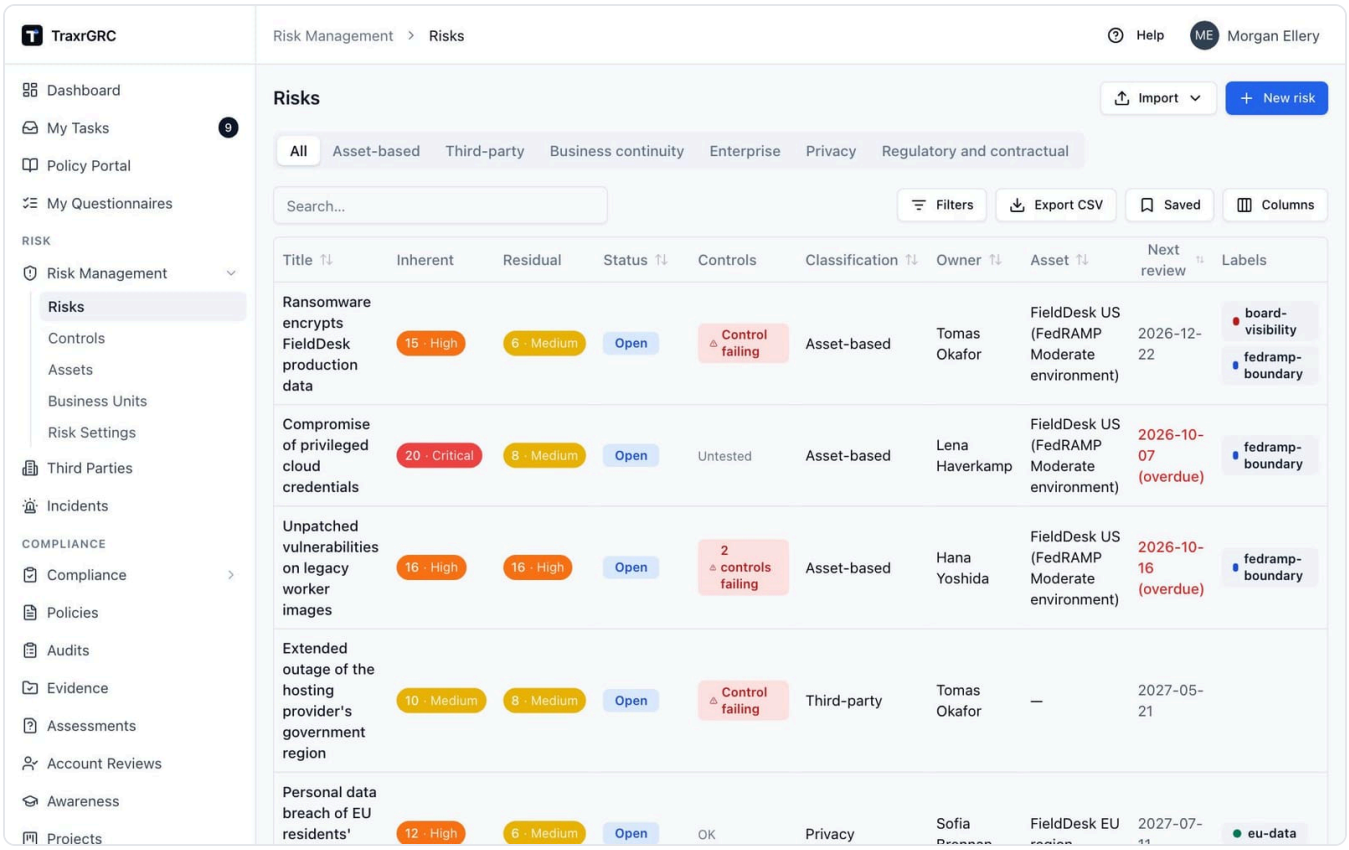


Figure 3.1. The risk register on the All tab, with inherent and residual score chips, control health and overdue review dates.

An **open** risk is flagged as having failing controls when a control linked from one of its treatments failed its latest audit or maintenance, or is overdue for one (see Controls). Rejected treatments and retired controls do not count, and closed or archived risks are never flagged.

A treatment whose control is **retired** earns no reduction: a retired control no longer treats anything, so retiring a control rescores every risk it treats, and the treatment's reduction badge reads "no credit: control retired". Making the control active again restores the credit. The flag is calculated every time it is shown and never changes the risk's status, scores or treatments; it tells you which risks to re-check, the way **At risk** does for compliance.

3.3 Creating a risk

1. Click **New risk** on the Risks page.
2. Complete the **New risk** dialog:

Field	What it means
Title	The risk name (required).
Description	Free text explaining the risk.
Classification	Which register tab the risk belongs to. Pre-filled with the classification of the tab you were viewing.
Asset	Optional link to a catalog asset (see Assets).
Owner	Optional named user responsible for the risk.
Owner group	Optional group responsible for the risk. You can set a user, a group, or both.
Inherent likelihood	A value from your likelihood scale, shown as "Name (value)", for example *Rare (1)* to *Almost Certain (5)*.
Inherent impact	A value from your impact scale, for example *Insignificant (1)* to *Severe (5)*.
Review frequency (months)	How often this risk should be reviewed: 1 to 120, default 12.
Schedule first review	Shown only when creating, ticked by default. The label shows the due date offset, for example "Schedule first review (due in 12 months)". When ticked, a first review is scheduled at the frequency from today and **Next review** is set to it.

3. Click **Create risk**.

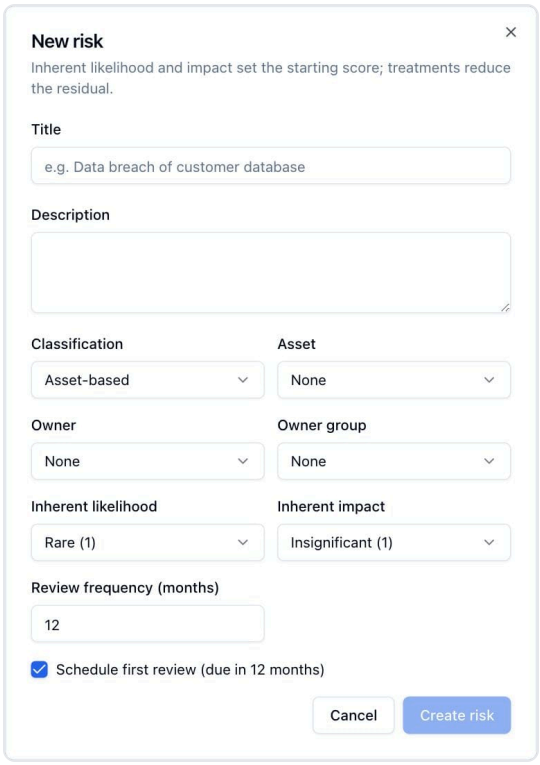


Figure 3.2. The New risk dialog with classification, owner, inherent likelihood and impact, and review frequency.

Creating a risk writes an entry to the audit log. The new risk's residual magnitudes start equal to its inherent ones, so residual equals inherent until mitigations are implemented.

A risk also has a **Business unit**, shown on its detail page. The dialog does not set it; it is set through a CSV import (the `business_unit` column) or the REST API.

3.4 How scoring works

- **Inherent score.** The value of the calculation-matrix cell at your chosen likelihood and impact. With the default 5 × 5 matrix the cell score is likelihood × impact (1 to 25).
- **Impact axes.** When Risk Settings defines impact axes (for example participant safety, data integrity, contractual and privacy), a risk scores each axis on the impact scale and its impact is the **highest** axis, never an average, so a severe consequence on one axis is not offset by a negligible one on another. With no axes, impact is one value.
- **Residual score.** Start from the inherent likelihood and impact, then subtract each *implemented* **mitigate** treatment's "Reduce likelihood by" and "Reduce impact by" values, scaled by the treatment's **effectiveness**: E1 Effective takes the full reduction, E2 Partially effective half (rounded down, so residual is never understated), E3 Largely ineffective and E4 Absent nothing. A treatment not yet assessed counts in full, as before effectiveness existed. Subtractions add up, and each magnitude stops at the lowest value on its scale, so it can never fall below 1 (or your

minimum). The resulting likelihood and impact pair is looked up in the matrix to give the residual score. Treatments with any other strategy, or any status other than *implemented*, do not affect the score.

- **Appetite bands.** Each band covers scores from its minimum score upward; the highest band whose minimum the score reaches wins (ties on minimum resolve by sort order). The band supplies the color and name shown on the score chips. A score with no matching band shows as plain gray "Unscored". The default bands are:

Band	Minimum score	Over appetite
Low	1	No
Medium	6	No
High	12	Yes
Critical	19	Yes

- **Over appetite.** Bands flagged *over appetite* mark unacceptable scores. An open risk whose residual score lands in an over-appetite band shows a red banner on its detail page ("This risk is over the organization's risk appetite (residual score 16, High band)."), and appears in the **Risks over appetite** dashboard widget. When the score *crosses into* an over-appetite band, the notification worker sends the risk owner a "Risk over appetite" notification. A recalculation that leaves the risk where it was sends nothing.

Scores recalculate automatically whenever something relevant changes: editing inherent likelihood or impact, adding, editing or deleting treatments, saving the matrix or bands, or changing magnitude values. The **Recalculate all risks** button on Risk Settings recomputes everything on demand.

3.5 The risk detail page

Click a risk title to open it. The header shows the title, a status badge, a red **Control failing** or **N controls failing** badge when the risk is flagged, and the classification. Under the header you may see:

- the over-appetite banner, for an open risk whose residual score is in an over-appetite band;
- a red panel listing each failing control with a link to it and the reason, such as "Audit failed 2026-09-02" or "Maintenance overdue since 2026-08-01". Controls you cannot read are not named; the panel says only how many there are.

Two cards show the **Inherent score** and **Residual score** as large numbers, each with its band chip and the magnitude names behind it (for example "Likely × Major"). Below them sit the description and the label picker (labels are managed in Administration & Settings). A row of summary cards shows the **Owner** (and owner group), the **Asset** (and its business unit), the risk's **Business unit**, and **Next review**.

The next review date turns red and is marked "overdue" when an open risk has passed it, and the review cadence appears underneath ("every N months").

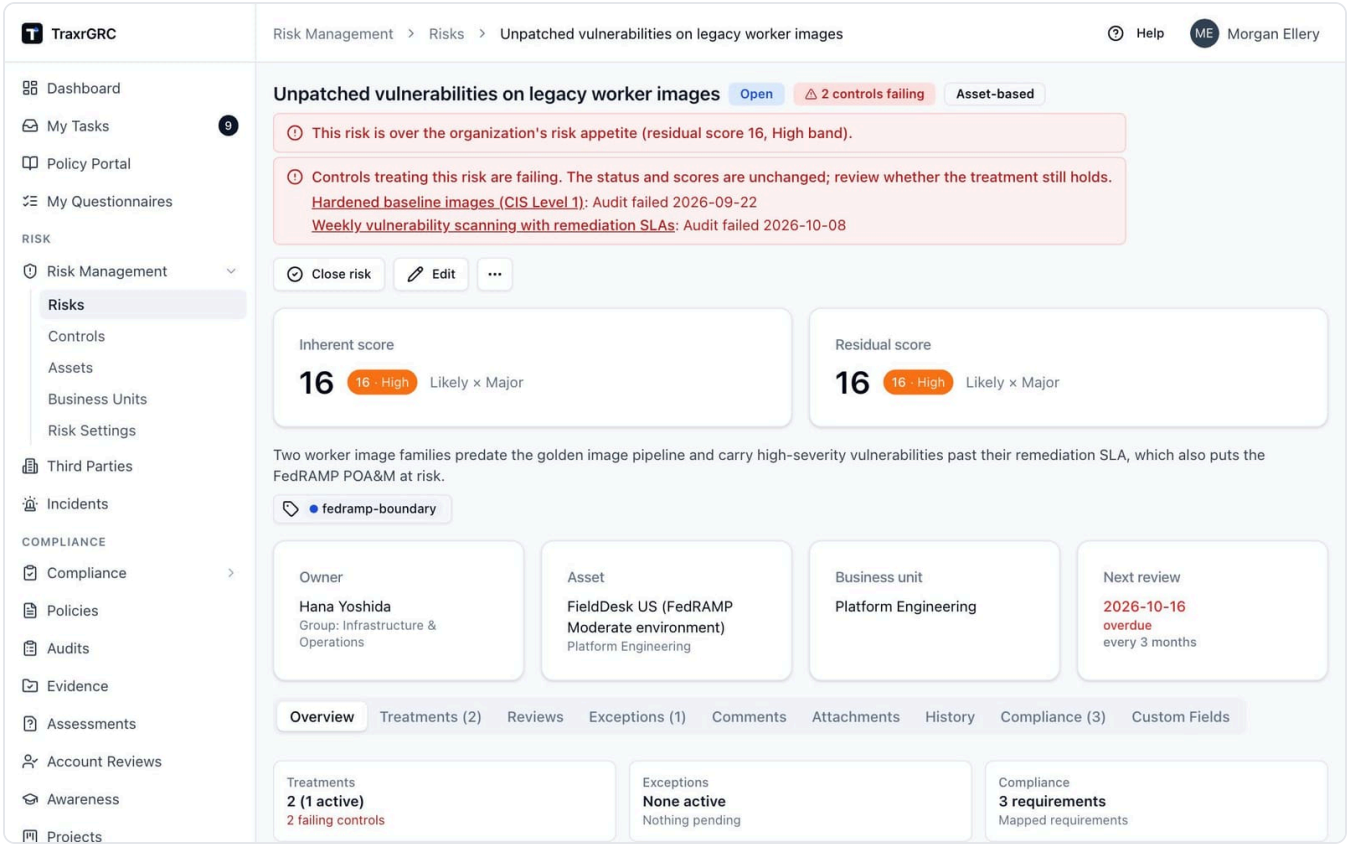


Figure 3.3. A risk over appetite with failing controls: the red banner, the failing-controls panel and the inherent and residual score cards.

Below the cards, one row of tabs covers the rest of the record: **Overview**, **Treatments**, **Reviews**, **Exceptions**, **Comments**, **Attachments** and **History**, then **Compliance (N)** (the compliance requirements this risk is mapped to; see Compliance Management; shown when you can read compliance) and **Custom Fields** (administrator-defined fields for risks; see Administration & Settings; editing requires risks:update on the record). The Treatments and Exceptions tab names show their counts. The **Overview** tab opens with three summary cards, each of which opens its tab when clicked: **Treatments** (how many, how many are active, and how many controls are failing), **Exceptions** (how many are active and when the next one expires, or how many await a decision) and **Compliance** (how many requirements the risk is mapped to; shown when you can read compliance). Below them it shows the created and updated times, the inherent and residual magnitudes as values ("L4 × I4"), and the completed-review history.

3.5.1 Status and actions

The action buttons depend on the status and your rights: risks:update on the record for **Close risk**, **Reopen**, **Archive** and **Edit**, and risks:delete for **Delete**, which is in the **More actions** menu (**...**).

Status	Meaning	Badge color	Allowed moves
open	Live risk, actively managed.	Blue	**Close risk**
closed	Accepted as resolved or no longer applicable.	Gray	**Reopen**, **Archive**
archived	Retired from the working register; terminal state.	Darker gray	none

Edit opens the **Edit risk** dialog with the current values. Saving recalculates the scores when likelihood or impact changed, and re-anchors the next review date when the frequency changed (see Reviews).

Delete asks for confirmation ("Delete risk?"), then removes the risk together with its treatments, reviews, exceptions, comments, attachments (files included), label assignments and custom-field values, and writes an audit-log entry. This cannot be undone.

3.6 Treatments

The **Treatments** tab lists the actions taken on the risk. Each treatment card shows strategy and status badges, the linked control, reduction badges ("–L1 –I2") for mitigate treatments, the description, the owner and the deadline. Strategies are color-coded: mitigate blue, transfer purple, avoid red, accept gray.

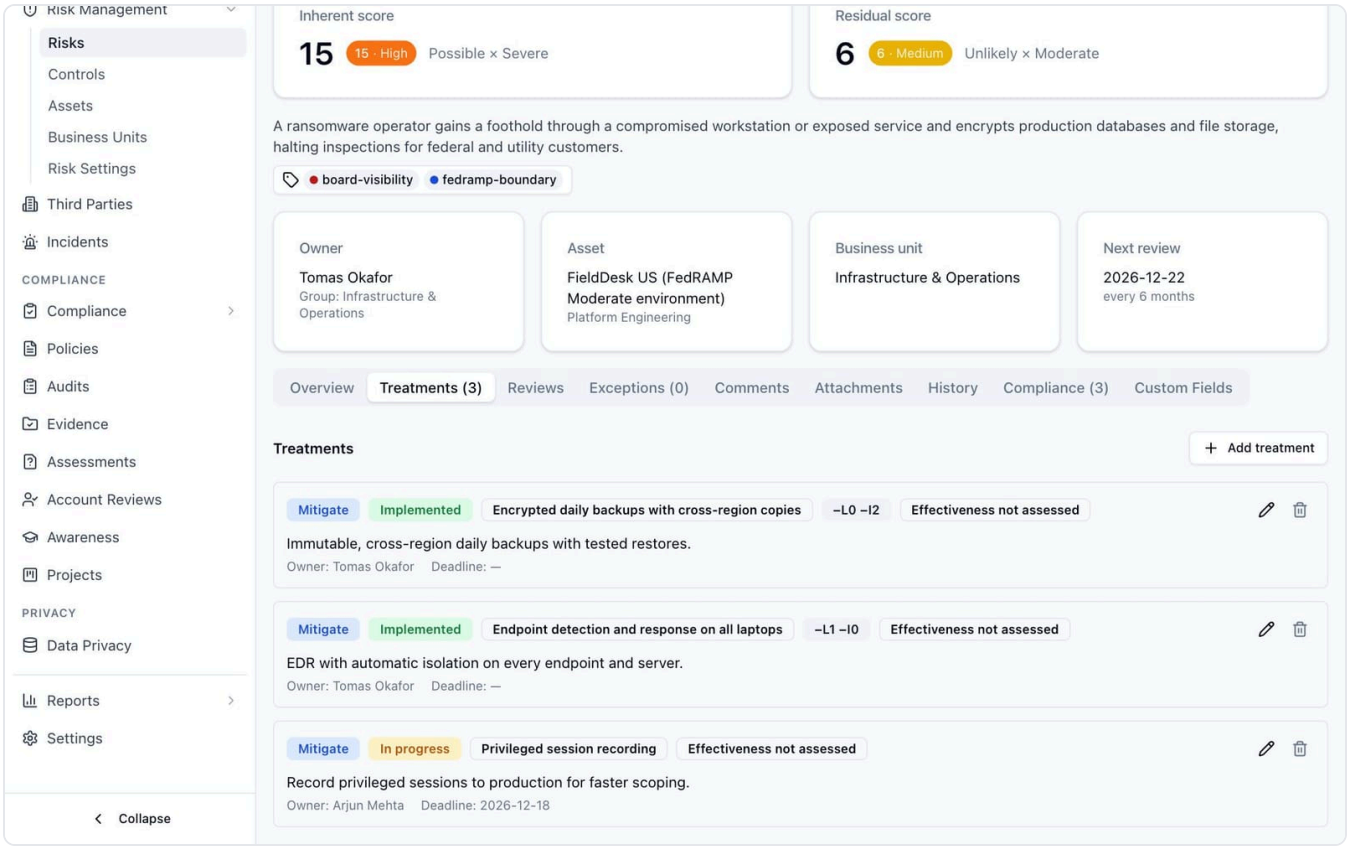


Figure 3.4. The Treatments tab of a risk, with implemented and in-progress mitigations, their linked controls and reductions.

To add a treatment:

1. Click **Add treatment**.
2. Complete the dialog:

Field	What it means
Strategy	mitigate (reduce the score), accept, transfer, or avoid.
Status	proposed, accepted, in_progress, implemented, or rejected.
Owner	Optional responsible user.
Linked control	Optional link to a catalog control. Controls are shared with compliance analyses.
Reduce likelihood by / Reduce impact by	0 to 4. Editable only for mitigate; how far this treatment lowers each magnitude once implemented.
Effectiveness	Mitigate only: how well it works in practice, E1 Effective, E2 Partially effective, E3 Largely ineffective, E4 Absent, or not assessed. It scales the reduction (see How scoring works). Assess it from the linked control's results; when that control is failing or overdue and the treatment is rated E1 or E2, the treatment shows **Control failing: review effectiveness**. The control's results never change the rating or the score by themselves.
Deadline	Optional target date.
Description	What the treatment actually is (required).

3. Click **Add treatment**.

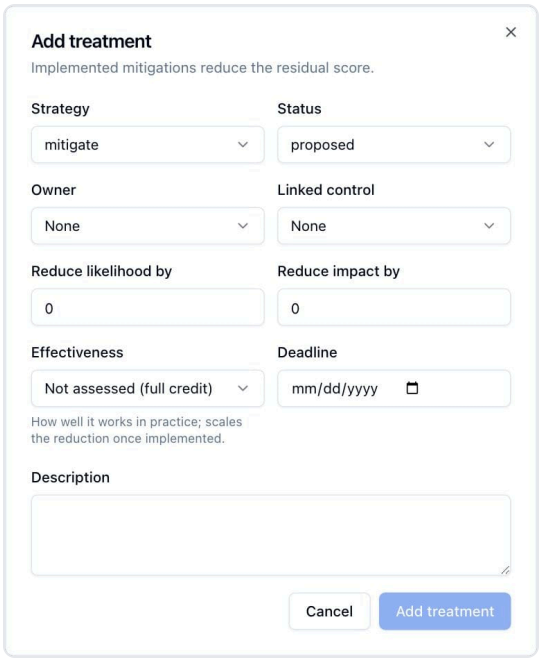


Figure 3.5. The Add treatment dialog with strategy, status, linked control and the likelihood and impact reductions.

Treatment statuses:

Status	Meaning	Badge color
proposed	Suggested, not yet agreed.	Gray
accepted	Agreed, not yet started.	Blue
in_progress	Being implemented.	Amber
implemented	In place. Only implemented mitigations reduce the residual score.	Green
rejected	Will not be done. Ignored by the failing-controls flag.	Red

Use the pencil icon on a card to edit a treatment and the bin icon to delete it. Adding, editing or deleting a treatment re-scores the risk immediately, which can move it into or out of an over-appetite band and can therefore trigger the over-appetite notification described above.

3.7 Reviews

The **Reviews** tab shows **Upcoming reviews** first, with past-due dates in red, and the completed-review **History** below it (result badge, time, reviewer and notes).

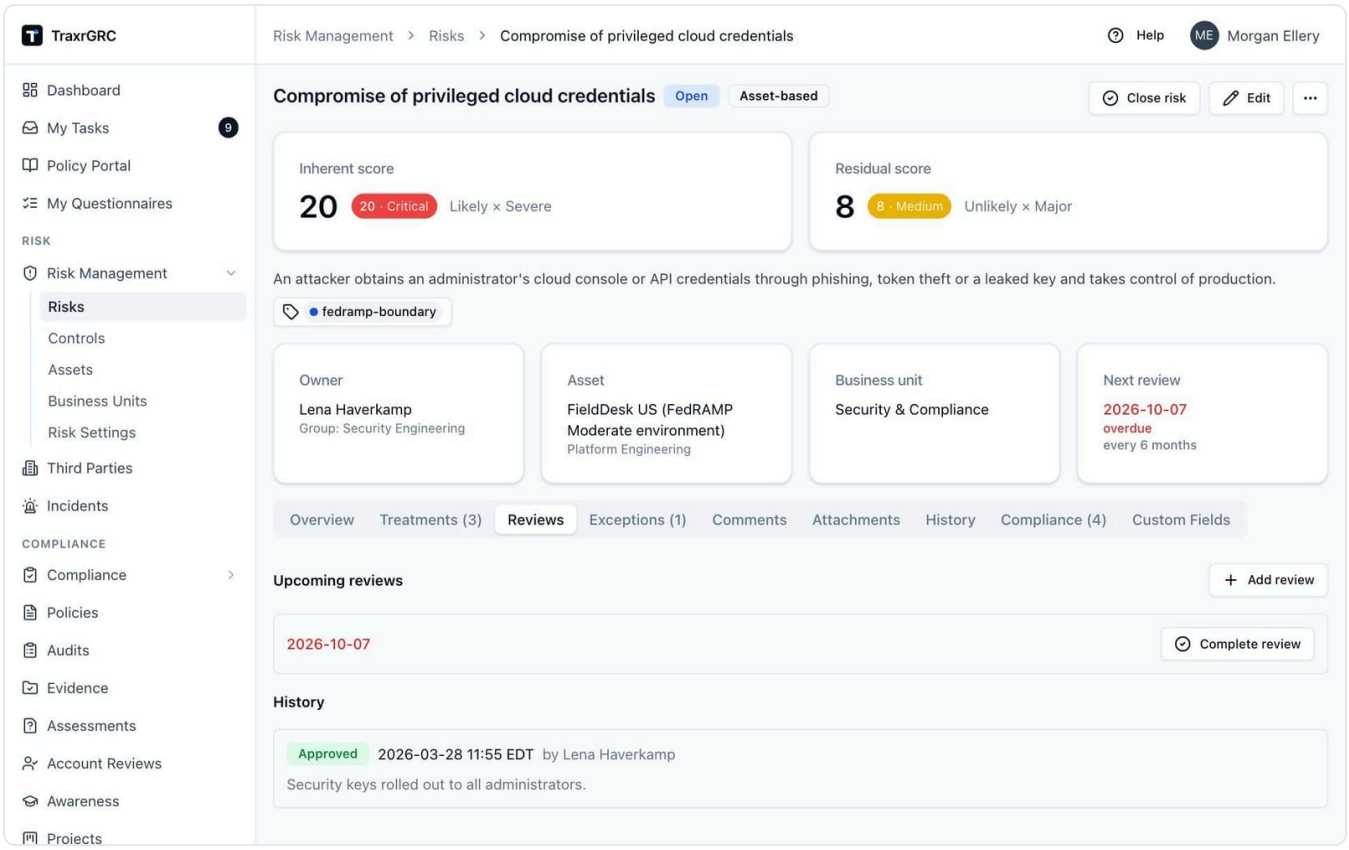


Figure 3.6. The Reviews tab of a risk with an overdue scheduled review and the history of completed reviews.

To complete a review:

1. Click **Complete review** on a scheduled review. The button appears only when you hold risks:review.
2. Choose a **Result** and add **Notes**:

Result	Meaning	Badge color
Approved	The risk was reviewed and is fine as it is.	Green
Changes required	The review found problems; follow-up is needed.	Amber
Accepted over appetite	The risk stays over appetite by explicit decision, once an exception request is approved (see below).	Red

3. Click **Complete review**.

Accepting a risk over appetite raises an exception request. Choosing **Accepted over appetite** opens three more fields: **Justification** (required), **Accepted until** (the expiry, required) and **Approver**

(optional; the list offers only active people who can see the risk, never you). Completing the review records the result and creates the request on the risk's **Exceptions** tab, linked to the review. Until it is decided the risk shows "Acceptance pending approval" above its tabs, and the review in the history reads the same. The request is decided like any other risk exception (see Exceptions): approved, the risk is accepted until the expiry date, with the usual expiry reminders and lapse; rejected, the risk shows "Acceptance rejected: plan a treatment" and you, the reviewer, get the email `risk.acceptance.rejected`. Reviews recorded with this result before the request was required keep their result and have no request.

Completing a review records you as the reviewer, schedules the next review at the risk's review frequency from today, updates the **Next review** date shown everywhere, and writes an audit-log entry. Completing reviews needs the separate risks:review right; risks:update does not include it.

To book an extra review without completing one:

1. Click **Add review** (requires risks:update).
2. Pick a **Scheduled date**.
3. Click **Schedule review**.

Next review moves to the earliest open review if the new one is sooner. Changing the review frequency on the risk re-anchors all open reviews from the last completed review (or the creation date).

The notification worker sends review reminders: "Risk review coming due" before the date and "Risk review overdue" once it has passed, individually or bundled into the daily digest when digests are enabled (see Notifications in Administration & Settings).

3.8 Exceptions

The **Exceptions** tab records time-boxed acceptances for the whole risk or for one treatment.

1. Click **Request exception** (requires risks:update).

2. Complete the dialog:

Field	What it means
Title	Short name of the exception (required).
Description	Why the exception is needed (required).
Linked treatment	One treatment, or **Whole risk** (the default).
Owner	Optional user responsible for the exception.
Approver	Optional user who must decide it. The list offers only active users who can see the risk, since nobody else could decide it, and never you.
Expires at	The date the exception runs out (required).

3. Click **Request exception**.

An exception starts as **requested** (amber) and is decided with **Approve** (green) or **Reject** (red). Exceptions follow a separation of duties:

- The person who requested an exception can never approve or reject it, administrators included, unless an administrator has relaxed separation of duties (see Sessions under Administration & Settings). A decision made that way carries an amber **Self-approved** or **Self-decided** badge.
- When the request names an approver, only that person (who must be able to see the record) or an administrator can decide it.
- When it names none, deciding needs the **Review** right on Risks (risks:review, the same right that completes risk reviews).

The exception card shows the owner, the designated approver, who requested it, who decided it, and "Linked to a treatment" when it covers one treatment. The **Approve** and **Reject** buttons appear only for someone allowed to decide. The decision is recorded with who made it and when, and an exception cannot be re-decided. Past its expiry date, an open exception displays as **expired** (gray) and the decision buttons disappear; a request past its expiry can still be rejected through the API but never approved. The nightly worker then records the exception as expired, with an audit entry, so the risk register's open-exception count, My Tasks, the API and notification rules on the **expired** status all agree. The same applies to compliance and policy exceptions. The notification worker sends "Risk exception expiring" reminders before an open exception runs out.

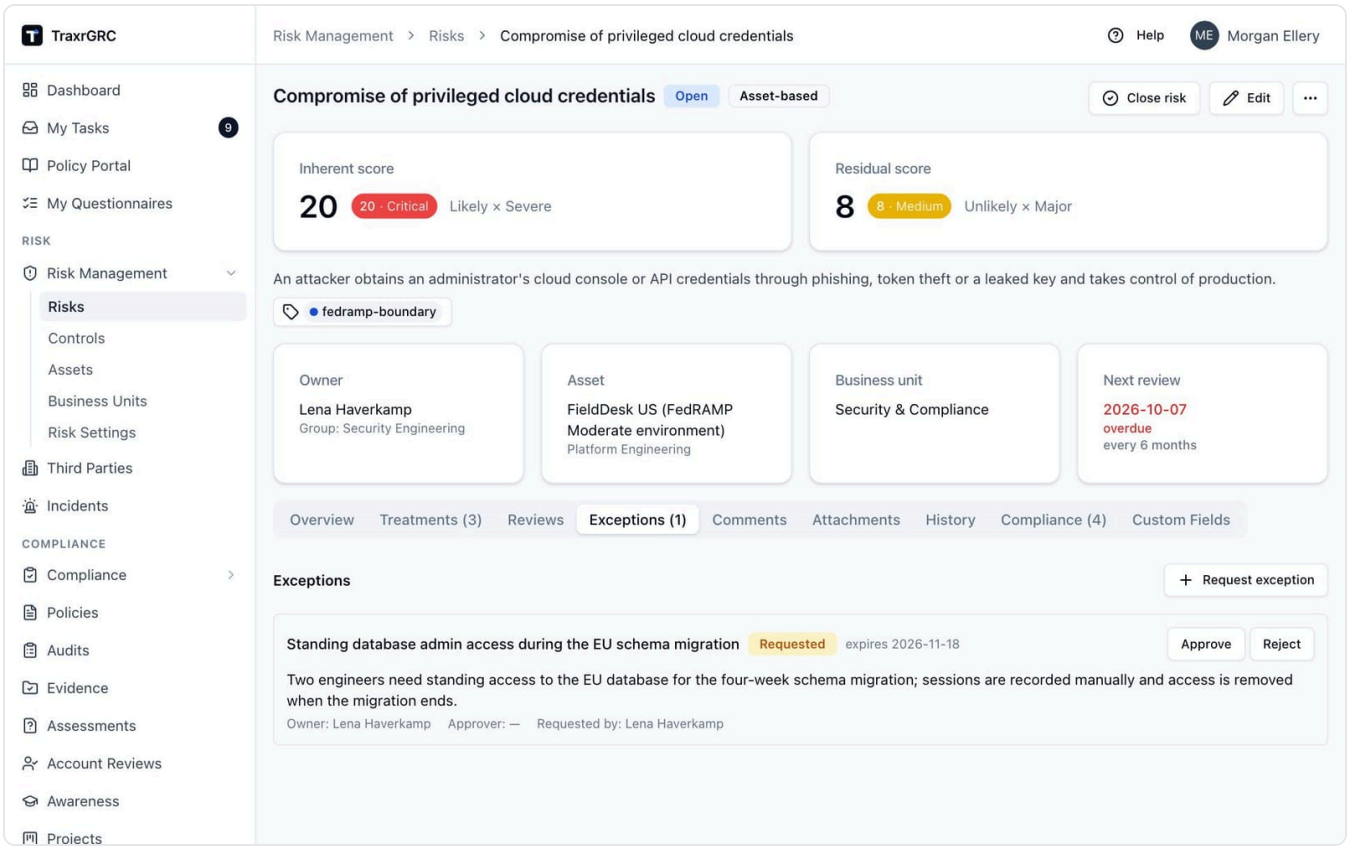


Figure 3.7. The Exceptions tab of a risk with a requested exception awaiting a decision, showing the Approve and Reject buttons.

3.9 Comments, attachments and history

- **Comments**: discussion thread on the risk. Posting requires risks:update on the record.
- **Attachments**: upload supporting files such as evidence or diagrams. Requires risks:update.
- **History**: the audit-log entries for this risk: action, actor, time and which fields changed.

3.10 Controls

Risk Management → Controls is the control library shared by risk treatments, compliance requirements and audit findings. The list has these columns:

Column	What it shows
Name	Click to open the control page.
Type	policy, process, technical, physical or other. Filterable.
Status	planned (gray), active (green) or retired (gray). Filterable.
Testing	The calculated testing status (see below).
Next test	The next scheduled audit or maintenance date.
Owner	The owning user.
References	How many treatments and compliance analyses link to the control, for example "2 treatments · 3 analyses".
Labels	The control's labels. Filter by typing part of a label name.

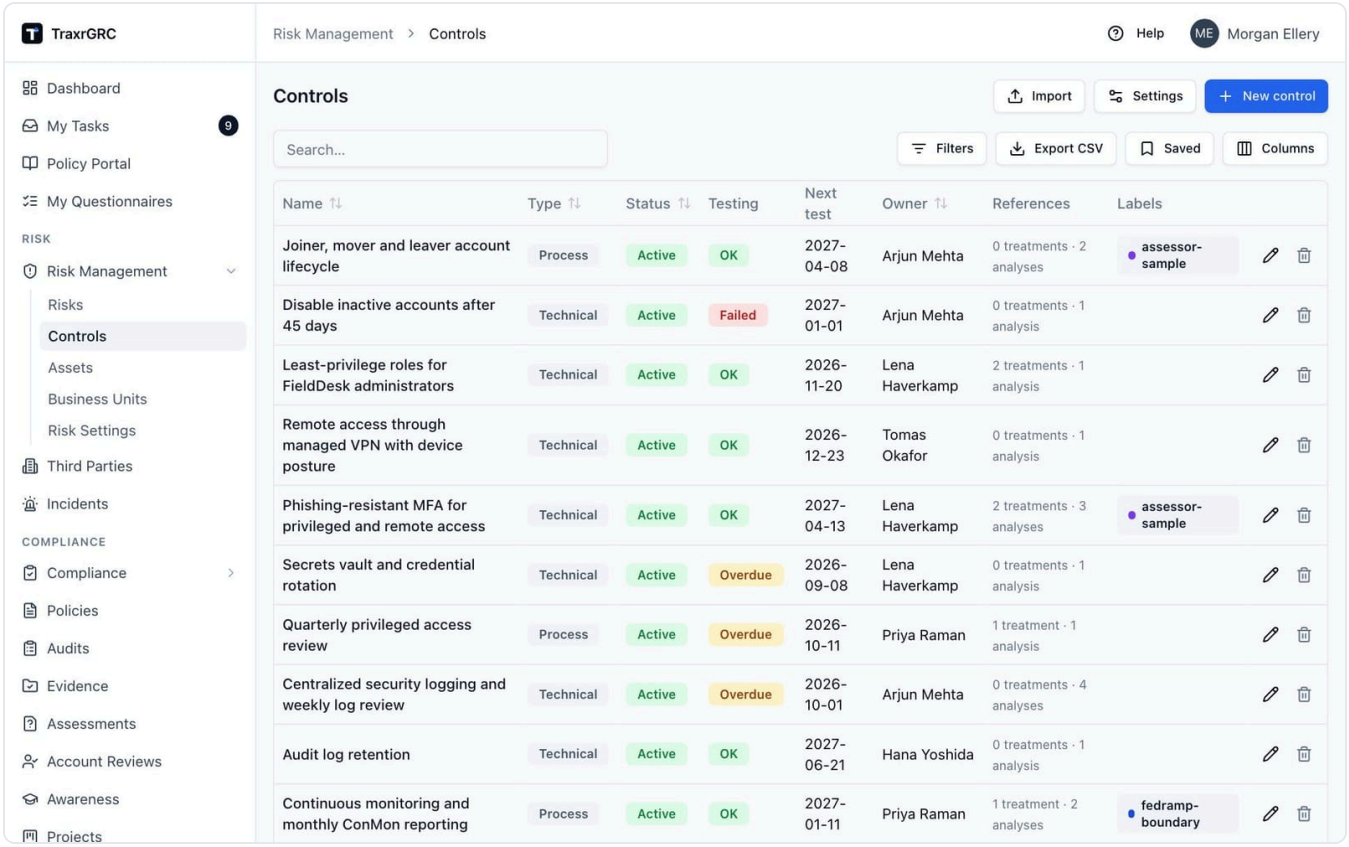


Figure 3.8. The control library with each control's type, status, testing status, next test date and references.

New control (controls:create) opens a dialog with **Name**, **Type**, **Status**, **Owner**, **Owner group** and **Description**. The pencil and bin icons on each row edit (controls:update) and delete (controls:delete) a control. **Import** loads controls from CSV (see Importing from CSV in Getting Started). Deletion is

blocked while any treatment or compliance requirement references the control (the confirmation dialog warns you), and for good once a test result has been recorded: results are evidence, so retire the control instead.

3.10.1 The control page

Click a control's name to open it. The header shows the name with type, status and testing-status badges, the description, and the **Labels** picker (see Labels in Getting Started; attaching needs controls:update on the control).

The description is laid out for reading. Each line you typed starts a new paragraph, and so do two spaces after a sentence. A line that begins with an enhancement code and a title followed by a colon, such as "AC-2(1) Automated System Account Management: ...", starts its own section with the code and title in bold, and the text after the colon starts with a capital letter. A paragraph longer than about 1,200 characters is split at sentence ends into paragraphs of a few sentences each. The stored text is not changed; **Edit** shows it exactly as saved. Cards show the owner (and owner group) and, for **Audit** and **Maintenance** separately, the last result, the next due date (red with "(overdue)" when it has passed) and the schedule ("Every 6 months after completion", "Each year on 15 Apr" or "Not scheduled").

Below the cards are seven tabs:

Tab	What it holds
Audits	The **Audit metric** and **Success criteria**, the **Scheduled** audits and the **Results**.
Maintenance	The **Maintenance task**, the scheduled maintenance and the results.
Used by (N)	The risk treatments, compliance requirements and audit findings that link to the control.
Comments	Discussion on the control.
Attachments	Evidence files.
History	Audit-log entries for the control and its tests.
Custom Fields	Administrator-defined fields for controls (see Administration & Settings). Editing requires controls:update on the control.

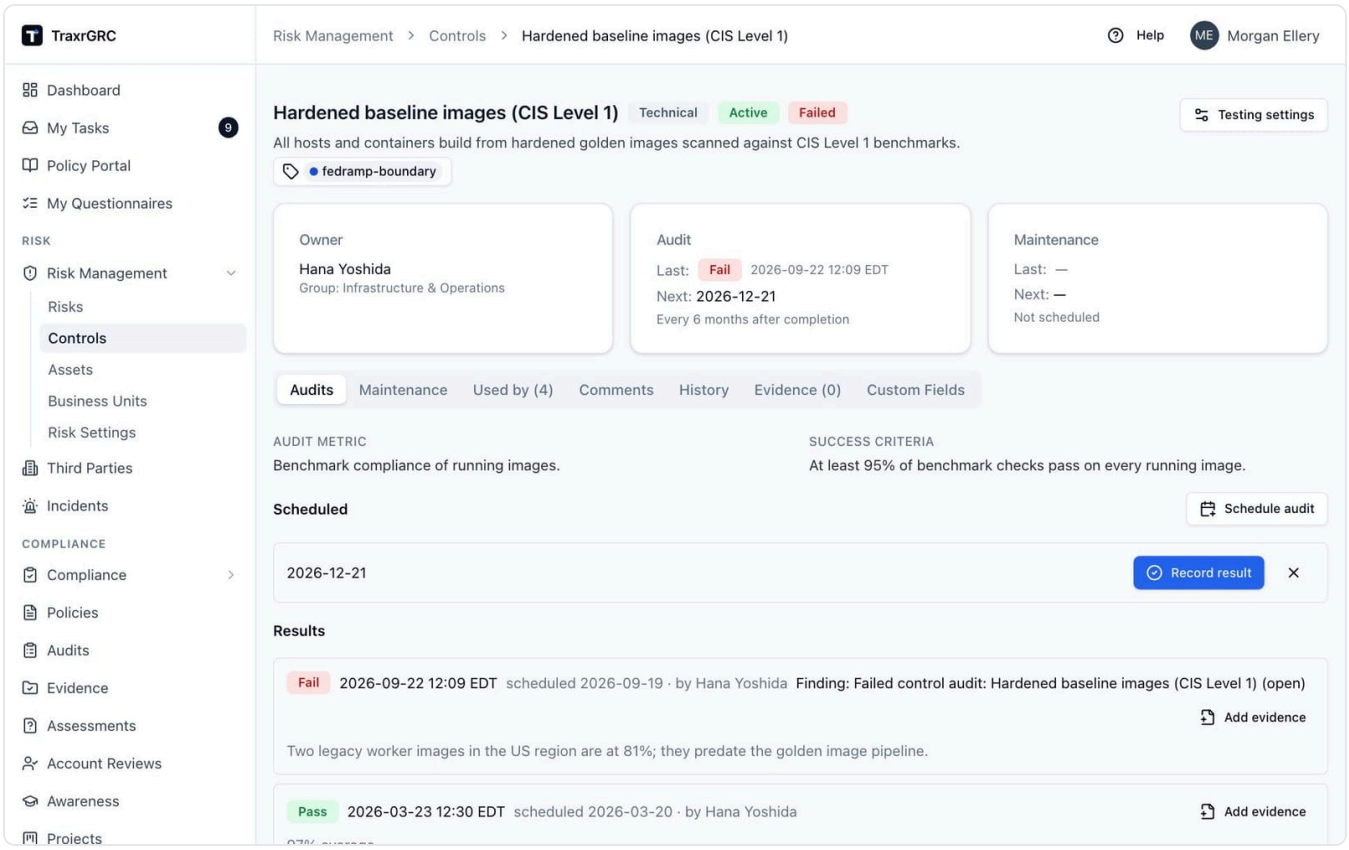


Figure 3.9. A control page with a failed audit, the audit metric and success criteria, and the finding raised from the failed result.

The calculated testing status is:

Status	When
Failed	The latest audit or maintenance failed.
Overdue	A scheduled test is past its date.
OK	Every scheduled kind's latest result passed.
Untested	None of the above, for example no result yet.
Retired	The control is retired.

3.10.2 Scheduling and recording control tests

Testing settings (controls:update) sets the schedule for audits and maintenance separately, plus the audit metric, success criteria and maintenance task. Each schedule is one of:

- **Not scheduled**.
- **Every N months after completion** (1 to 60). A late test pushes the schedule back.

- **Fixed dates each year**: month and day, for example `03-01, 09-01`. 29 February is not allowed, and a test done early or late falls back onto the next fixed date.

Switching a schedule on books the first test; changing an active schedule moves the next open test.

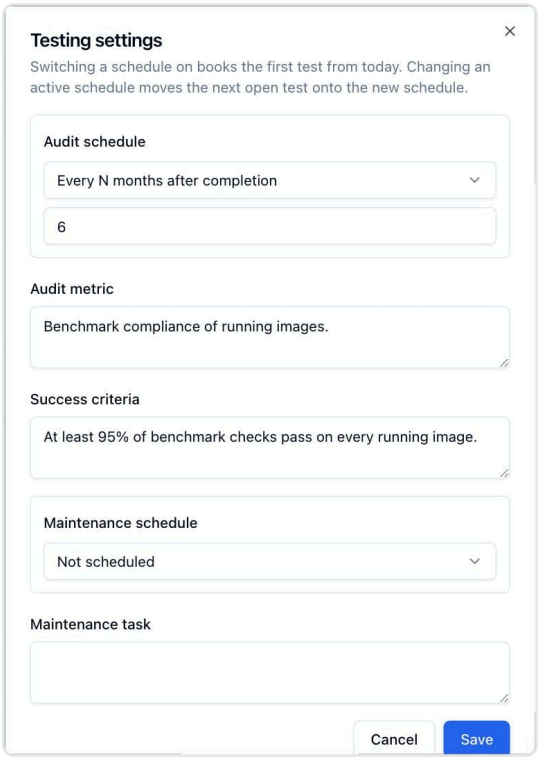


Figure 3.10. The Testing settings dialog with the audit schedule, audit metric, success criteria and maintenance schedule.

On the **Audits** or **Maintenance** tab:

- **Schedule audit** or **Schedule maintenance** adds a one-off test alongside the recurring schedule. Not offered for retired controls.
- **Record result** on a scheduled test opens a dialog that shows the success criteria (for audits). Choose **Pass** or **Fail** (there is no default, so a result is always a decision), add notes and click **Record result**. Recording a result books the next test, and the confirmation offers **Add evidence** to record the evidence for it (see Evidence). Each recorded result also has an **Add evidence** button.
- The **×** button removes a test that is still open.

A failed audit emails the control owner. Each result shows the time, the scheduled date and the tester, plus a link to the finding raised for it, if any. While a finding raised by a failed audit is still open, the control gets no second one: a later failure adds the comment "Failed again on *date*" with its notes to that finding's thread, and a passing re-audit adds "Passed the re-audit on *date*; ready to verify and close", for whoever runs the audit plan to close it. These comments are marked **Recorded by the app** and cannot be deleted; the finding's description stays as the auditor wrote it.

3.10.3 Control testing settings

The **Settings** button on the Controls page (settings:read to view, settings:update to change) opens **Control testing settings**. **Raise findings for failed audits in** names an audit plan that receives a finding for every failed control audit. The finding is owned by the control owner, due by the re-audit date, linked to the control, and shown next to the failed result on the control page. Choose **Off: do not raise findings** to stop raising them; deleting the chosen plan also turns it off. You must create an audit plan before you can turn this on.

3.11 Assets

Risk Management → Assets lists the things you protect, with **Name**, **Business unit**, **Owner**, **Risks** (how many risks link to the asset) and **Labels** columns. Assets have no page of their own, so labels are attached in the list: with assets:update on an asset, its **Labels** cell is the label picker; otherwise it shows the labels read-only. Each asset has a name, description, optional business unit and optional user or group owner; risks can then be linked to it.

New asset (assets:create) opens the asset dialog; the pencil icon on a row edits an asset (assets:update) and the bin icon deletes it (assets:delete). **Import** loads assets from CSV.

Tick **A system that holds or processes personal data** to mark an asset as a system. It then takes a **Hosting country** and the **Categories of personal data held** (type a category and press Enter), shows a "System" badge with its hosting country in the list, comes first when a Data Liaison picks systems, and carries its hosting country on the data map (see Data Privacy).

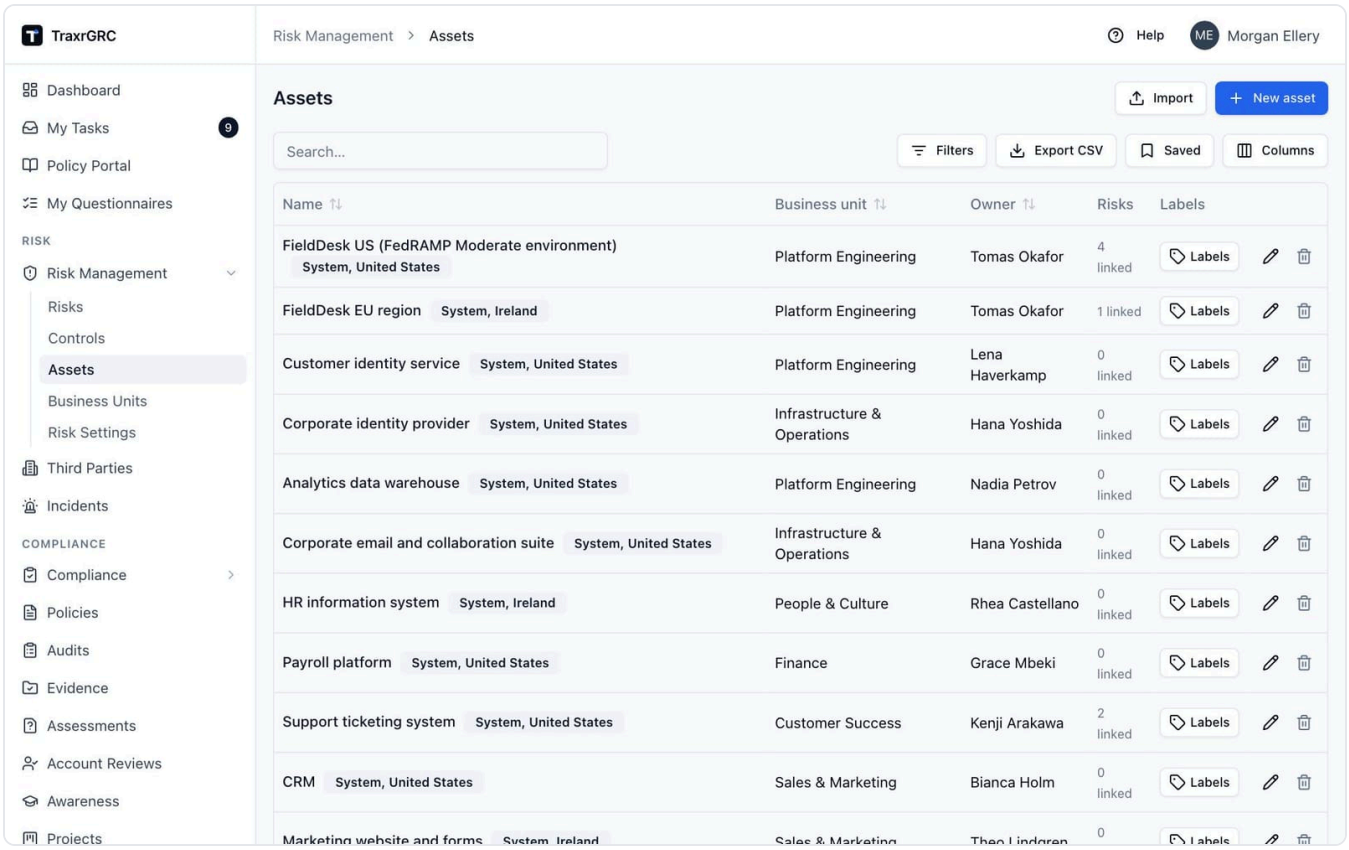


Figure 3.11. The asset catalog, with systems that hold personal data marked by a System badge and their hosting country.

Deleting an asset keeps its risks but clears the asset link on them, and removes the asset's comments, attachments (files included), labels and custom-field values.

3.12 Business Units

Risk Management → Business Units holds your organizational hierarchy, with **Name**, **Parent**, **Owner** and **Assets** (count) columns. A unit can have a parent (forming a tree), an owner and a description. Create units with **New business unit** (businessUnits:create) and edit or delete them from the row icons.

When editing, a unit cannot be parented to itself or any of its descendants: the parent list leaves those out, and the server rejects cycles outright. Deletion is blocked while child units exist. Linked assets and risks are kept (they lose the reference), and the unit's comments, attachments (files included), labels and custom-field values are removed with it.

3.13 Risk Settings

Risk Management → Risk Settings configures the scoring engine. Viewing requires settings:read; the edit controls require settings:update (creating and deleting records also uses settings:create and

settings:delete). The screen has four panels plus a maintenance button:

- **Classifications**: the register tabs. Each has a name, a unique slug (derived from the name if left blank), an optional description and a sort order. Add one with **New classification**. Deletion is blocked while any risk uses the classification.
- **Magnitudes**: the likelihood and impact scales, side by side, each with its own **Add** button. Each entry has a kind (Likelihood or Impact), a name (for example *Rare* or *Severe*), a numeric value, a sort order and an optional description. Names appear in the risk form as "Name (value)". Values must be unique within a kind. A value in use by any risk or matrix cell cannot be changed or deleted until those references are removed.
- **Calculation matrix**: one editable score per likelihood and impact cell. Each cell is tinted live with the color of the appetite band its current score falls into. Click **Save matrix** to keep your changes; saving recalculates every risk immediately.
- **Appetite bands**: name, minimum score, color (preset swatches or any custom hex), an **Over appetite** switch ("Risks in this band count as appetite breaches.") and a sort order. Add one with **New band**. A bar above the table previews the bands in minimum-score order. Creating, updating or deleting bands recalculates all risks.
- **Impact axes**: the axes impact is scored on, each with a name and an optional description. Every axis uses the impact magnitudes. Adding an axis starts every risk at its current impact on it, so no score moves until someone rescores an axis. Removing one asks for confirmation, says how many risks have it as their highest axis, and lowers those risks to their next highest axis with residual rescored; each lowered risk's history records it and the audit log keeps the removed scores. While axes exist, the risk form shows one impact per axis and the highest becomes the risk's impact. The API, CSV import and other paths that send a single impact set every axis to it when the value differs from the current impact; sending back the same value changes nothing.
- **Recalculate all risks**: recomputes the residual score of every risk using the current matrix, bands and treatments, and fires over-appetite notifications for any risk that crosses into an over-appetite band. Use it as a maintenance sweep after bulk changes made outside the normal flows.

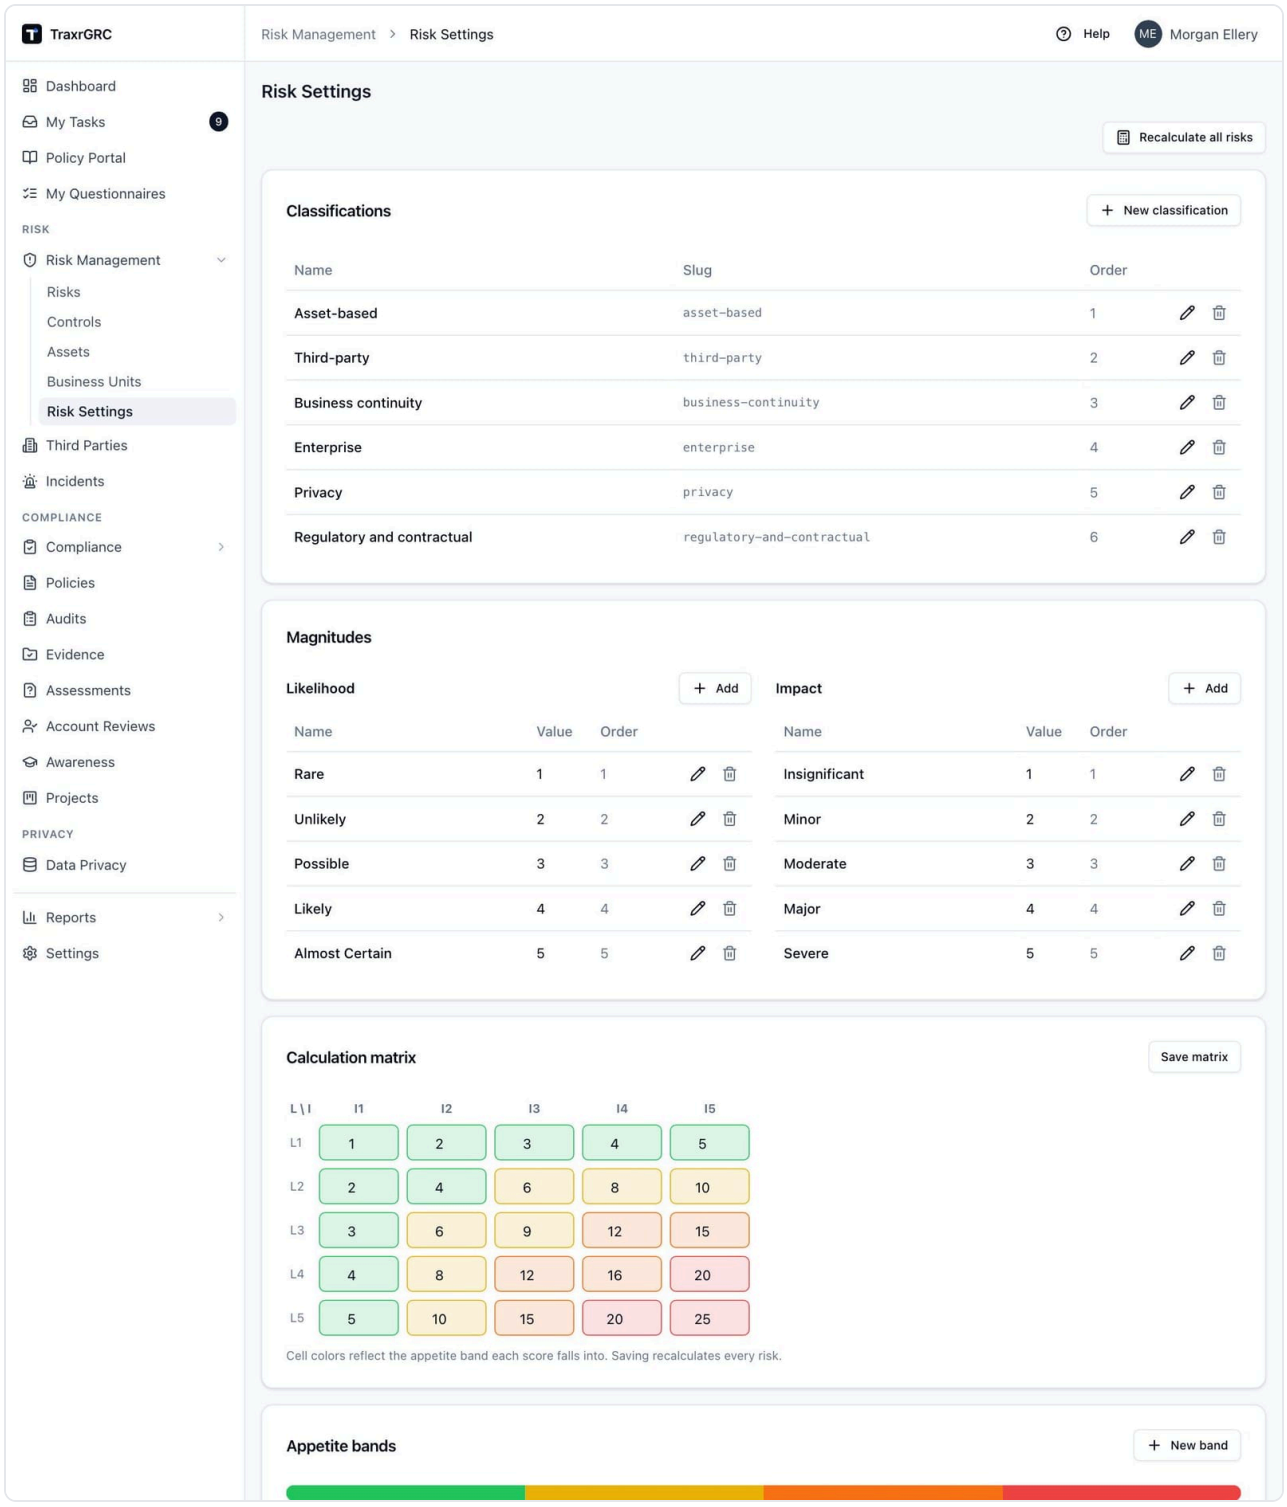


Figure 3.12. The Risk Settings screen with classifications, likelihood and impact magnitudes and the color-tinted calculation matrix, with the appetite bands below.

3.14 Who can do what (Risk Management)

Action	Permission
See the Risks list and risk detail	risks:read (at 'own' level you see only risks you or your groups own)
Create a risk	risks:create
Edit a risk, change its status, add/edit/delete treatments, schedule extra reviews, request exceptions, comment, upload	risks:update: 'all' for every record, 'own' only for records you or your groups own
Complete a review; decide an exception that names no approver (never one you requested)	risks:review (same all/own levels)
Decide an exception that names you as approver (never one you requested)	risks:read on the risk
Delete a risk	risks:delete (all/own)
Manage controls and record control tests	controls module rights (create/update/delete, with the same all/own levels on update/delete)
Control testing settings (findings for failed audits)	settings:read to view, settings:update to change
Manage assets	assets module rights
Manage business units	businessUnits module rights
View Risk Settings; change classifications, magnitudes, matrix and bands; recalculate	settings:read to view; settings:update to edit (settings:create and settings:delete for adding and removing records)

Administrators bypass all of the above.

CHAPTER 4

Compliance Management

Compliance Management is where you assess your organization against a standard or framework. You load the framework's clause structure as a *package*, record a status and owner for every clause, map the controls, policies and risks behind each one, link the projects that remediate its gaps, raise time-boxed exceptions, and export a board-ready PDF report. Compliance managers and GRC analysts maintain packages and run assessments; requirement owners record their analysis; approvers decide exceptions; auditors and executives read the workspace, the gap view and the report.

The **Compliance** entry in the sidebar expands to these pages:

Page	Screen title	What it is for
Analyses	Compliance analyses	The assessment workspace: one row per requirement, per package.
Packages	Compliance packages	The frameworks and their clause trees.
Gaps	Compliance gaps	The POA&M view of every open gap and its remediation projects.
POA&M	Plan of action and milestones	The POA&M register: permanent items from requirement status, audit findings and entries.

Access is controlled by the `compliance` module permissions (see Who can do what, at the end of this chapter).

4.1 Packages

A package is one framework: a named set of items (usually domains at the top level and clauses beneath them) with a standard reference, version and description. **Compliance → Packages** lists every package:

Column	What it shows
Name	Click to open the package.
Standard	The standard reference, for example *NIST SP 800-53*.
Version	The version of the standard.
Items	How many items the package holds, headings included.
Analyses by status	A stacked mini-bar of requirement statuses, with "+ N additional" beneath it when the package has additional items. With compliance:read at 'own' level it counts only the analyses you own.
Created	When the package was created.

Each row has a download icon that exports the package (see Exporting a package), and, with the matching rights, pencil and bin icons to edit or delete it. **New package** (compliance:create) registers an empty package with a **Name** (required), **Standard**, **Version** and **Description**.

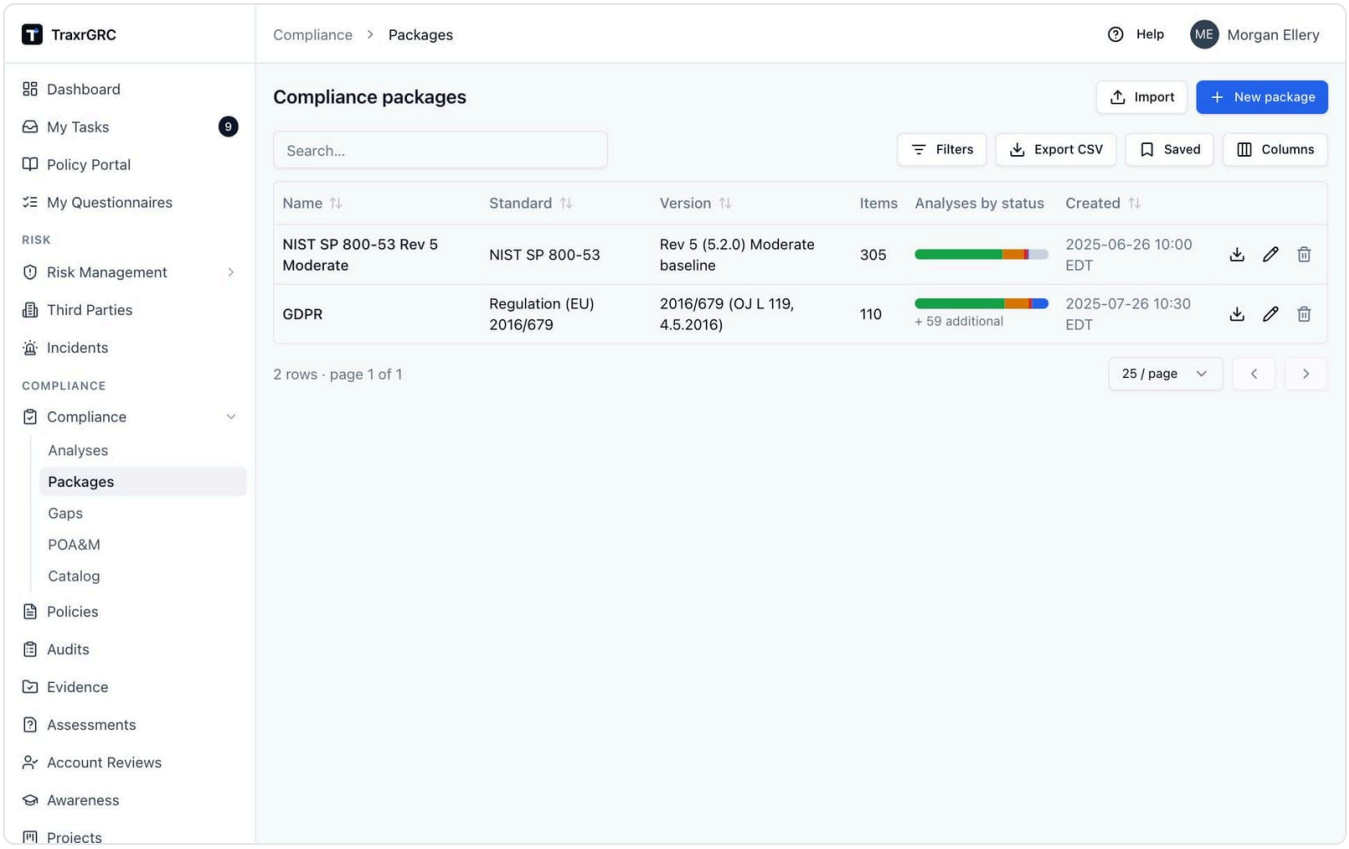


Figure 4.1. The compliance packages list with each package's standard, version, item count and stacked status bar.

4.1.1 The package page

Click a package name to open it. The header shows the name, a badge with the standard and version, the creation time, the item count and the description. Three buttons sit beside it:

- **Report** generates the PDF compliance report (see The compliance report).
- **Import crosswalk** (compliance:update at 'all') links this package's requirements to another package's (see Crosswalks between frameworks).
- **Edit metadata** (compliance:update) changes the name, standard, version, description and the provider package it inherits common controls from (see Common controls and inheritance).
- **Add item** (compliance:create) adds a top-level item.

Under the description the header says which package this one inherits common controls from, which packages inherit from it, and how many crosswalk links connect it to each other package.

Below the header is the item tree, indented by depth and ordered by sort order and code. Each row shows the code, the name (click it to open the item) and a kind badge for headings and additional items. The row icons are:

Icon	Action	Permission
Up and down arrows	Move the item among its siblings.	compliance:update
Plus	Add a child under the item.	compliance:create
Pencil	Edit the item.	compliance:update
Bin	Delete the item (see Deleting packages and items).	compliance:delete

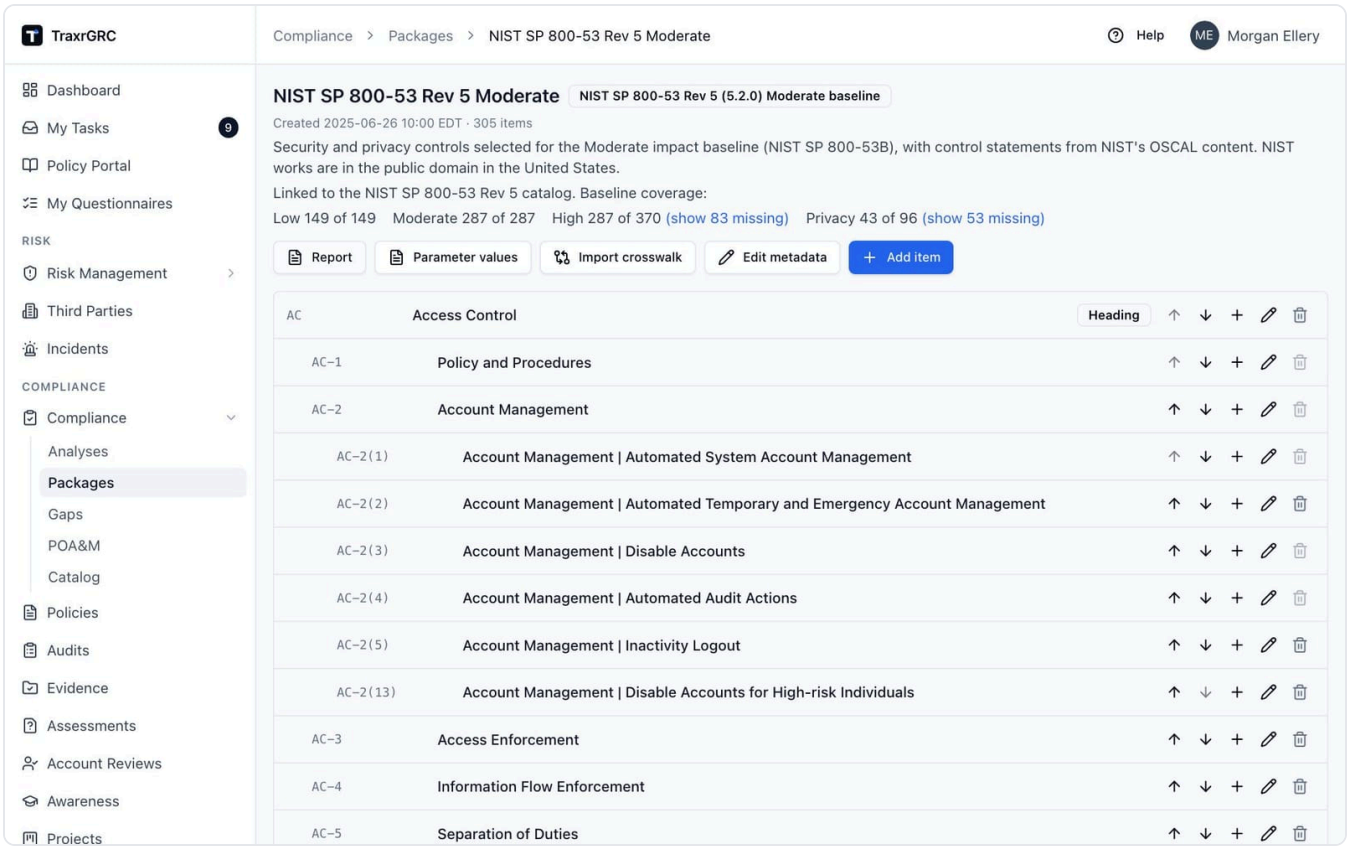


Figure 4.2. The NIST SP 800-53 package page with its item tree: family headings and the controls and enhancements beneath them.

Every item has a **Code** (for example *A.5.1*), a **Name**, an optional **Description**, and a **Counts as** setting. The code must be unique within the package. The dialog is titled **Add top-level item**, **Add child under** followed by the parent's code and name, or **Edit item**.

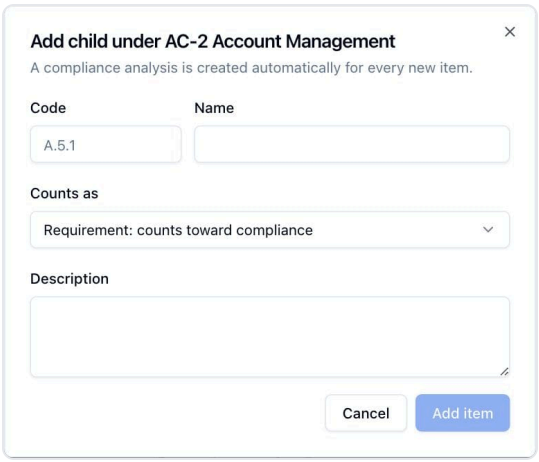


Figure 4.3. The Add child dialog with the code, name, Counts as and description fields.

Counts as sets one of three kinds, shown as a badge in the tree:

Kind	Option in the dialog	Meaning
Requirement	Requirement: counts toward compliance	The default. Counts toward compliance: the status counts, the compliance percentage and the dashboard. No badge.
Additional	Additional: implemented, not required, reported separately	Implemented but outside the required set (for example, controls the organization documents but does not have to meet). Reported in its own section with its own counts and percentage; never part of the compliance percentage. Marked with an "Additional" badge.
Heading	Heading: structure only, not counted	Structure only, such as a family or domain title. Counted nowhere and left out of the analysis workspace; its children stand in its place.

A compliance analysis is created automatically for every new item.

4.2 Importing a package

Importing builds a whole package, metadata plus the item hierarchy, from a file.

1. On the Packages page, click **Import** (requires compliance:create).
2. In the **Import compliance package** dialog, choose a CSV or JSON file:
   - **CSV**: a header row with `code` and `name` columns plus optional `description`, `parent_code` and `kind` (`requirement`, `additional` or `heading`; blank means requirement) columns. Each row after the header is one item; `parent_code` points at another row's code to build the tree. A CSV carries no metadata, so you fill in the **Package name** (required), **Standard** and **Version** in the dialog.
   - **JSON**: an object of the form `{ "name": "…", "standard": "…", "version": "…", "description": "…", "items": [{ "code": "…", "name": "…", "description": "…", "parentCode": "…", "kind": "…" }] }`. Name and a non-empty items array are required; the metadata is read from the file.
3. Check the preview. The dialog shows either the parsed tree (indented, first 20 items) or a list of validation errors.
4. Click **Commit import**.

The whole file is rejected when any problem is found: empty codes or names, duplicate codes, references to a `parent_code` that does not exist in the file, and parent cycles. Nothing is written until the file is completely clean. Each imported item gets its own "Not reviewed" analysis.



Figure 4.4. The Import compliance package dialog, before a file is chosen.

TraxrGRC does not bundle licensed frameworks such as ISO 27001. Import clause structures only from exports you are licensed to use (a purchased toolkit, or your own earlier export; see below).

4.3 Exporting a package

The download icon on every package row produces a JSON file in exactly the import shape above. Exports carry each item's kind and re-import cleanly, so you can copy a package between workspaces or keep versioned snapshots of a framework. A package linked to a catalog exports its `catalog` key and each requirement's own organization-defined values as `parameters` (a map from parameter id, such as `ac-02_odp.01`, to text; values inherited from a provider are not exported, and values on analyses you cannot read are left out without notice, so a copy made from your export lacks them; a value the control no longer has, after the package's catalog or an item's code changed, is left out too). Importing a JSON file with `parameters` stores them on the new analyses; every id must be one of the control's parameters in the named catalog, and headings take none, or the whole file is refused. CSV imports carry neither.

4.4 Common controls and inheritance

Most organizations implement many controls once, at the corporate level, and every system inherits them: NIST calls this control origination. A system package can name a **provider** package of the same framework, for example *NIST SP 800-53 Moderate: Trial platform* inheriting from *NIST SP 800-53 Moderate: Corporate*. Each requirement of the system package then has an **Origination**:

Origination	Effective status
System-specific	The system package's own status (the default).
Common (inherited)	The provider's status for the requirement with the same code. The system package's own status is kept but does not count.
Hybrid	Both parts apply, so the worse of the system package's own status and the provider's. Not reviewed counts as worse than partially compliant, since an unreviewed part may hide anything.

If the provider has no requirement with the same code, the inherited part counts as Not reviewed and the requirement page says why.

To set it up:

1. Open the system package and click **Edit metadata**.
2. Choose the provider in **Inherits common controls from** and save.
3. Mark requirements Common or Hybrid on each requirement page, or many at once with **Set origination...** in the analysis workspace.

The effective status is worked out each time it is shown and is never stored, so a change in the provider reaches every system package straight away. It is what the analysis workspace badges, the package counts, the dashboard's compliance donut, the compliance status report and the PDF report use; the filter and sort in the workspace still use the status recorded in the package. Like the control health roll-up, it is the same for every viewer: someone who cannot read the provider's analysis still sees the inherited status.

Inheritance is one level deep: a provider has no provider of its own, and a package that others inherit from cannot inherit itself. A provider cannot be removed while any requirement is still Common or Hybrid (set them back to System-specific first), and a package cannot be deleted while another inherits from it.

4.5 The NIST SP 800-53 catalog

The application ships NIST SP 800-53 Rev 5 (release 5.2.0) as reference data: all 1,196 controls and enhancements (182 of them withdrawn), their families, statements and organization-defined parameters, and their membership of the Low, Moderate, High and Privacy baselines from NIST SP 800-53B. A package whose **Reference catalog** is set to *NIST SP 800-53 Rev 5* (in **Edit metadata**) matches its requirement codes, such as `AC-2(1)`, against it. Packages loaded before 2026-10-18 whose standard names 800-53 Rev 5 were linked automatically.

A catalog package shows:

- **Baseline coverage** on the package page: "High 287 of 370", for example, with a link to the missing controls in the catalog.
- **Baseline badges** on each requirement page.
- A **Statement and organization-defined values** card on each requirement page (see below), and a **Parameter values** CSV export on the package page.

4.5.1 Organization-defined values

Many controls leave a value to the organization: "Review accounts for compliance with account management requirements [Assignment: organization-defined frequency]". The card shows the control statement with each value you have set highlighted in green and each unset one in amber, then a field per parameter: a text box, a list for a single choice, or tick boxes for "one or more". Click **Save values**

(compliance:update on the analysis). Clearing a field removes the value. Every change is recorded on the requirement's history.

A Common or Hybrid requirement whose package inherits from a provider shows the provider's value for the same parameter where it has none of its own ("Inherited from the provider: ..."); setting a value here overrides it. Unlike the inherited status, an inherited value is shown only to people who can read the provider's analysis, since values are free text. Saving sends only the fields you changed, so a value someone else saved meanwhile is kept.

Parameter values on the package page downloads every parameter of every requirement with its value and source (own or provider), one row per parameter, for a system security plan. It needs compliance:read at 'all'.

4.5.2 The catalog page

Compliance → Catalog lists the whole catalog. Filter by **Family**, **Baselines** or text; tick **Include withdrawn controls** to see those too. Click a control's code to read its statement.

Choose a **Package** to add an **In the package** column with each control's effective status there, and a **Show** filter: all controls, those in the package, those missing from it, or gaps (missing, not reviewed, partially or not compliant). "Every High control we are not covering" is the Baselines filter set to High with Show set to gaps. Statuses of analyses you cannot read are not shown and never count as gaps.

With compliance:create and compliance:update at 'all', tick missing controls and click **Add to package**. Each is added as a requirement with its statement and an empty analysis, under its base control when the package holds it, otherwise under its family heading (created if needed), in catalog order among its siblings; items you reordered by hand keep their order. This is how a Moderate package grows to High without a second import.

4.6 Crosswalks between frameworks

A crosswalk links requirements in two packages, so a control implemented once against NIST SP 800-53 shows what it says about ISO/IEC 27001 or SOC 2. Each link has a direction (the source's status indicates the target's) and a **relation**: Equivalent, Partial or Supports.

The indication never sets the target's status. Published mappings are a general indication that a person must judge (NIST says so of its own), so the target requirement page shows the linked statuses beside the status you set: "Indicated by 4 requirements: 3 compliant, 1 partially compliant".

To import a crosswalk:

1. Open either package and click **Import crosswalk**.
2. Choose the **Other package**, and whether the file's source column holds this package's codes or the other package's.

3. Choose a CSV file with `source` and `target` columns (optional `relation`, blank meaning Supports, and `note`), or click the built-in **NIST SP 800-53 Rev 5 to ISO/IEC 27001:2022** mapping (NIST OLIR 155, 643 links, source: 800-53 codes).
4. Click **Import links**.

Codes must match each package's codes exactly. Rows whose code is not in a package are skipped and listed under the header, so one file serves any baseline; headings are never linked. A pair already linked takes the file's relation and note. Importing never removes links. The same CSV ships in the repository as `docs/crosswalks/nist-sp800-53r5-to-iso27001-2022.csv`. ISO text is copyrighted and not included: load the ISO package from your own licensed copy with codes like `5.2` and `A.5.15`.

On a requirement page the **Other frameworks** card lists what this requirement **Indicates** and what it is **Indicated by**, each with its package, relation and (effective) status. With compliance:update at 'all', **Link requirement** adds a link by package and code, and the cross beside a link removes it. Each link and removal is recorded in the audit log, and so is each link an import rewrites. Statuses of linked analyses you cannot read are counted, not shown, and left out of the summary. The built-in mapping sets the direction itself: it finds which of the two packages holds the 800-53 codes.

4.7 The analysis workspace

Compliance → Analyses (titled **Compliance analyses**) is where the assessment happens. It has one tab per package, in the order the packages were created. Each tab lists every requirement and additional item in the package as a row; headings are left out. Until a package exists, the page links you to the Packages page instead.

Column	What it shows
Code	The clause code (monospace).
Name	The clause name.
Scope	"Requirement", or an **Additional** badge. Filter on it to list one kind.
Status	The analysis status badge (see Analysis statuses).
Owner	The assigned assessor.
Controls	The mapped controls by name, or a count ("2 mapped") when you cannot read controls.
Control health	The worst testing status across the mapped controls ("Controls: Failed", "Controls: OK" and so on, or "No controls"), an **At risk** badge, and the number of mapped policies and risks. Filter on it to list only at-risk requirements.
Evidence	How many evidence records support the requirement: those on its **Evidence** tab, linked to it directly or through its mapped controls and their tests, not counting superseded ones. A red badge counts the expired ones and an amber one the expiring ones. Filter on it for **Expired or expiring** or **No evidence**. Shows "—" without `evidence:read`.
Exceptions	A count badge for open exceptions, suffixed "· expiring" (amber) when one runs out within the reminder window or "· expired" (red) when one is past expiry.
Last reviewed	When the status last changed.

Clicking any cell opens the item's page. The search box matches code, name and description; the status column filters by value. **Export CSV** downloads the current view. **Import** creates compliance exceptions in bulk from CSV (see Importing from CSV in Getting Started).

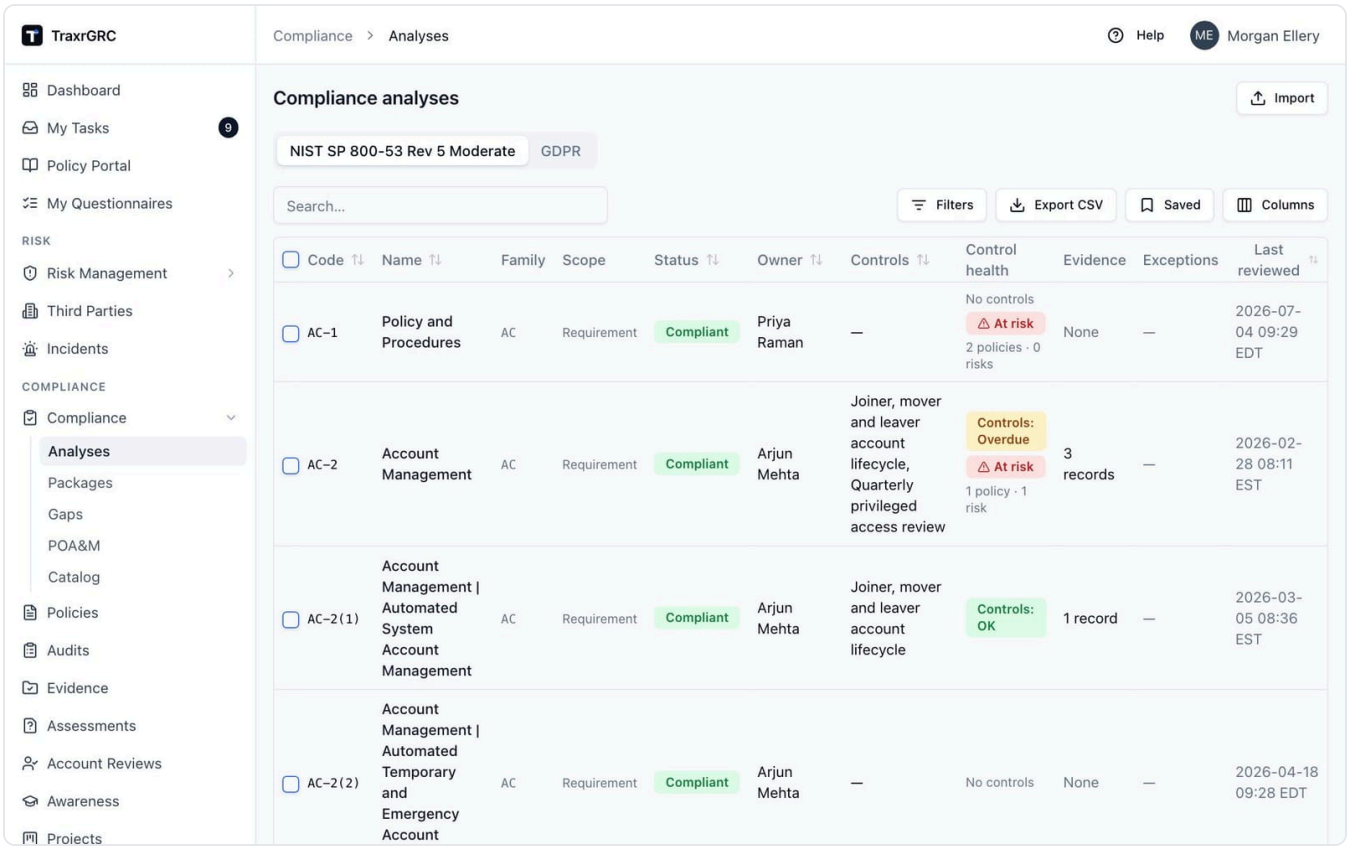


Figure 4.5. The Compliance analyses workspace on the NIST SP 800-53 tab, with status, owner, mapped controls, control health and exception badges.

4.7.1 Filtering by family

In a package with family or domain headings at the top of its tree, the workspace has a **Family** column and filter: pick a family, such as *AC Access Control*, to list only the requirements beneath it. Requirements at the top level of the tree belong to no family and never match.

4.7.2 Bulk updates

With compliance:update each row has a checkbox, and the header has a select-all box. To change many analyses at once:

1. Tick the rows you want. The toolbar shows how many are selected.
2. Choose a value from **Set status...** or **Assign owner...**, or, in a package that inherits from a provider, **Set origination...**. The change is applied as soon as you pick it.

A message reports how many analyses were updated. This is the quickest way to mark a whole domain "Not applicable".

4.8 Analysis statuses

Status	Meaning	Badge color
Not reviewed	No assessment recorded yet (the default).	Gray
Compliant	The clause is met.	Green
Partially compliant	The clause is met in part; gaps remain.	Amber
Not compliant	The clause is not met.	Red
Not applicable	The clause does not apply to your organization.	Blue

Saving an analysis with a *changed* status re-stamps **Last reviewed** to that moment. Edits that leave the status untouched (notes, owner, mappings) do not move the stamp.

A requirement is **At risk** when it is marked compliant or partially compliant but its evidence says otherwise: a mapped control is failed, overdue or retired, or a mapped policy is unpublished or past its review date. The flag is calculated every time it is shown and never changes the status you set; it tells you which claims to re-check.

4.9 Working on an item

An item's page opens with a breadcrumb of the package and every ancestor, then the code and name with the status badge, an **At risk** badge when it applies, a kind badge for headings and additional items, and the standard badge. The clause description follows. When it was imported with its list run together on one line ("a. ...; b. ...: 1. ..."), each item is shown on its own line and indented by level, a., then 1., then (a); the stored text is not changed. A **Sub-items** card lists the item's children with their own status badges; click through to work down the tree.

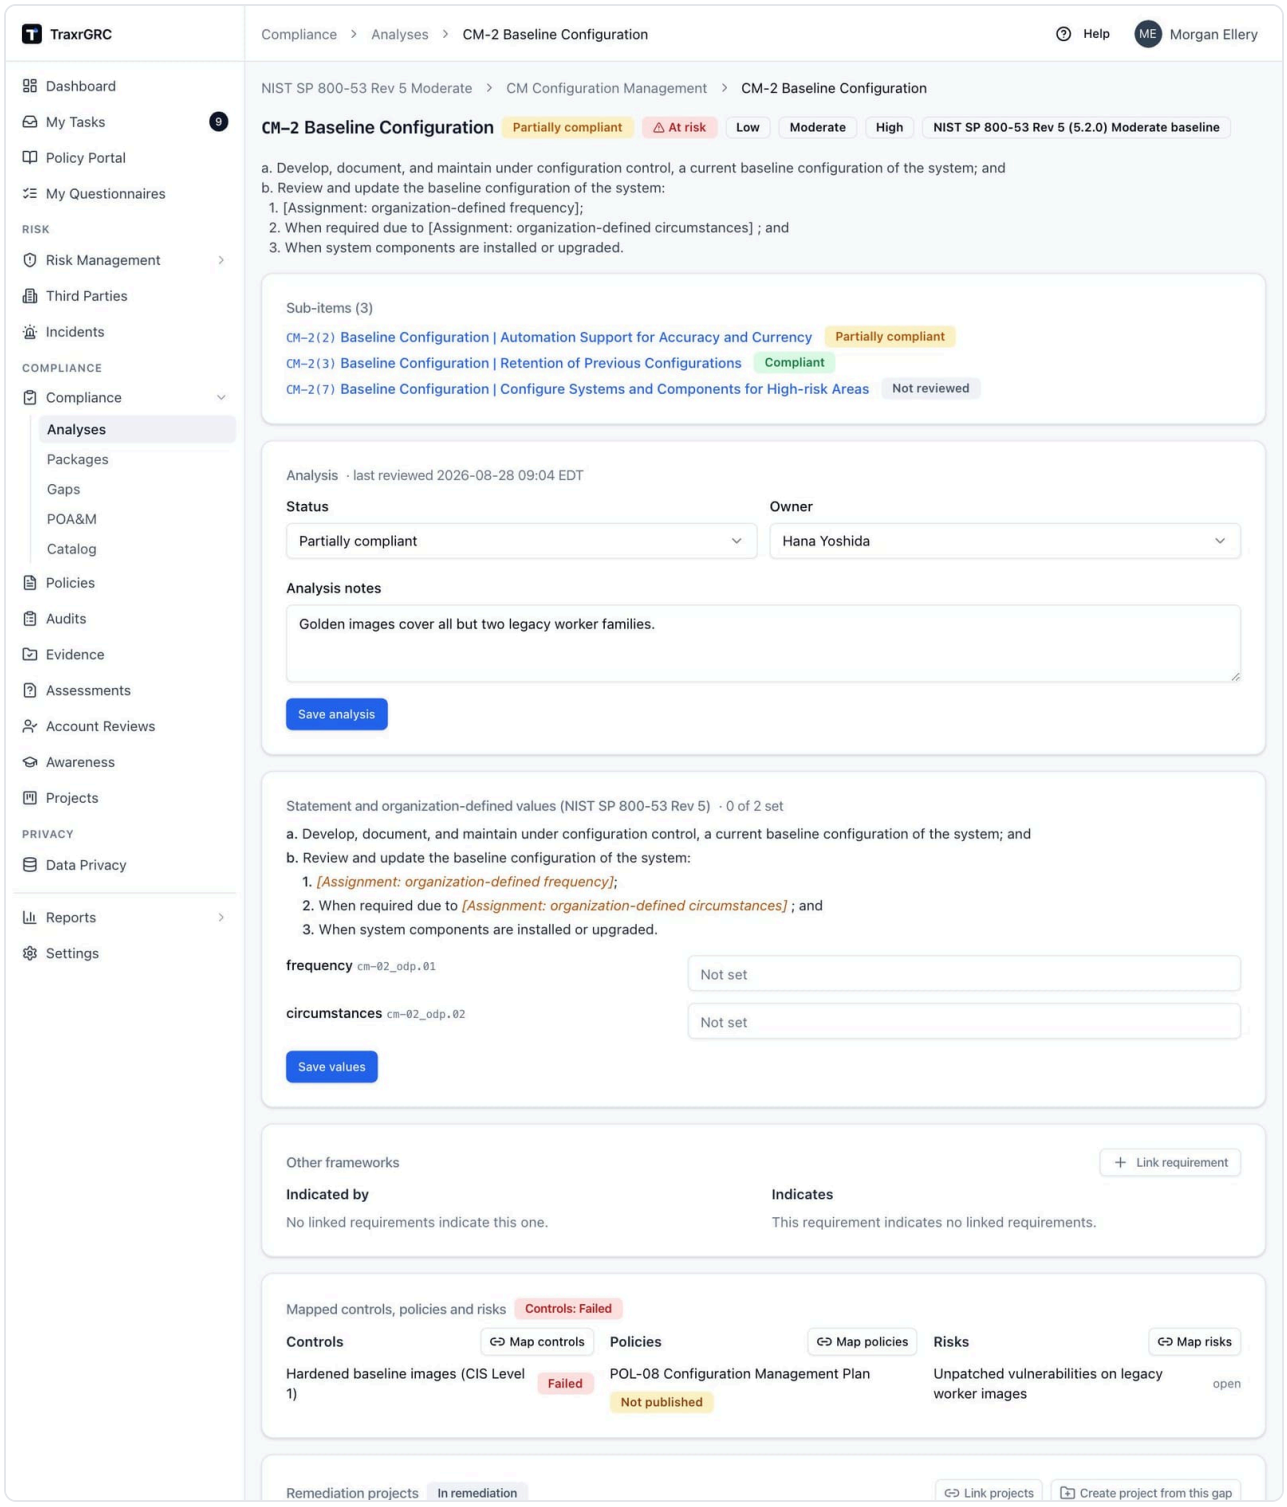


Figure 4.6. A partially compliant requirement at risk: the Analysis card, the mapped controls, policies and risks, and an overdue remediation project.

4.9.1 Recording the analysis

The **Analysis** card shows when the status last changed and holds the assessment. With compliance:update:

1. Pick a **Status** from the five values above.
2. In a package that inherits from a provider, pick an **Origination** (see Common controls and inheritance). For Common and Hybrid requirements the card shows the **Effective status**, how it was reached, and a link to the provider's requirement; the page header shows the effective status.
3. Optionally assign an **Owner**.
4. Write up your reasoning in **Analysis notes**.
5. Click **Save analysis**.

Without compliance:update the fields are shown read-only.

4.9.2 Mapping controls, policies and risks

The **Mapped controls, policies and risks** card links the requirement to the controls that implement it, the policies that mandate it and the risks it addresses. The card header shows the controls' combined health ("Controls: Failed", for example).

1. Click **Map controls**, **Map policies** or **Map risks** (compliance:update, plus read on that module).
2. Search the checklist and tick or untick entries. Each change saves straight away.

Controls show their testing status and policies show whether they are published and in date. Entries in modules you cannot read are counted but not listed, and are left untouched when you change the others. Policy and risk pages have a **Compliance** tab listing the requirements they are mapped to; a control's **Used by** tab does the same. Deleting a policy or risk removes its mappings and records that on each affected requirement's history.

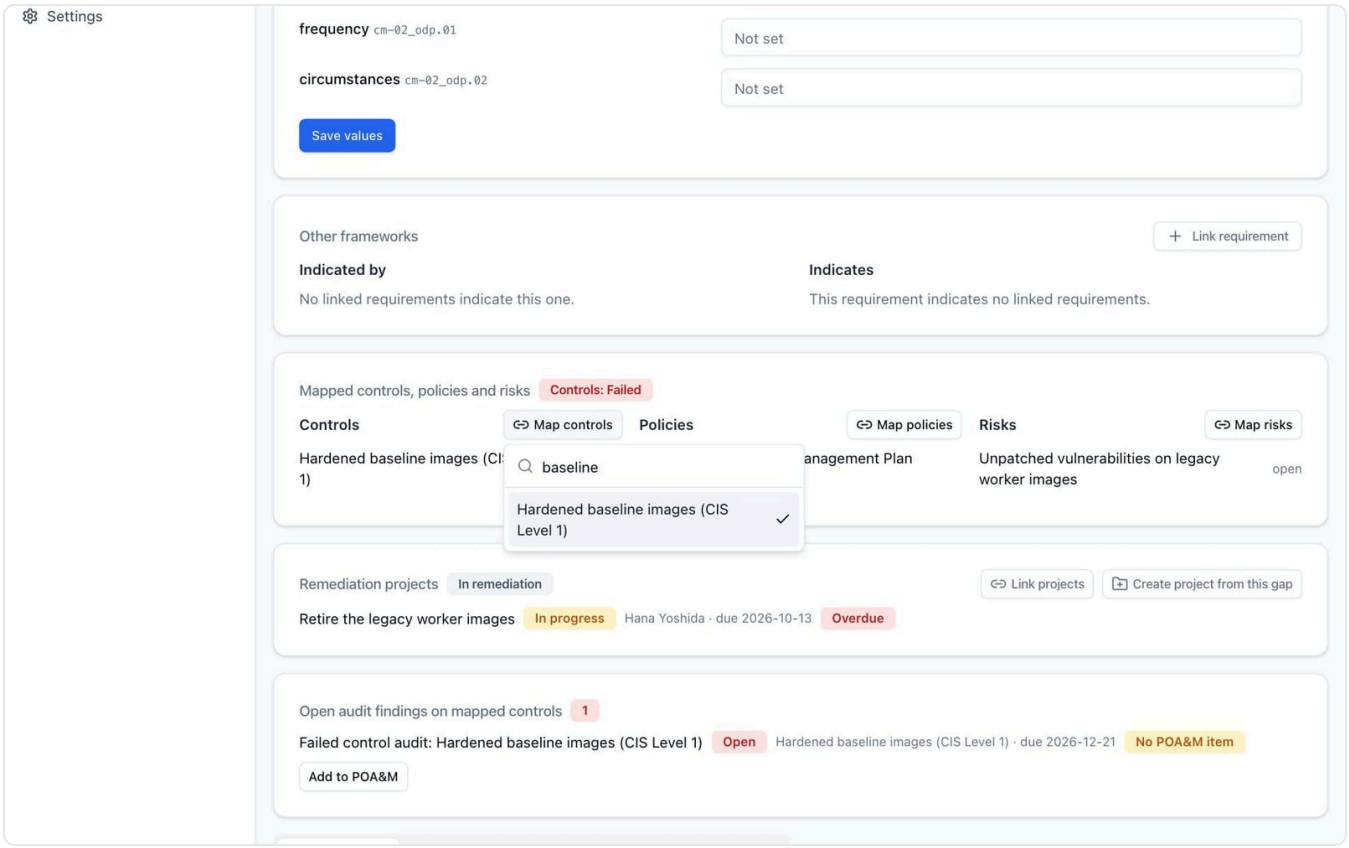


Figure 4.7. The Map controls checklist open on a requirement: a searchable list of every control, where ticking or unticking one saves immediately.

4.9.3 Open audit findings on mapped controls

When an audit finding that is still open names one of the requirement's mapped controls, the card **Open audit findings on mapped controls** lists it: its title (a link to its audit plan), status, control and deadline (red when past), and the POA&M item tracking it, or a **No POA&M item** badge with **Add to POA&M** for users with compliance:create. Findings you cannot read are counted. The card appears only when there is at least one such finding; it is worked out from the current mappings each time the page opens, so mapping or unmapping a control changes it at once.

4.9.4 Remediation projects

The **Remediation projects** card lists the projects that close the requirement's gap, each with its status, owner (or owner group), deadline and an **Overdue** badge when the deadline has passed and the project is not completed. On a partially compliant or not compliant requirement the header shows the coverage:

Coverage	Meaning
No project	No project is linked (red).
In remediation	At least one linked project is still running.
Projects completed	Every linked project is done but the status still reports a gap, so the requirement needs reassessing.

Link projects (compliance:update, plus projects:read) opens a searchable checklist of the projects you can read. Ticking or unticking one saves straight away, and the change is recorded on the history of both the requirement and the project. Projects you cannot read are counted ("N more you do not have access to") but not named, and stay linked when you change the others.

To start remediation from the requirement:

1. On a partially compliant or not compliant requirement, click **Create project from this gap** (requires projects:create).
2. The **New project** dialog opens filled in from the requirement: the title "Remediate *code name*", a description naming the package, the clause and its status followed by your analysis notes, and the requirement's owner.
3. Set the deadline (it becomes the planned completion date in the gap view) and adjust anything else.
4. Click **Create project**. The new project is linked to the requirement.

Linking or creating a project never changes the analysis status.

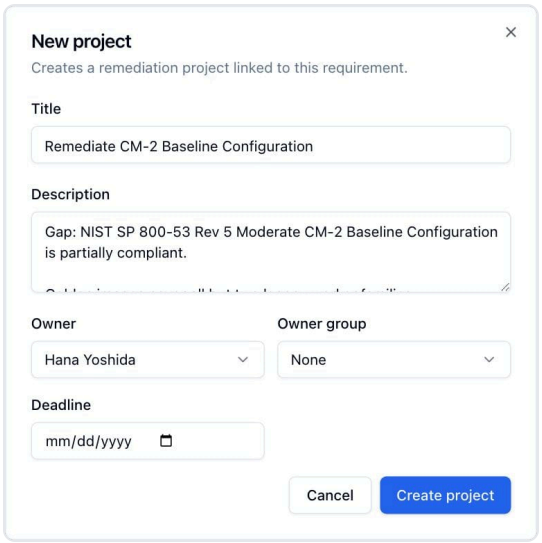


Figure 4.8. The New project dialog opened from a gap, pre-filled with the requirement's code, name, status and analysis notes.

4.9.5 Record tabs

Below the cards, five tabs cover the rest of the record:

Tab	What it holds
Exceptions (N)	Time-boxed deviations from the clause; see Exceptions on a clause.
Evidence	Evidence records supporting the requirement, directly or through its mapped controls and their tests (see Evidence on other records). Adding evidence requires evidence:create.
Comments	Discussion on the analysis. Posting requires compliance:update.
History	Audit entries for the analysis and its exceptions: action, actor, time and changed fields.
Custom Fields	Administrator-defined fields on the analysis (see Administration & Settings). Editing requires compliance:update.

4.10 The gap view (POA&M)

Compliance → Gaps (titled **Compliance gaps**) lists every requirement whose effective status is **Partially compliant** or **Not compliant** across all packages, one row per requirement, like a plan of action and milestones. A requirement inherited from a provider package (common) is a gap when the provider's requirement is; a hybrid one when the worse of the two is. The Status column shows that effective status. Additional items and headings are left out, and with compliance:read at 'own' level you see only analyses you own.

Column	What it shows
Package	The package. Filter on it to see one framework.
Code, Requirement	The clause; the code opens the requirement.
Status	Partially compliant or Not compliant (the severity). Filterable.
Owner	The requirement's owner.
Coverage	**No project**, **In remediation** or **Projects completed** (see Remediation projects). Filter on **No project** for uncovered gaps.
Projects	Each linked project you can read, with its status, owner, deadline and an **Overdue** badge; projects you cannot read are counted ("1 more you do not have access to").
Planned completion	The latest deadline among the linked projects not yet completed.
Overdue	How many linked projects are overdue. Filter on **Overdue projects** to see gaps whose remediation is late.

Coverage, planned completion and the overdue count take every linked project into account, including ones you cannot open, so every viewer sees the same picture of a gap. Without projects:read, linked projects are counted only. **Export CSV** downloads the current view (at most 5,000 rows, with a notice when there are more) with one row per gap and the projects, their owners and statuses joined into single cells. There is no PDF of this view.

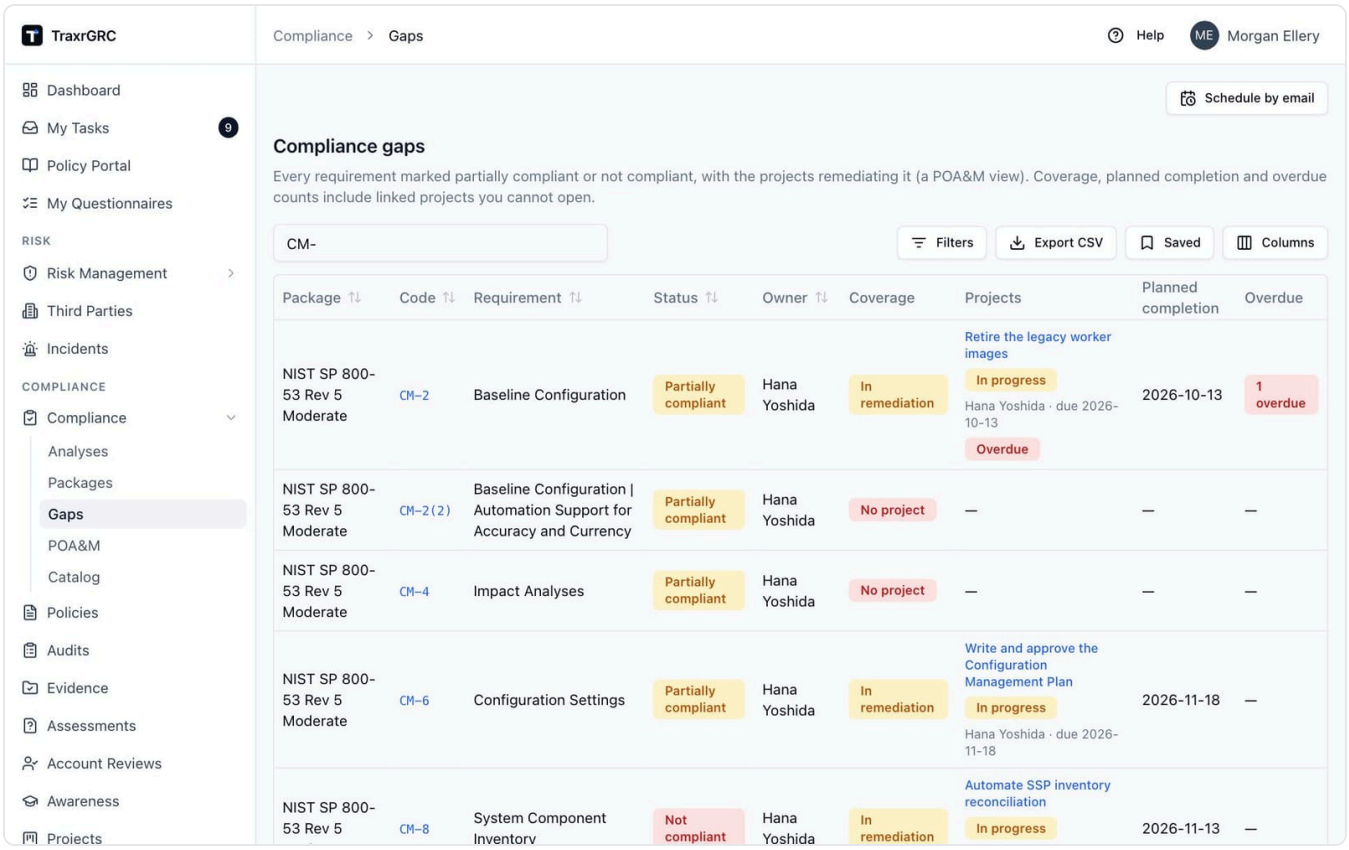


Figure 4.9. The Compliance gaps view with each gap's coverage, linked projects, planned completion and overdue count.

4.11 The POA&M register

Compliance → POA&M is the plan of action and milestones as a register of its own. Each item has a permanent id (POAM-0007) that is never reused, even after the item closes, so references in reports and emails stay valid.

Items come from three sources:

Source	How it opens	How it closes
Requirement status	On its own, when a requirement's own status becomes **Partially compliant** (moderate) or **Not compliant** (high).	On its own, when the status changes to anything else, the item stops being a requirement, or the requirement is deleted. If the gap returns, the same item reopens under its number.
Audit finding	**Add to POA&M** on the finding's row of its audit plan. An open finding moves to **In remediation**.	By a person, with a note, or on its own when whoever runs the audit plan closes the finding.
Entered	**New item**, for a weakness no requirement or finding records: a scan, a penetration test, one you found yourself.	By a person, with a note.

If a requirement's status worsens while its item is open (partially to not compliant), the item rises to high and its due date moves to the high window; a severity a person raised is never lowered.

A requirement set to **Common** opens no item: the weakness is the provider's, and the provider's own requirement carries the item. A **Hybrid** requirement opens one for its own part. Additional items and headings open none.

Due dates. An item is due a number of days after it opens, set per severity: 30 (high), 90 (moderate) and 180 (low) by default. A GRC Manager, or anyone whose role holds compliance:review at 'all' level (typically the CISO), changes them with **Due dates** on the register. New values apply to items opened, reopened or given a new severity from then on; existing due dates stay. On an item, a new severity moves its due date to that severity's window from the opening date, and **Due** can also be set directly.

On an item's page you edit the weakness and milestones text, severity and due date. A finding or entered item also takes a title, an owner and one remediation project, and closes with a note saying how the weakness was resolved; it reopens under the same id. A requirement-status item takes its title, owner and remediation projects from its requirement and cannot be closed by hand. Every change, including the automatic ones, is in the item's **History** and the audit log.

Milestones. Each item carries dated milestones (**Milestones** on its page: a description and a due date). Marking one **Done** records the day it was completed; a completed milestone stays on the plan and cannot be removed, reworded or moved (**Not done** undoes a completion made by mistake). Moving a milestone's date keeps the date it was first planned for, shown beside it, so the plan's changes stay visible. An item holds up to 50 milestones, and a closed item's plan is final until the item reopens. Every milestone change is in the item's History.

FedRAMP fields. An item also records **Controls** (such as "RA-5, SI-2"), **Detected by** (the scanner, test or audit that found it) and **Assets** (the hosts, components or services affected). A requirement-status

item's control is its requirement code and its source is the control assessment; its assets are yours to fill in. **Vendor dependency** says whether the fix waits on a vendor (Not recorded, No, or Yes). Yes asks for the vendor product the fix waits on and takes the date of the last vendor check-in (today or earlier); switching to No clears both.

FedRAMP export. FedRAMP export on the register downloads an .xlsx in the layout of FedRAMP's POA&M template (R3.0, December 2025): the **Open POA&M Items** sheet with the provider, offering, impact level and POA&M date at the top and the column headers on row 5, **Closed POA&M Items**, and an empty **Configuration Findings** sheet. The dialog prefills the provider with your organization's name and takes an optional package (requirement-status items of that package only). R3.0 has no separate milestone columns, so each item's milestones are listed in its **Remediation Plan** cell, as the template's guidance asks, and a moved milestone is noted in **Comments** with a closed item's closure note. An audit-finding item's detection date is the day the finding was recorded. Vendor Dependency, Vendor Dependent Product Name and Last Vendor Check-in Date come from the item's vendor dependency; an item where it is **Not recorded** leaves Vendor Dependency blank (the template makes it mandatory, so record it before you submit). The situational columns (risk adjustment, false positive, operational requirement, BOD 22-01, CVE) are left for you to fill in. Only items you can read are exported.

One weakness, one item. When a finding describes a weakness an open item already tracks, choose **Attach to an open item** in the Add to POA&M dialog instead of opening a second one. The audit plan's findings table shows each finding's POA&M id, marked "(closed)" once the item closes.

Findings and their items move together. Adding a finding to the POA&M (or attaching it to an open item) moves an open finding to **In remediation** and notes the item on the finding's comments. Closing the item notes "ready to verify and close the finding" there; the finding itself stays open until whoever runs the audit plan closes it, which verifies the fix. Closing the finding closes its open item ("The audit finding was verified and closed."), and reopening the finding reopens the item under its number with a new due date. While an open finding item tracks a weakness on a control, a requirement that control is mapped to opens no second, status-derived item for the same weakness.

Findings on mapped controls. When the finding names a control that is mapped to requirements, the dialog says which ("Its control is mapped to AC-2 (NIST SP 800-53 Rev 5), ...") and fills **Controls** with their codes and **Detected by** with the source (*Control audit: <control>* for a finding raised by a failed control audit, otherwise *Audit: <plan>*); change either before adding. The item still belongs to no package, like every finding item. While such a finding has no item, the POA&M page shows a notice ("2 open audit findings affect mapped requirements and have no POA&M item") with **Review them**, which opens Open findings filtered to those findings.

The register lists, filters (source, severity, status, package, overdue) and exports CSV like the gap view; with compliance:read at 'own' level you see items you or your groups own. Linked projects you cannot read are counted, and a finding you cannot read shows as "A finding you cannot open" (a finding needs audits:read on both the finding and its audit plan, as in CSV import). The register is also `GET /api/v1/compliance/poam` (see the API chapter) and, for AI assistants, the read-only `list_poam` and `get_poam_item` tools. The package PDF report lists the package's open requirement-status items with their milestones.

4.12 Exceptions on a clause

An exception documents a time-boxed deviation from a clause, for example accepting a known gap until a control is deployed.

1. On the item's **Exceptions** tab, click **Request exception** (requires compliance:update).
2. Complete the dialog:

Field	What it means
Title	Short name of the exception (required).
Description	Why the deviation is needed (required).
Owner	Optional user responsible for the exception.
Approver	Optional user who must decide it. You cannot name yourself.
Expires at	The date the exception runs out (required).

3. Click **Request exception**.

New exceptions are **requested** (amber) and are decided with **Approve** (green) or **Reject** (red). Exceptions follow a separation of duties:

- The person who requested an exception can never approve or reject it, administrators included, unless an administrator has relaxed separation of duties (see Sessions under Administration & Settings). A decision made that way carries an amber **Self-approved** or **Self-decided** badge.
- When the request names an approver, only that person (who must be able to see the record) or an administrator can decide it.
- When it names none, deciding needs the **Review** right on Compliance (compliance:review).

The exception card shows the owner, the designated approver, who requested it and who decided it; the **Approve** and **Reject** buttons appear only for someone allowed to decide. The decision is recorded with who made it and when, and is final. Once the expiry date has passed, an open exception displays as **expired** (gray). The notification worker sends "Compliance exception expiring" reminders before open exceptions run out, and expiring and expired exceptions color the exceptions badge in the analysis workspace.

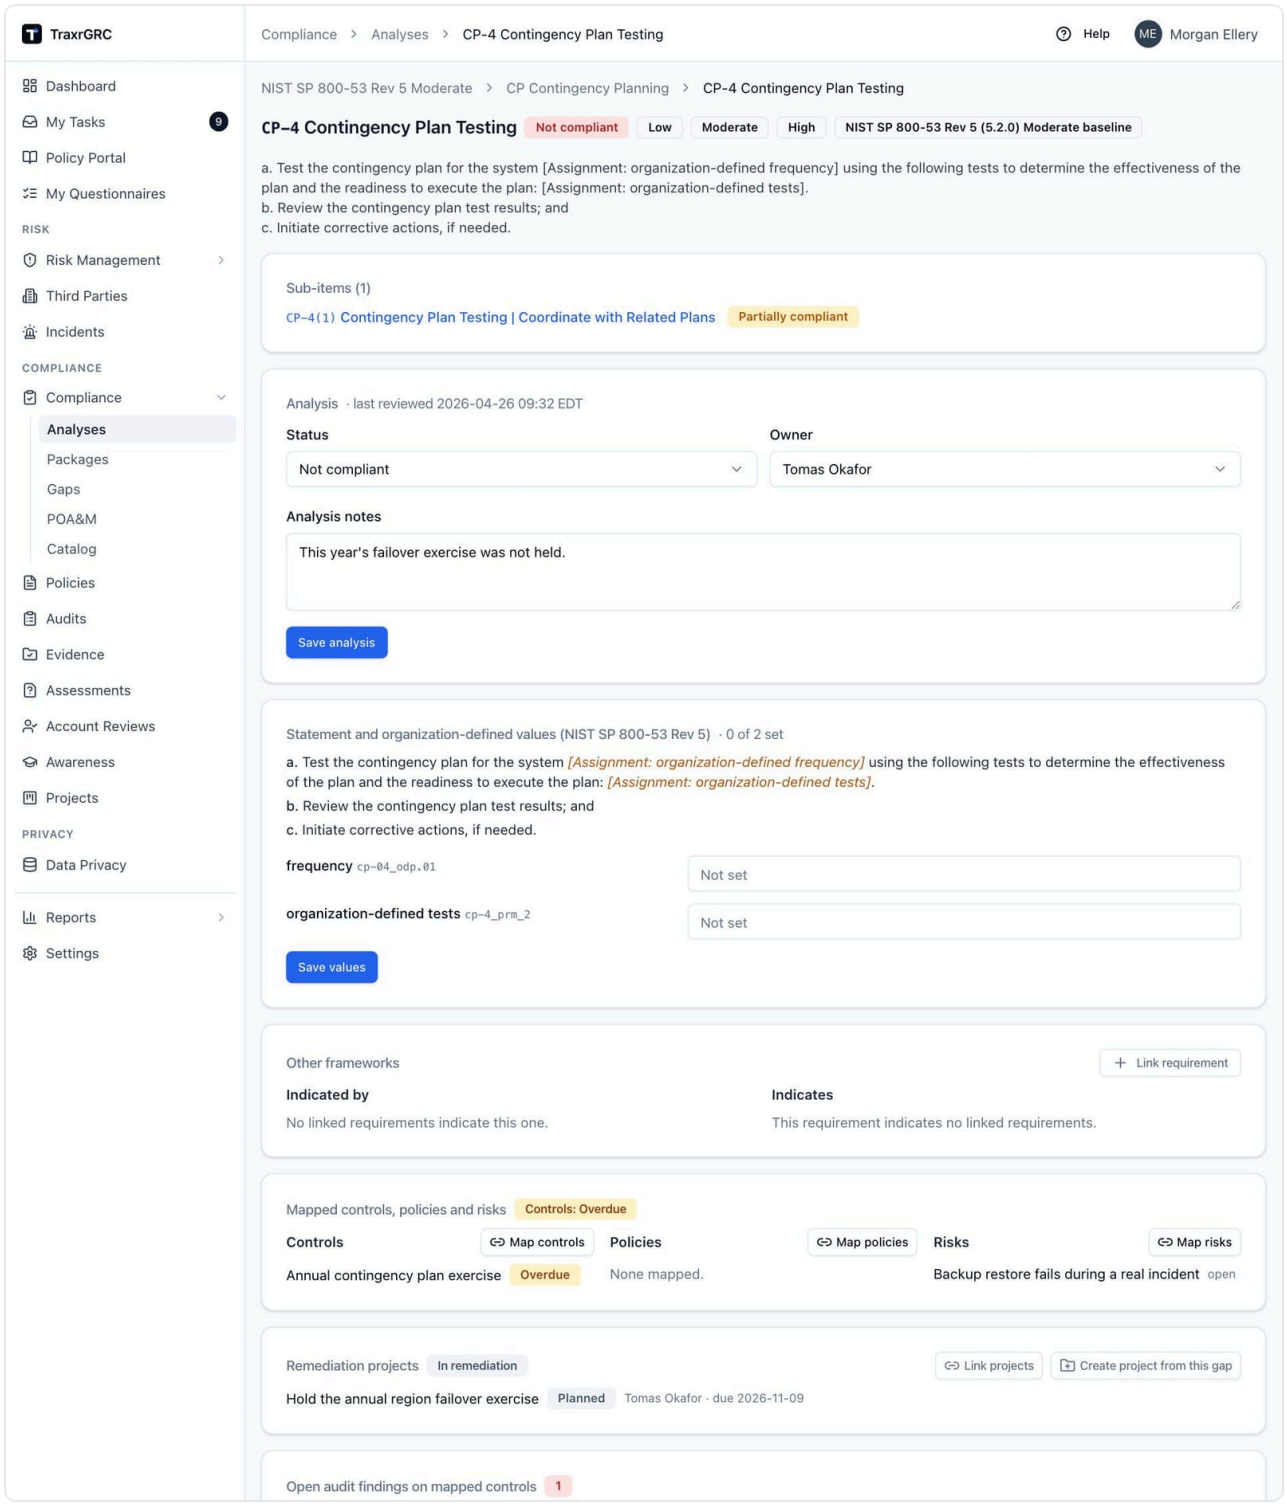


Figure 4.10. A requirement's Exceptions tab with a requested exception awaiting a decision.

4.13 The compliance report (PDF)

Click **Report** on a package page to download a PDF for that package. The button appears when you hold compliance:read at 'all' level. The report contains:

- A cover header with the organization name, package name, standard and version, and the generation time.
- An **Executive summary** of requirements only: the number of requirements, a count for each of the five statuses, a headline **Compliant (excl. N/A)** percentage (the share of applicable requirements that are compliant), and the number of **At risk** requirements. Headings and additional items are never counted here.
- One section per top-level domain, each with a clause table (clause, name, color-coded status, owner, exception count) covering that domain's whole subtree. At-risk clauses carry an **At risk** marker under their status. Headings appear only as section titles.
- **Additional implemented items**, when the package has any: their own counts, percentage, at-risk count and clause table.
- For a package linked to a catalog, **Organization-defined values**: how many of the package's parameters are set, then each requirement with parameters and every parameter's value, "Not set", or the provider's value marked Inherited.
- An appendix of **open exceptions** (requested or approved, with past-expiry ones shown as expired), sorted by expiry date.

PDFs (this report, policy downloads and the Article 30 register) are generated a couple at a time. When several people generate one at the same moment, the rest are told to try again in a few seconds.

4.14 Deleting packages and items

Deleting a **package** (compliance:delete) asks for confirmation ("Delete package?"), then removes it together with all of its items, analyses, exceptions, comments, attachments (files included), label assignments and custom-field values. It also removes its requirements' links to projects (the projects stay) and records that on each project's history, and its requirements' crosswalk links. It is refused while another package inherits common controls from it. This cannot be undone.

Deleting an **item** (compliance:delete) removes the item, its analysis and its whole subtree of sub-items with their analyses. Deletion is blocked while the item's analysis has mapped controls, policies, risks or projects (unmap them first) or any exceptions (remove them first); hover over the disabled bin icon to see why.

4.15 Who can do what (Compliance Management)

Action	Permission
See packages, the analysis workspace and item pages	compliance:read (at 'own' level you see only analyses you or your groups own)
Download the PDF report	compliance:read at 'all' level
Create packages, import, add items	compliance:create
Edit packages and items, reorder items, set statuses, bulk updates, map controls/policies/risks and link projects (also needs read on that module and on each newly linked record), request exceptions, comment, upload evidence, edit custom fields	compliance:update: 'all' for every record, 'own' only for analyses you or your groups own
See the gap view	compliance:read (projects are named only with projects:read on them)
See the POA&M register and its items	compliance:read (at 'own' level, items you or your groups own)
Add a finding or entered POA&M item	compliance:create, plus audits:read on the finding and its plan, and projects:read on the project
Edit, close, reopen a POA&M item, attach a finding, add, move, complete or remove its milestones	compliance:update on the item
Download the FedRAMP export	compliance:read (items in your read scope)
Set POA&M due dates by severity	compliance:review at 'all' level
Create a project from a gap	projects:create and compliance:read on the requirement
Decide an exception that names no approver (never one you requested)	compliance:review

Action	Permission
Decide an exception that names you as approver (never one you requested)	compliance:read on the analysis
Set a package's provider, import crosswalks, link and unlink requirements	compliance:update at 'all'
Set an analysis's origination or its organization-defined values	compliance:update on the analysis
Link a package to the catalog	compliance:update at 'all'
Add catalog controls to a package	compliance:create and compliance:update at 'all' (it reorders the package)
Download the parameter values CSV	compliance:read at 'all'
Delete packages and items	compliance:delete

Administrators bypass all of the above.

CHAPTER 5

Policy Management

Policies are the written rules your organization operates under: acceptable use, security standards, HR policy and anything else that needs a formal owner, a review cycle and proof of approval. TraxrGRC manages the full lifecycle of each policy: drafting in a built-in editor, routing changes through approval, publishing immutable version snapshots, scheduling periodic reviews, tracking time-boxed exceptions, and publishing policies to a portal where staff read and acknowledge them.

Policy owners, the GRC team and approvers work in the **Policies** area (sidebar, under **Compliance**). Every signed-in user, whatever their role, reads published policies on the **Policy Portal** (sidebar, under **Dashboard**), described at the end of this chapter.

5.1 The policies list

Open **Policies** from the navigation. The list shows every policy you have access to, with one row per policy:

Column	What it shows
Title	The policy name. Click it to open the policy.
Status	Draft, Published or Archived, as a colored badge (see table below).
Version	The latest published version (for example `v2`), or a dash if nothing has been published yet.
Owner	The owning user, or the owning group if no user is set.
Next review	The date the next review is due. If the date has passed and the policy is not archived, the date is shown in red with "(overdue)".
Exceptions	A badge with the number of exceptions, not counting rejected ones. The badge is red when any exception has expired, amber when one expires within the next 14 days, and gray otherwise. A dash means there are no exceptions.
Labels	The policy's labels. Filter by typing part of a label name.

Policy statuses:

Status	Meaning	Badge color
Draft	Working copy: not yet published, or changed since the last publish.	Gray
Published	A version snapshot is published and current.	Green
Archived	Retired from active use; excluded from review-due reminders.	Gray

Use the filter box to search titles and descriptions, and the Status column filter to narrow the list. You can sort by title, status, next review date and owner. Policies and policy exceptions can also be created in bulk with **Import** (see Importing from CSV in Getting Started).

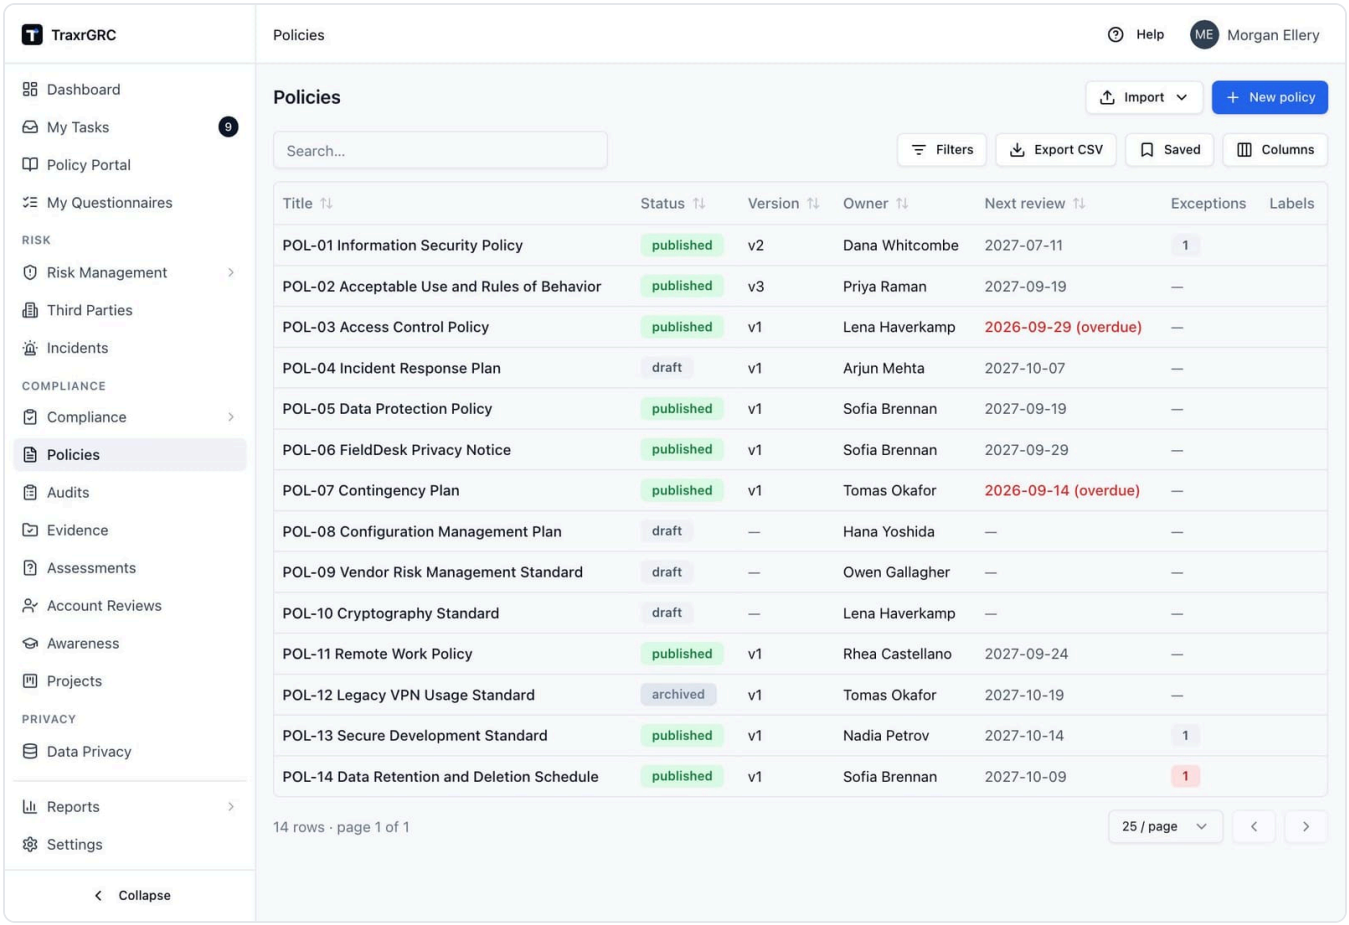


Figure 5.1. The policies list, showing published, draft and archived policies, overdue review dates in red and exception count badges.

5.2 Creating a policy

1. Click **New policy**.
2. Complete the dialog:

Field	What it means
Title	The policy name (required).
Description	Optional background, shown on the detail page.
Owner	The user responsible for the policy. Both owner fields are optional.
Owner group	The group responsible for the policy.
Review frequency (months)	How often the policy must be reviewed: 1 to 120 months; defaults to 12.

3. Click **Create policy**.

The new policy starts as a draft with empty content. You add the content after creation, from the policy's Content tab.

5.3 The policy detail page

Click a policy title in the list to open its detail page.

Header. The title, the status badge and, when applicable, an amber **Unpublished changes** badge with a banner explaining it. On the right are the action buttons: **Publish vN** (only when the next version is approved), **Archive** or **Reactivate**, **Edit**, and **Delete** in the **More actions** menu (**...**). Each appears only if you have the right to use it.

Labels. Below the description sits the **Labels** picker (see Labels in Getting Started; attaching needs policies:update on the policy).

Summary cards. Below the description, five cards show:

Card	What it shows
Owner	The owning user, with the owner group underneath.
Latest version	The latest published version and when it was published.
Effective	The effective date of the published version and, for an approval recorded outside the app, its signatories.
Next review	The next review due date, with "overdue" in red when it has passed, and the review frequency.
Review frequency	How often the policy is reviewed, in months.

Tabs. One row of tabs covers the rest of the page:

- **Content**, **Versions**, **Approvals**, **Reviews**, **Exceptions**, **Comments**, **Attachments** and **History**: the policy itself. The Versions, Approvals and Exceptions tabs show a count in brackets.
- **Compliance (N)**: the compliance requirements this policy is mapped to, with the package and the analysis status of each. Shown only if you can read the Compliance module. Mappings are made from the requirement's page (see Compliance Management).
- **Portal & acknowledgements**: putting the policy on the Policy Portal and asking people to acknowledge it.
- **Custom Fields**: administrator-defined extra fields for policies.

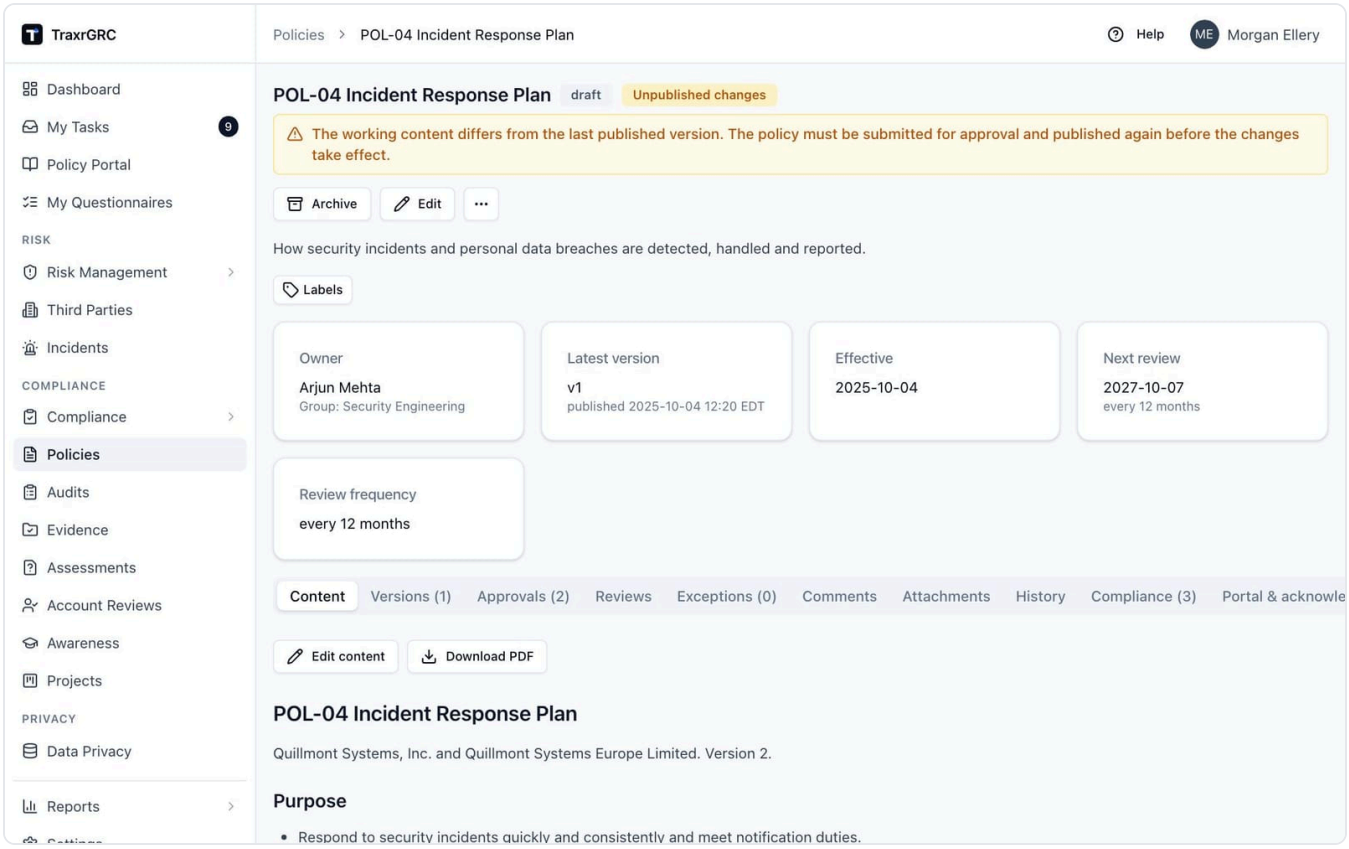


Figure 5.2. A policy with a published version and unpublished changes: the amber badge and banner, the summary cards and the Content tab.

5.4 Writing content (Content tab)

The Content tab is where the policy text lives.

- If the policy has content, it is shown as a formatted document (headings, lists, tables, bold, italic and code styling are rendered). Otherwise the tab reads "No content yet. Use Edit content to write the policy."
- Click **Edit content** to open the editor: a **Markdown source** pane on the left and a live **Preview** on the right. Markdown support covers headings, paragraphs, lists, tables, `**bold**`, `*italic*`

and `` `code` ``.
- The toolbar above the source pane formats the selected text: **Heading**, **Bold**, **Italic**, **Code**, **Bulleted list** and **Numbered list** (click again to remove the formatting), and **Table**, which inserts a two-column table to fill in. Ctrl+B and Ctrl+I (⌘B and ⌘I on a Mac) work in the source pane too.
- Click **Save content** to store the working copy.

A table is written as rows of cells separated by `|`, with a separator row of dashes under the header row; colons in the separator align a column (`:---` left, `:---:` centered, `---:` right). Leave a blank line before and after it:

```
| Role | Approves | Reviews |
| --- | --- | --- |
| CISO | Policies | Exceptions |
| DPO | DPIAs | Records of processing |
```

Saving changed content on a published policy puts it back into Draft (see Unpublished changes below) and withdraws any approval for a version not yet published.

Download PDF produces a PDF of the content: for a published policy it renders the latest published snapshot, for a draft the current working content labeled "draft". The PDF is A4 with the title and a "Version … · Generated …" line (in your time zone) at the top, and a footer on every page carrying the title, version and page number; the file name combines the policy title and version. Tables print with their header row repeated on each page they span.



Figure 5.3. The content editor, with the Markdown source on the left and the live preview on the right.

5.5 The draft → approval → publish workflow

Publishing is gated so that no policy version goes live without a recorded approval. Versions are numbered sequentially: the first publish creates version 1, and each later publish is one higher than the highest published snapshot.

An approval covers the exact text that was submitted. If the content changes after a request is submitted or approved, and that version has not been published yet, the approval is withdrawn

automatically: it shows as Rejected with the note "Withdrawn automatically: the policy content changed after submission." Submit the new text for approval again.

5.5.1 Submitting for approval

1. Open the policy and go to the **Approvals** tab.
2. Click **Submit version N for approval**, where N is the next version number (highest published + 1).
3. An approval request appears as a card, "Version N awaiting approval", showing who requested it and when.

Only one request can be pending at a time: the button is disabled while a request is awaiting a decision, and the server also refuses a new request while one is open. If the next version is already approved, publish it instead of submitting again. Archived policies cannot be submitted for approval.

5.5.2 Approving or rejecting

1. On the Approvals tab, find the pending card and click **Approve** or **Reject**.
2. In the dialog ("Approve version N" or "Reject version N"), optionally add a **Note**. For an approval, this note becomes the change note stored on the published version snapshot.
3. Click **Approve** or **Reject** in the dialog. The decision is recorded with your name and the time, and the request's status changes to Approved (green) or Rejected (red).

You cannot approve a request you submitted yourself; another user with permission must make that decision, unless an administrator has relaxed separation of duties (see Sessions under Settings). An approval made that way carries an amber **Self-approved** badge so reviewers can see the control was waived. The requester may still reject their own request. An approval is also refused if the content changed after the request was submitted; reject it and submit again. When the real approval happens outside the app, record it instead (see below). Deciding an approval writes an entry to the audit log.

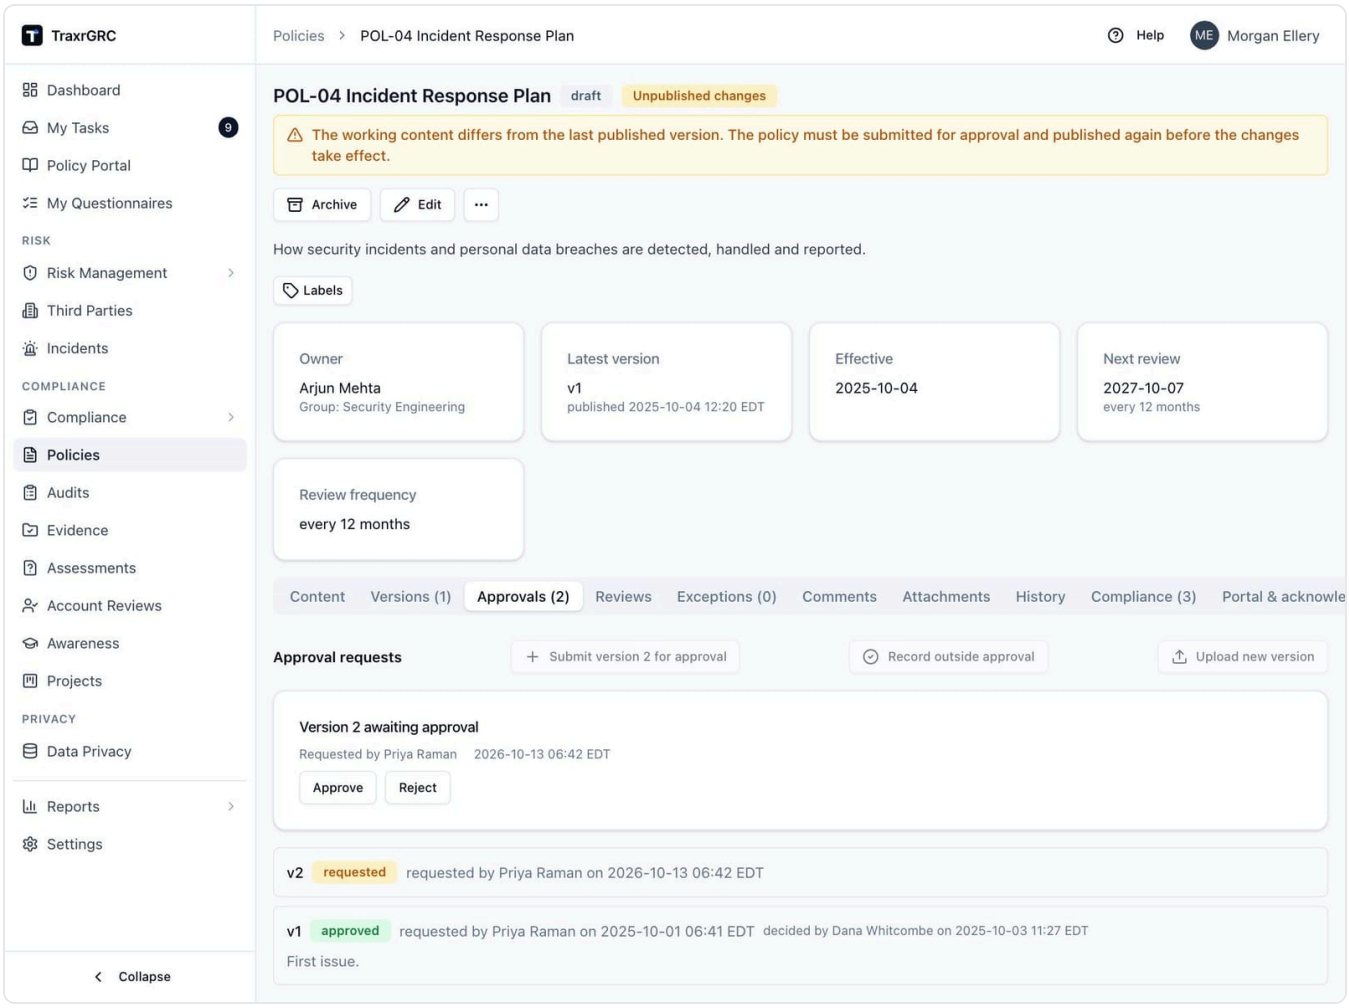


Figure 5.4. The Approvals tab with a version awaiting approval, its Approve and Reject buttons, and earlier approvals listed below.

5.5.3 Publishing

1. Once a request for the exact next version is approved, click **Publish vN** in the page header.

Publishing checks all of the following and refuses with an explanatory message otherwise:

- the policy is not archived,
- the content is not empty,
- no approval request for the next version is still pending,
- an **approved** approval exists for **exactly that version** (an approval for an older version does not count), and
- the content is the same text that was approved.

On success, TraxrGRC creates an immutable version snapshot of the current working content (with the approval note as its change note, and you as the author), sets the policy status to Published, and records the publish date and effective date. Previous approvals stay visible in the Approvals tab as history.

5.5.4 Recording an approval made outside the app

Many organizations approve policies outside any tool: a board signs them, or a controlled-document process collects signatures. Record that approval instead of re-approving the policy in the app. Because it publishes with no second person approving in the app, it is off until an administrator turns on **Settings → Policies → Allow recording outside approvals**, and it needs both update and **Review** on Policies.

1. On the Approvals tab, click **Record outside approval**.
2. Enter the **Effective date**.
3. Under **Signatories**, enter each signatory as the document states them: the capacity they signed in (for example *Board of Directors Approval*), their name, and optionally their title. Click **Add signatory** for each further signatory.
4. Add a **Reference** if there is one, such as board minutes.
5. Click **Record and publish**. The next version is published at once.

The record names the signatories as the approvers and you as the person who recorded it. In the Approvals tab it reads "approved outside the app, effective …", followed by the signatories, the reference and "Recorded by …". The second-person rule for in-app approvals does not apply. The policy's **Effective** card shows the date and signatories of the published version, and each version keeps its own. The button is unavailable while an in-app approval request is pending; reject that request first.

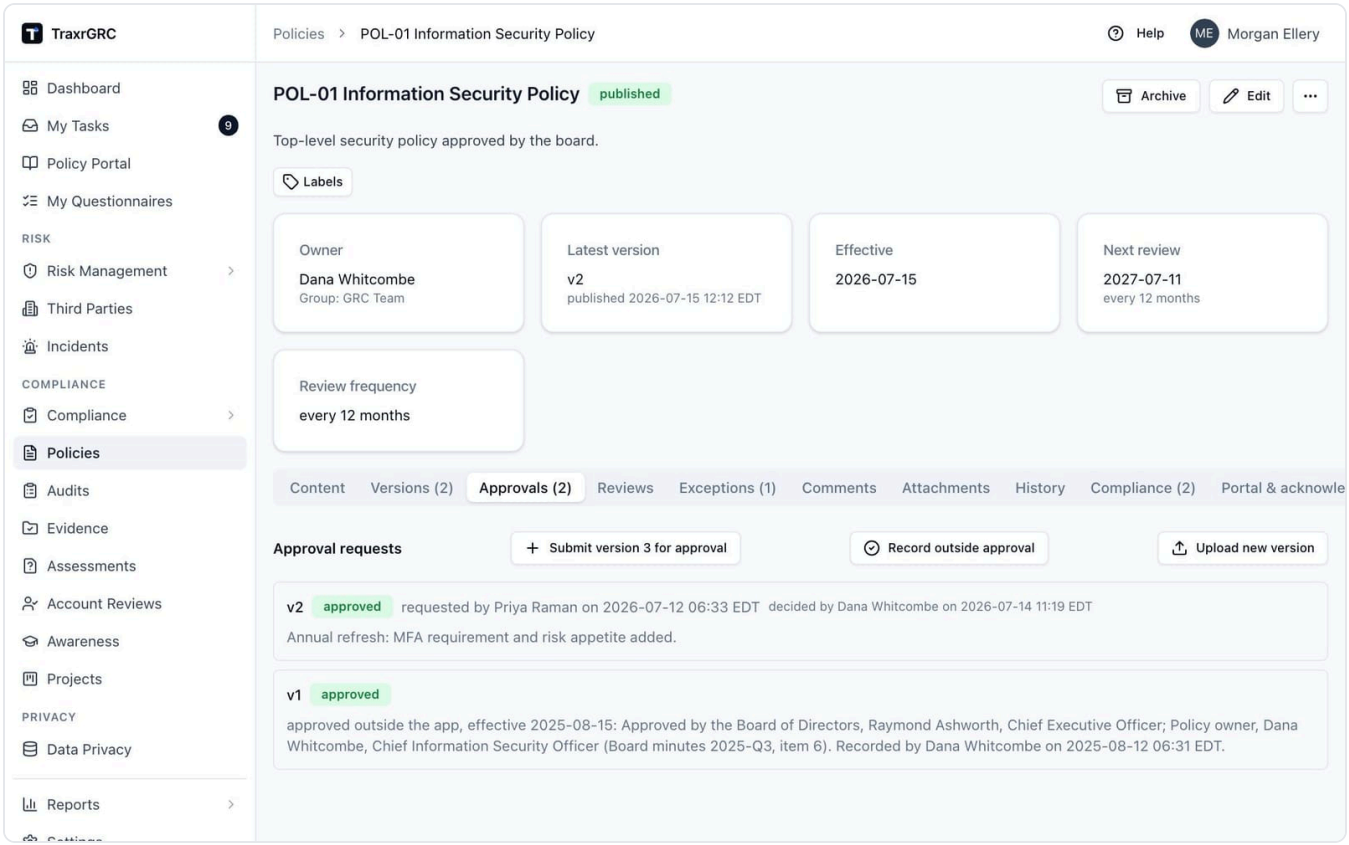


Figure 5.5. The Approvals tab of a policy whose first version was approved by the board outside the app and whose second version was approved in the app.

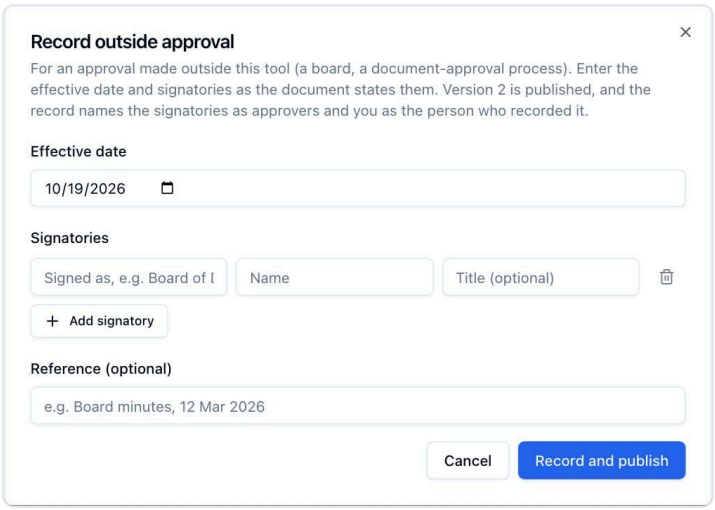


Figure 5.6. The Record outside approval dialog, with the effective date, signatory rows and reference.

5.5.5 Uploading a reissued document

When a controlled document is reissued, upload it as the next version instead of pasting its text. This is an outside approval too, so it needs the same setting and rights; without them the button and the dashboard upload link do not appear.

1. On the **Approvals** tab, click **Upload new version**. Alternatively, use the upload button beside the policy on the **Out-of-date documents** dashboard widget, which opens the same dialog.
2. Under **Document**, choose the file: `.docx`, `.pdf`, `.md` or `.txt`, up to 25 MB. A Word document's body text is read paragraph by paragraph, and a PDF's text page by page, so a PDF's running headers and footers appear on each page; formatting and images are not kept in the content. A scanned PDF has no text to read and is refused; upload the Word document or a PDF saved from it.
3. Enter the **Effective date**. The signatories start as those of the current version; change them if the reissue was signed by others. Add a **Reference** if there is one.
4. Click **Upload and publish**.

The document's text replaces the working content, the file is attached to the policy, and the next version is published as an outside approval (see above), all in one step: if any part fails, nothing changes and the file is not kept. In the Versions tab that version gets a **Source document** button that downloads the file as uploaded. A file of another type, one with no text, or one whose text is longer than the 200,000 characters a policy holds, is refused and nothing changes.

5.5.6 Unpublished changes

Editing the content of a published policy, or restoring an older version into the working copy, moves the policy back to **Draft**. The header then shows an amber **Unpublished changes** badge and a banner explaining that the working content differs from the last published version and must be submitted for approval and published again before it takes effect. The published snapshot itself is not modified; the Policy Portal and readers of the current version keep seeing it until the new version is published.

5.6 Versions tab

Every published snapshot is listed here, newest first, with its version number, change note, author and date; the newest carries a "latest" badge. Each version offers three icon buttons:

- **View content**: a read-only dialog showing that snapshot's content.
- **Restore into working draft**: after you confirm, copies the snapshot's content into the working copy and puts the policy into Draft (an archived policy stays archived). The restored change must be approved and published before it takes effect.
- **Download PDF**: a PDF of that exact version, styled like the Content-tab PDF, with the version number in the header and footer.

A version published by upload also has **Source document**, which downloads the uploaded file.

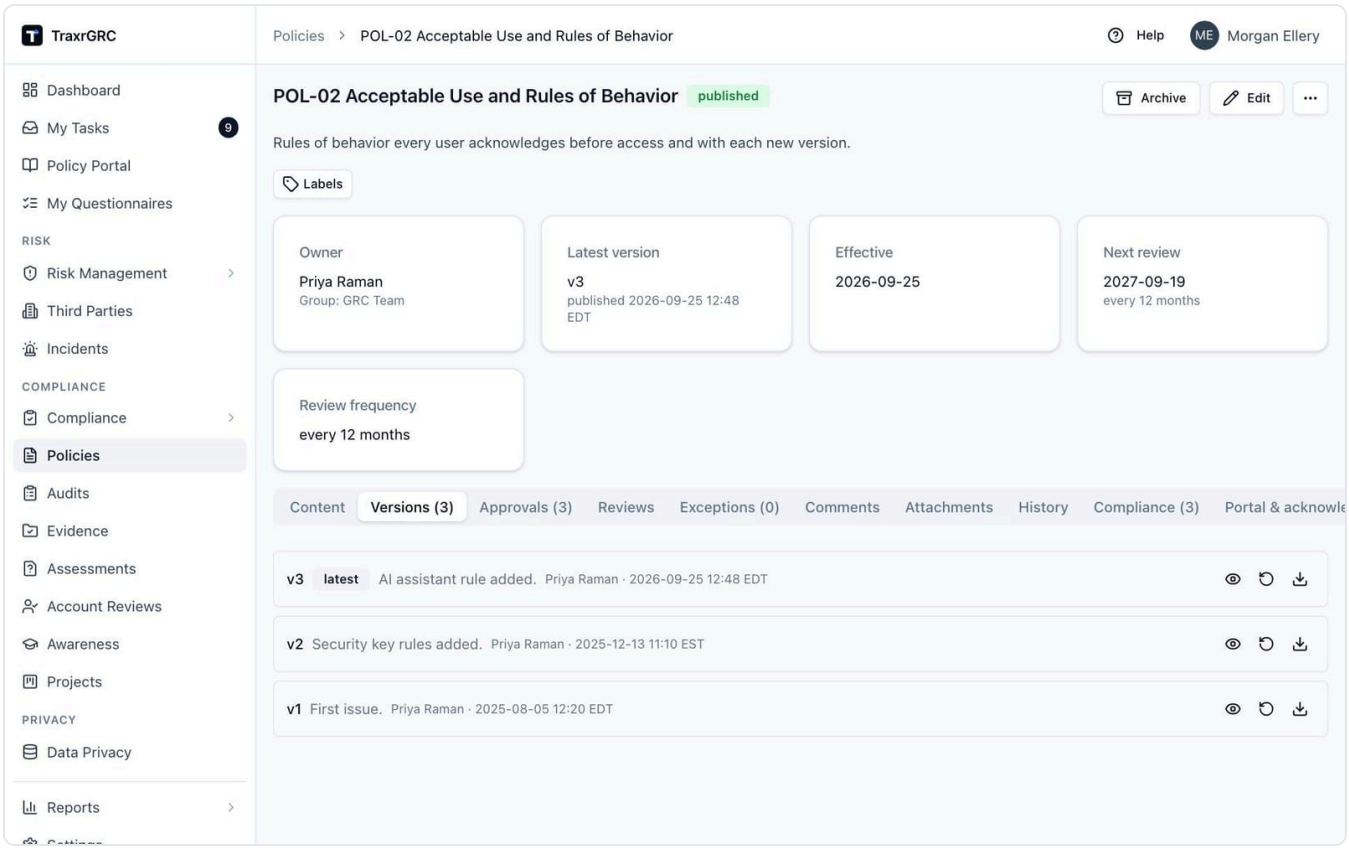


Figure 5.7. The Versions tab of a policy with three published versions, each with its change note and the view, restore and download buttons.

5.7 Reviews tab

Reviews are the recurring check that a policy is still accurate and being followed. The tab has two sections: **Upcoming reviews** and **History**.

Completing a review:

1. Under **Upcoming reviews**, click **Complete review** on the due entry. Past-due dates are shown in red.
2. In the dialog, choose the **Result** (Approved or Changes required) and add optional **Notes**.
3. Click **Complete review**.

Publishing a policy schedules its first review one review interval after the version takes effect, unless a review is already scheduled, so every published policy has a Next review date. A review is overdue once that date has passed on a published policy; drafts and archived policies are never overdue. The policy list, the dashboard, My Tasks, compliance evidence and the AI assistant all use this rule.

The review is stamped with the completion time, the result and your name, and moves into the **History** section. Completing a review automatically schedules the next one at the policy's review frequency (counted from the completion date) and updates the Next review date on the policy and in the list. Completing a review also writes an audit log entry.

Result	Meaning	Badge color
Approved	The policy is still accurate.	Green
Changes required	The review found changes to make; edit the content and take it through approval.	Amber

Ad-hoc reviews:

1. Click **Add review**.
2. In the **Schedule review** dialog, pick a **Scheduled date**.
3. Click **Schedule review**.

This inserts an extra review without completing one. The policy's Next review date is set to the earliest incomplete review.

Changing the review frequency (via **Edit**) re-anchors the schedule: the next review date is recomputed from the last completed review (or the policy's creation date if none) plus the new frequency, and all incomplete reviews move to that new date.

A daily background sweep emails the policy owner when a review is due soon or overdue, according to the notification settings (see Administration & Settings).

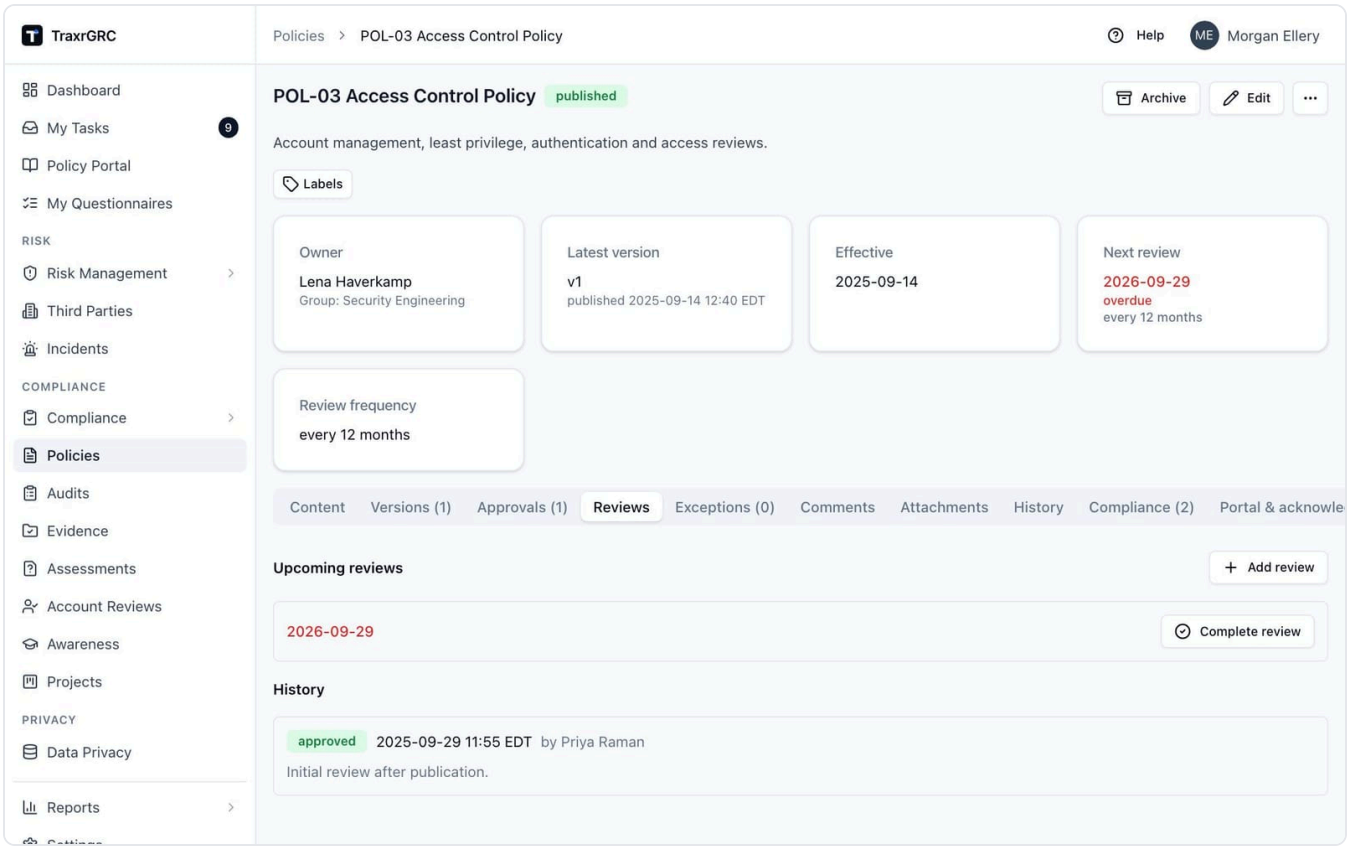


Figure 5.8. The Reviews tab of a policy whose review is overdue, with the past-due date in red and a completed review in History.

5.8 Exceptions

Exceptions document a time-boxed acceptance of a deviation from a policy, for example a team running an unsupported system until a migration finishes.

1. Go to the **Exceptions** tab and click **Request exception**.
2. Complete the dialog:

Field	What it means
Title	Short name for the exception (required).
Description	What is being accepted, and why (required).
Owner	The user responsible for the exception.
Approver	The user expected to decide it. Leave it as None to let anyone with the Review right decide.
Expires at	The date the exception ends (required).

3. Click **Request exception**.

The exception is created with the status **Requested** (amber). While it is requested and not yet expired it can be decided with **Approve** or **Reject** on its card.

Exceptions follow a separation of duties:

- The person who requested an exception can never approve or reject it, administrators included, unless an administrator has relaxed separation of duties (see Sessions under Settings). A decision made that way carries a **Self-approved** or **Self-decided** badge.
- When the request names an approver, only that person (who must be able to see the policy) or an administrator can decide it.
- When it names none, deciding needs the **Review** right on Policies (`policies:review`).
- You cannot name yourself as the approver of your own request.

The exception card shows the description, the owner, the designated approver, who requested it and who decided it; the Approve and Reject buttons appear only for someone allowed to decide. The decision is recorded with who made it and when. Statuses:

Status	Meaning	Badge color
Requested	Awaiting a decision.	Amber
Approved	The acceptance is in force.	Green
Rejected	The request was declined.	Red
Expired	Past its expiry date; computed automatically.	Gray

"Expired" is dynamic: the stored status stays Approved or Requested, but once the expiry date passes the card displays an **expired** gray badge and the Approve/Reject buttons no longer appear. A daily background sweep emails the exception owner (or the policy owner) when an exception is about to expire, per the notification settings.

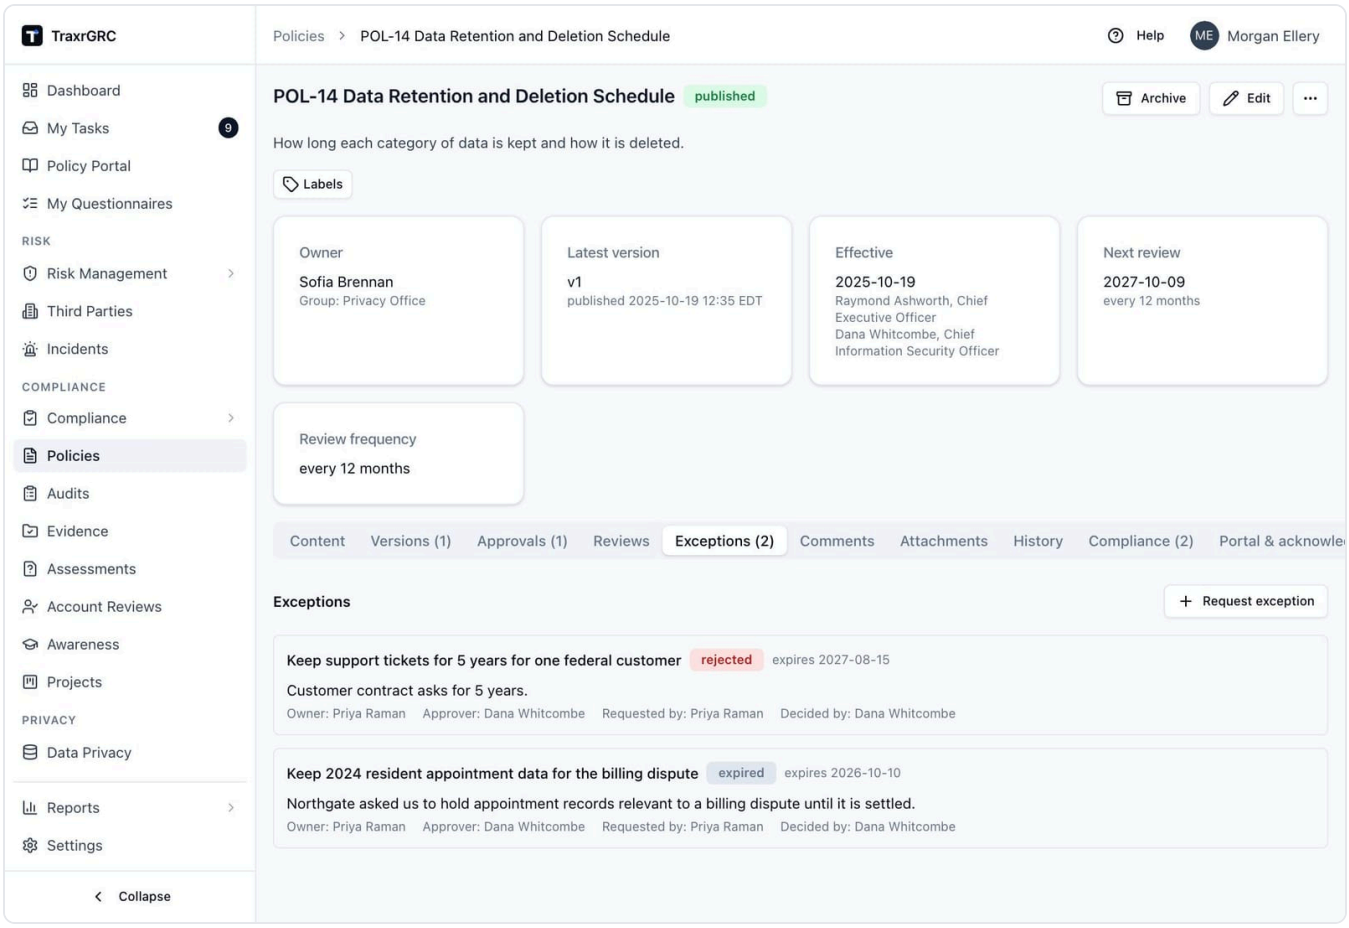


Figure 5.9. The Exceptions tab with an expired exception and a rejected one, each showing its owner, approver, requester and decider.

5.9 Archive and reactivate

- Click **Archive** to retire a draft or published policy. Archived policies cannot be submitted for approval, published or put on the Policy Portal, no longer appear as overdue, and are excluded from review and exception reminder emails. Their acknowledgement requests show as Closed.
- Click **Reactivate** on an archived policy to bring it back as a **Draft**, ready for the approval and publish cycle again.

Both transitions are recorded in the audit log.

5.10 Comments, Attachments, History and Custom Fields

- **Comments**: a discussion thread on the policy. Available to anyone who can update the policy.
- **Attachments**: upload supporting files (for example signed approvals or annexes). Available to anyone who can update the policy. Documents uploaded as new versions appear here too.
- **History**: the last 100 audit entries for this policy. Each entry shows the action, who performed it, when, and which fields changed. Actions include create, update (also used for archive and

reactivate), submit-for-approval, approve, reject, withdraw-approval, publish, record-external-approval, restore, review, portal-publish, portal-unpublish, request-acknowledgement, acknowledgement-due-date, close-acknowledgement, delete-acknowledgement-request and delete. Exception decisions and individual acknowledgements are recorded against their own records and appear in the administrator's audit log, not on this tab.
- **Custom Fields**: the Custom Fields tab, last in the row, holds administrator-defined extra fields for policies (see Administration & Settings).

5.11 Policy portal and acknowledgements

The **Policy Portal** (sidebar, under Dashboard) is where every signed-in user, whatever their role, reads the policies in force and acknowledges the ones they are asked to (for example PL-4 rules of behavior). Nothing is on the portal until someone puts it there, policy by policy.

5.11.1 What the portal shows

The portal home page has up to three sections:

- **Policies to acknowledge**: the policies you are asked to acknowledge, with the version and the due date ("Due ...", or in red "Overdue, was due ..."), and an **Acknowledge** button that opens the policy at its acknowledgment. When there is nothing to do it reads "Nothing to acknowledge right now."
- **Policies**: every policy on the portal, with its title and description, version and effective date. Type in the search box and click **Search** to match titles and descriptions.
- **My acknowledgements**: the policies and versions you have acknowledged, with the time of each.

Click a policy to open it. The page shows the title, the version number and effective date, the description, a **PDF** button for that version and, when the version came from an uploaded document, a button carrying the **source document**'s file name, followed by the policy text. It never shows the working draft, earlier versions, approvals, signatories, reviews, exceptions, owners, comments or other attachments. A policy that is archived, was never published, or is not on the portal answers "Page not found" to staff. When a published policy is edited (and so back in draft), the portal keeps showing the last published version until the next one is published.

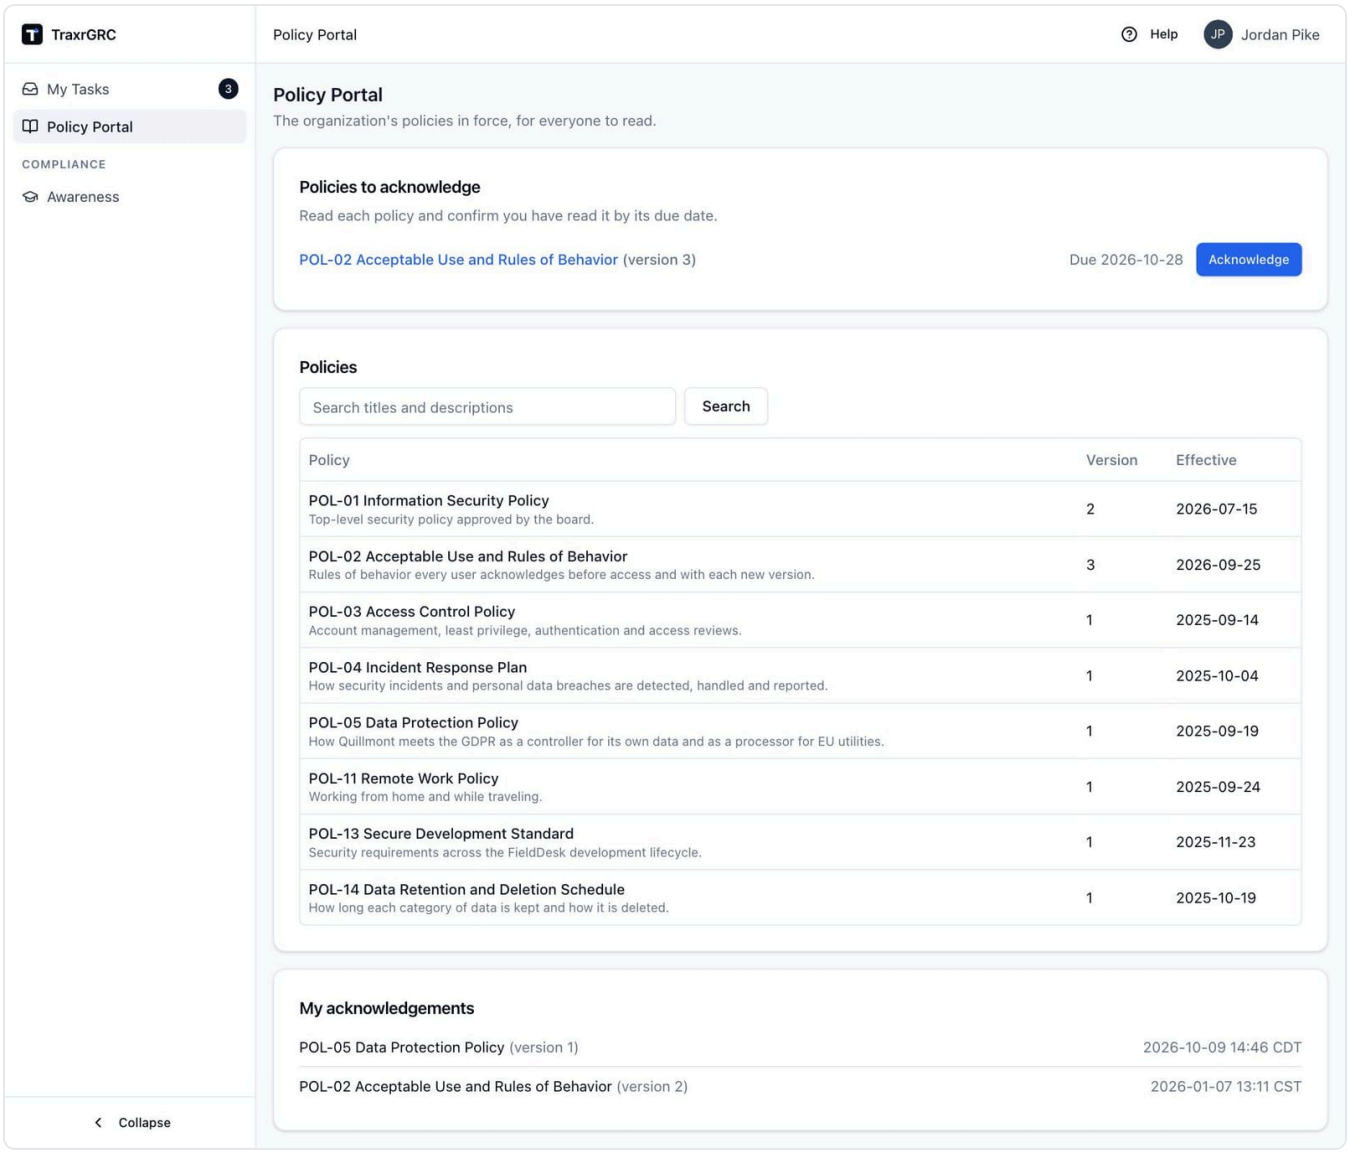


Figure 5.10. The Policy Portal as a staff member sees it: one policy to acknowledge, the policies on the portal and the acknowledgements already given.

5.11.2 Acknowledging a policy

When you are asked to acknowledge a policy:

1. Open it from **Policies to acknowledge** on the portal, your dashboard or My profile, or from the link in the email.
2. Read the policy. A note at the top links to the acknowledgment, and the panel below the text says which version you are asked to acknowledge and by when.
3. Tick the attestation checkbox in that panel.
4. Click **Acknowledge**.

The panel is replaced by "You acknowledged version N on ...".

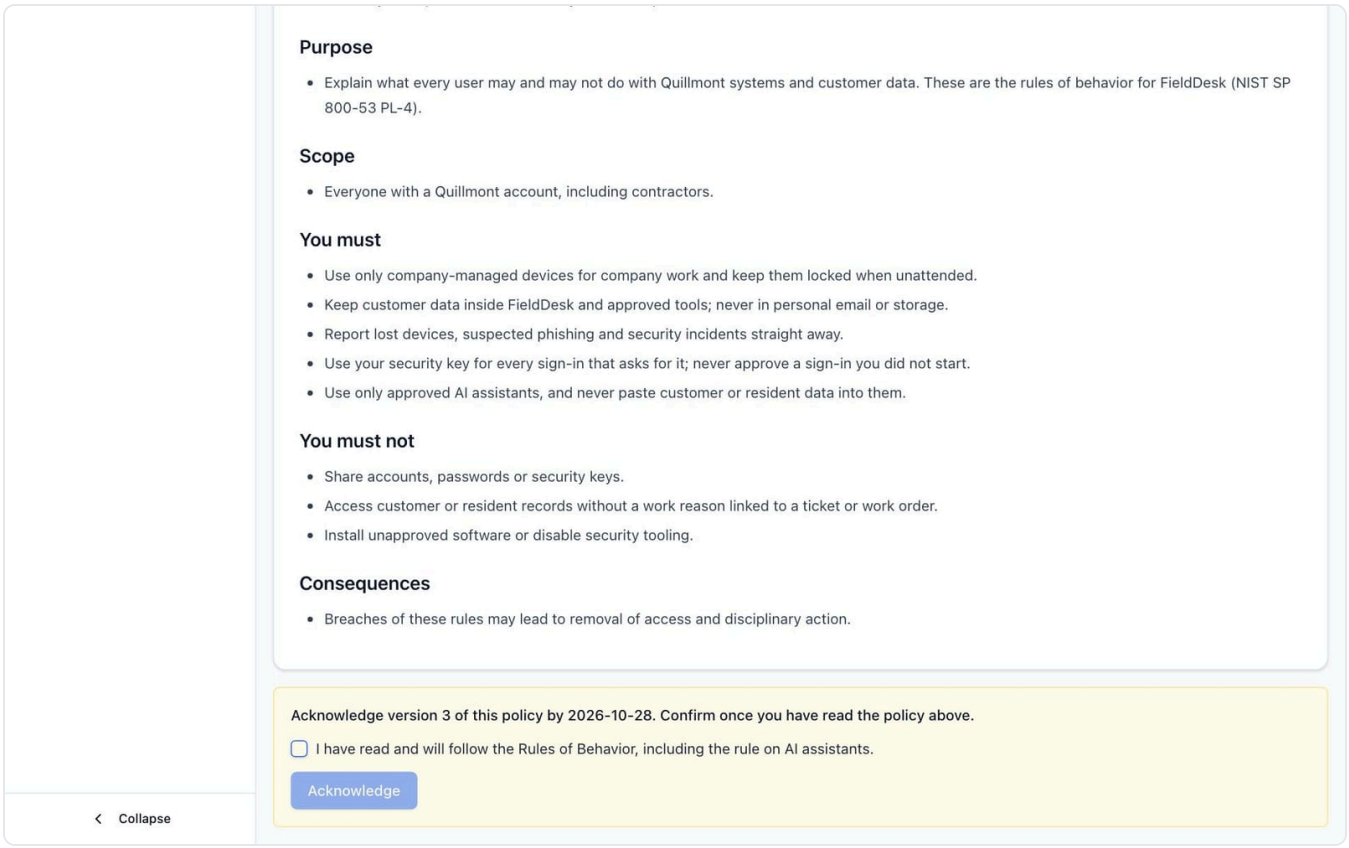


Figure 5.11. The acknowledgment panel below a policy's text on the portal: the due date, the attestation checkbox and the Acknowledge button.

5.11.3 Putting a policy on the portal

1. Open the policy and choose the **Portal & acknowledgements** tab.
2. Under **Policy portal**, click **Publish to the portal**.

The badge changes from "Not on the portal" to "On the portal", and a **View on the portal** link appears. **Take off the portal** reverses it. Both need `policies:update` and `policies:review` on the policy, since putting it on the portal discloses the text to everyone, and the policy needs a published version and must not be archived. A policy cannot be taken off while one of its acknowledgement requests is active: close the request first. Both actions are recorded on the policy's history.

5.11.4 Asking for acknowledgement

Request acknowledgement on the same tab needs `policies:update` on the policy (so its owner or a manager) and the policy on the portal. It asks for the current published version.

1. Under **Acknowledgement requests**, click **Request acknowledgement**.
2. Complete the dialog, "Request acknowledgement of version N":

Field	What it means
Everyone (every active user)	Every active user, including people who join later.
Users, Groups, Roles	When Everyone is off: any mix of named users, members of the groups, and members of groups carrying the roles. Membership is read live, so someone added to a group later is asked too.
Due date	A calendar date, 30 days ahead by default; the request is overdue from the next day. It cannot be in the past.
Attestation (optional)	What each person ticks to confirm, for example "I have read and will comply with the rules of behavior". Left blank, they confirm "I have read and understood this policy."

3. Click **Send request**.

Each person asked gets an email with a link, a reminder in the week before the due date and one notice once it has passed (never within a day of the first email; each once per due date). Deactivated accounts are left out. The request appears for them under **Policies to acknowledge** on the portal, their dashboard and My profile.

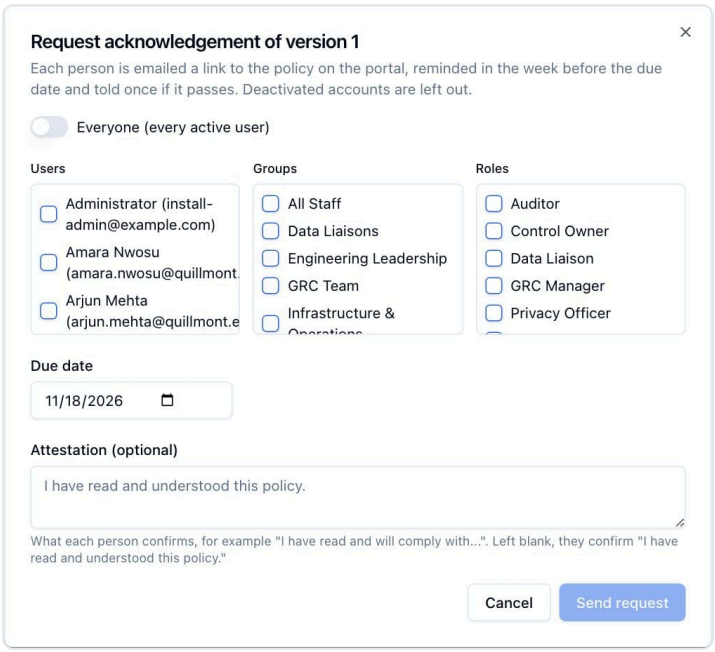


Figure 5.12. The Request acknowledgement dialog, with the audience choices, the due date and the attestation.

5.11.5 What is recorded

Each acknowledgement stores the person, the request, the policy version, the SHA-256 of that version's text, the attestation text they confirmed and the time (stored in UTC, shown in each viewer's zone). It is written once, never changed or deleted, and audited. One acknowledgement per person per version: someone asked twice for the same version acknowledges once. If the text they were shown is no longer the version the request is for, the acknowledgement is refused and they are asked to reload.

5.11.6 Managing requests

Requests are listed on the tab with the version, audience, due date, state (**Request** column), progress (**Acknowledged** column, for example "3 of 40", with how many are overdue and how many acknowledged but are no longer in the audience) and who requested it and when.

State	Meaning	Badge color
Active	Takes acknowledgements and sends reminders.	Green
Superseded	A newer version has been published since; it no longer takes acknowledgements or sends reminders.	Amber
Closed	Someone closed it, or the policy was archived.	Gray

On an active or superseded request:

- **Due date** opens an inline date field; change it and click **Save**.
- **Close** closes the request.

A request nobody has acknowledged through also has **Delete**.

A new version. Publishing a new version does not ask anyone again by default: earlier acknowledgements stay as evidence for the version they were given for, and open requests become Superseded. To ask for the new version, click **Ask again for version N** on the earlier request: the dialog opens with its audience and attestation, ready to send with a new due date.

People who leave. Deactivated accounts drop out of the audience and the outstanding and overdue counts; their acknowledgements stay and are listed (marked "(deactivated)", or "No longer in the audience" for someone who left a group). A user who has acknowledged a policy cannot be deleted, only deactivated, and a policy with acknowledgements cannot be deleted, only archived.

5.11.7 Completion and evidence

The **Completion** table on the tab lists everyone asked and everyone who acknowledged. Use the **Request** selector to show one request or **All requests**. Columns:

Column	What it shows
Name	The person, marked "(deactivated)" or "No longer in the audience" where applicable.
Email	Their email address.
Version	The version they were asked to acknowledge.
Due	The request's due date.
Status	**Acknowledged** (green), **Outstanding** (gray) or **Overdue** (red).
Acknowledged	The time they acknowledged, in your time zone.
Content hash	The first characters of the SHA-256 of the text they acknowledged; hover for the full hash.

Filter by status, search by name or email, and export it to CSV. It needs `policies:read` on the policy. For a view across policies, and to email it on a schedule, use the **Policy Acknowledgements** report in the Report Center. A notification rule on **Deadline approaching** or **Deadline overdue** for **Policy** records also fires for requests with people still to acknowledge (see Notification rules).

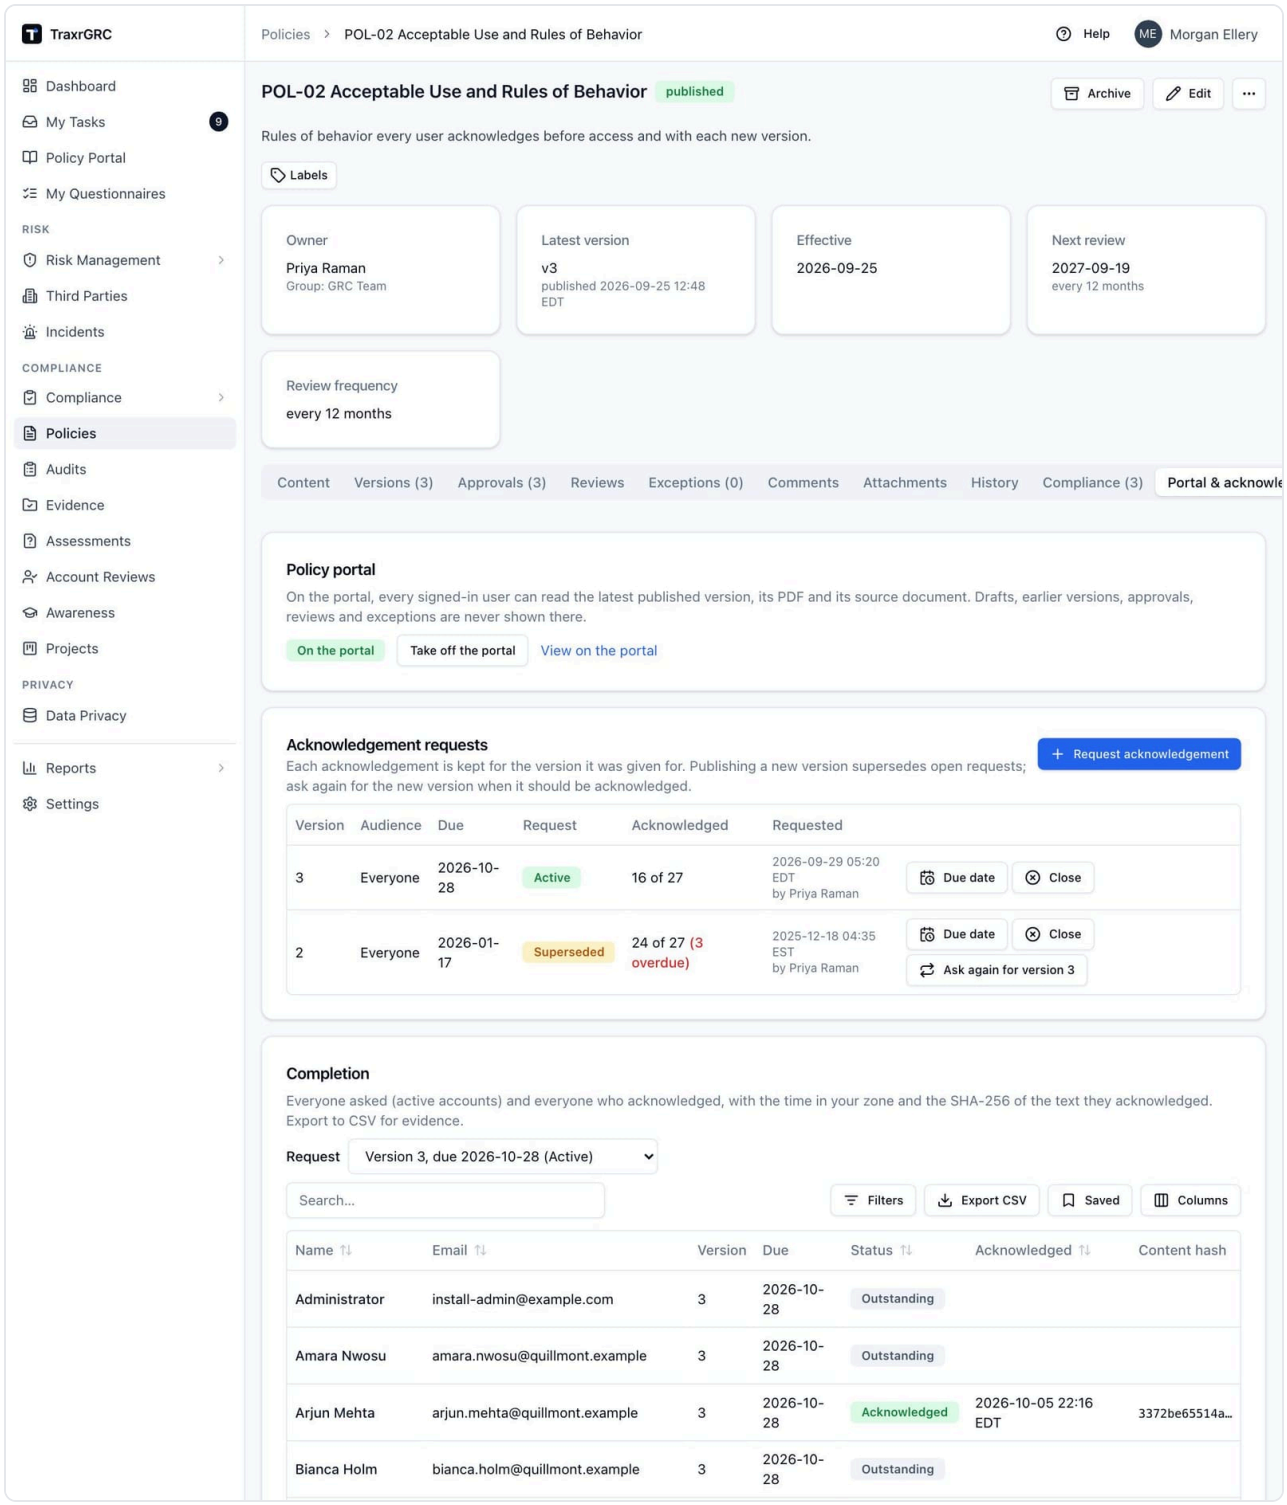


Figure 5.13. The Portal & acknowledgements tab: the policy on the portal, an active request and a superseded one with Ask again, and the completion table.

5.12 Deleting a policy

1. Open **More actions** (**…**) in the header and click **Delete**.

2. Confirm with **Delete** in the "Delete policy?" dialog.

> **Warning:** Deleting a policy permanently removes it together with its versions, approvals, reviews, exceptions, acknowledgement requests nobody answered, comments, attachments and custom-field values. This cannot be undone.

A policy that anyone has acknowledged cannot be deleted; archive it instead.

5.13 Who can do what

Access is controlled by the **policies** module rights (see Administration & Settings):

Action	Permission needed
View the list and a policy	`policies:read`
Create a policy	`policies:create`
Edit metadata or content, submit for approval, decide approvals, publish, restore versions, complete or schedule reviews, request exceptions, archive/reactivate, comment, attach files	`policies:update`; at the `own` level you must be the owner or in the owner group
Decide an exception that names no approver (never one you requested)	`policies:review`
Decide an exception that names you as approver (never one you requested)	`policies:read` on the policy
Record an outside approval or upload a new version	`policies:update` and `policies:review`, with **Settings → Policies → Allow recording outside approvals** on
Delete a policy	`policies:delete`; at the `own` level you must be the owner or in the owner group
Read the Policy Portal and acknowledge what you are asked to	Signed in; no policies right needed
Put a policy on the portal or take it off	`policies:update` and `policies:review` on the policy
Request acknowledgement, change a request's due date, close or delete it	`policies:update` on the policy
See a policy's requests and completion	`policies:read` on the policy
See the Compliance tab	`compliance:read`

Administrators bypass these checks, except that they too cannot decide an exception they requested while separation of duties is enforced.

CHAPTER 6

Third-Party Management

Third parties are the external organizations your business depends on or serves (vendors, partners and customers) and the risks they introduce. GRC Platform keeps a register of these organizations with their key contacts, links them to the risks they give rise to, tracks each vendor's data protection position as a processor (data processing agreement, sub-processors and transfer impact assessments), and records the full history of changes.

Vendor managers and procurement teams maintain the register; risk owners link vendor risks to it; the privacy team and DPO use its data protection tab. This chapter covers the **Third Parties** area (sidebar, under **Risk**).

6.1 The register

Open **Third Parties** from the navigation. One row per organization:

Column	What it shows
Name	The organization's name. Click it to open the record.
Type	Vendor (blue), Partner (purple), Customer (teal) or Other (gray). Filterable.
Status	Active (green), Offboarding (amber) or Terminated (gray). Filterable.
Data protection	Flags that need attention (see below). Empty when nothing needs attention.
Owner	The owning user, or the owning group if no user is set.
Contacts	How many contacts are recorded.
Open risks	How many linked risks currently have the Open status.
Created	When the record was created.
Labels	The third party's labels. Filter by typing part of a label name.

The Data protection column shows only the states that need attention:

Flag	Meaning
DPA missing (red)	The data processing agreement is required but not started, or still in negotiation.
DPA expiring (amber)	A signed DPA expires within 60 days.
DPA expired (red)	A signed DPA's expiry date has passed.
TIA review due (amber)	At least one transfer impact assessment is due for review within 30 days or is overdue.

Use the filter box to search names, descriptions and websites, and the Type and Status filters to narrow the list. You can sort by name, type, status, owner and creation date. Third parties can also be created and updated in bulk with **Import** (see Importing from CSV in Getting Started).

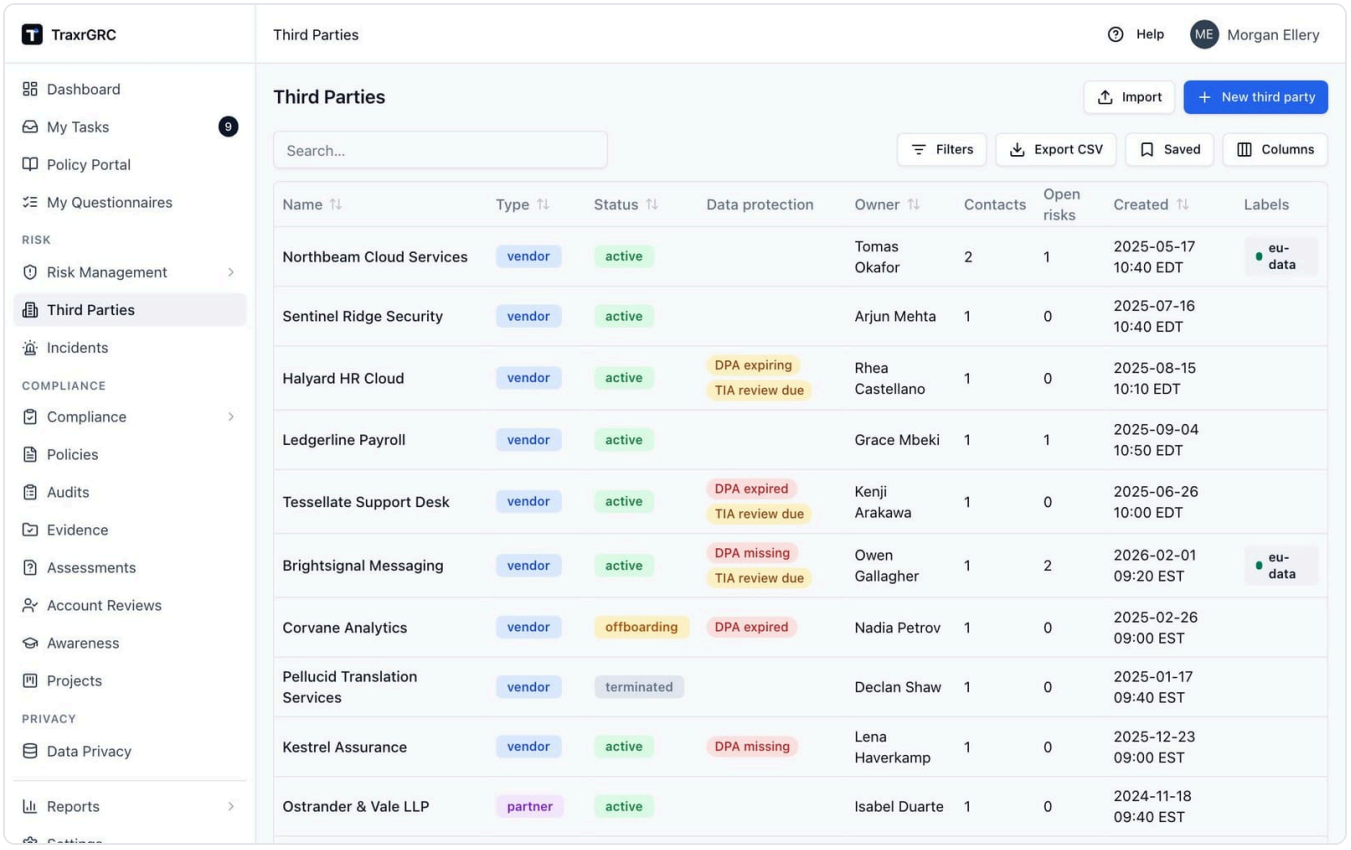


Figure 6.1. The third-party register, with type and status badges and data protection flags for missing, expiring and expired DPAs and TIA reviews due.

6.2 Creating a third party

1. Click **New third party**.

2. Complete the dialog:

Field	What it means
Name	The organization's name (required; must be unique).
Description	Optional background.
Type	Vendor, Partner, Customer or Other; defaults to Vendor.
Website	Optional; must be a valid `http://` or `https://` address when filled in.
Owner	The user responsible for the record. Both owner fields are optional.
Owner group	The group responsible for the record.

3. Click **Create third party**.

The new record starts with the status Active. To change the status later, click **Edit**: the Edit dialog adds a **Status** field with the three values. Click **Save changes** to save. Status changes are saved with the rest of the record and are recorded in the audit log.

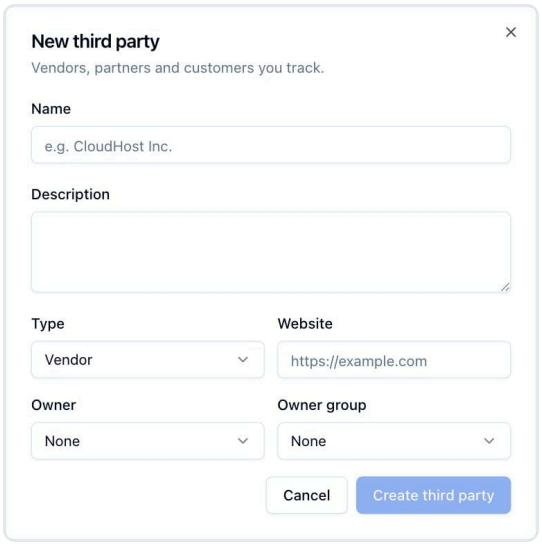


Figure 6.2. The New third party dialog, with the name, description, type, website and owner fields.

Status	Meaning
Active	The relationship is current.
Offboarding	The relationship is ending.
Terminated	The relationship has ended. Terminated vendors get no DPA or TIA reminder emails.

6.3 The detail page

The header shows the name, type and status badges, the website as a link when one is set, the description and the **Labels** picker (see Labels in Getting Started; attaching needs thirdParties:update on the record), followed by the **Edit** button and the **More actions** menu (**...**) holding **Delete**. Four cards summarize the owner (with the owner group), the number of contacts, the number of open risks and the creation date.

One row of tabs covers the rest of the page:

- **Contacts**, **Risks**, **Comments**, **Attachments** and **History**: the organization itself.
- **Data protection**: the vendor's data processing agreement, sub-processors and transfer impact assessments.
- **Custom Fields**: administrator-defined extra fields for third parties (see Administration & Settings).

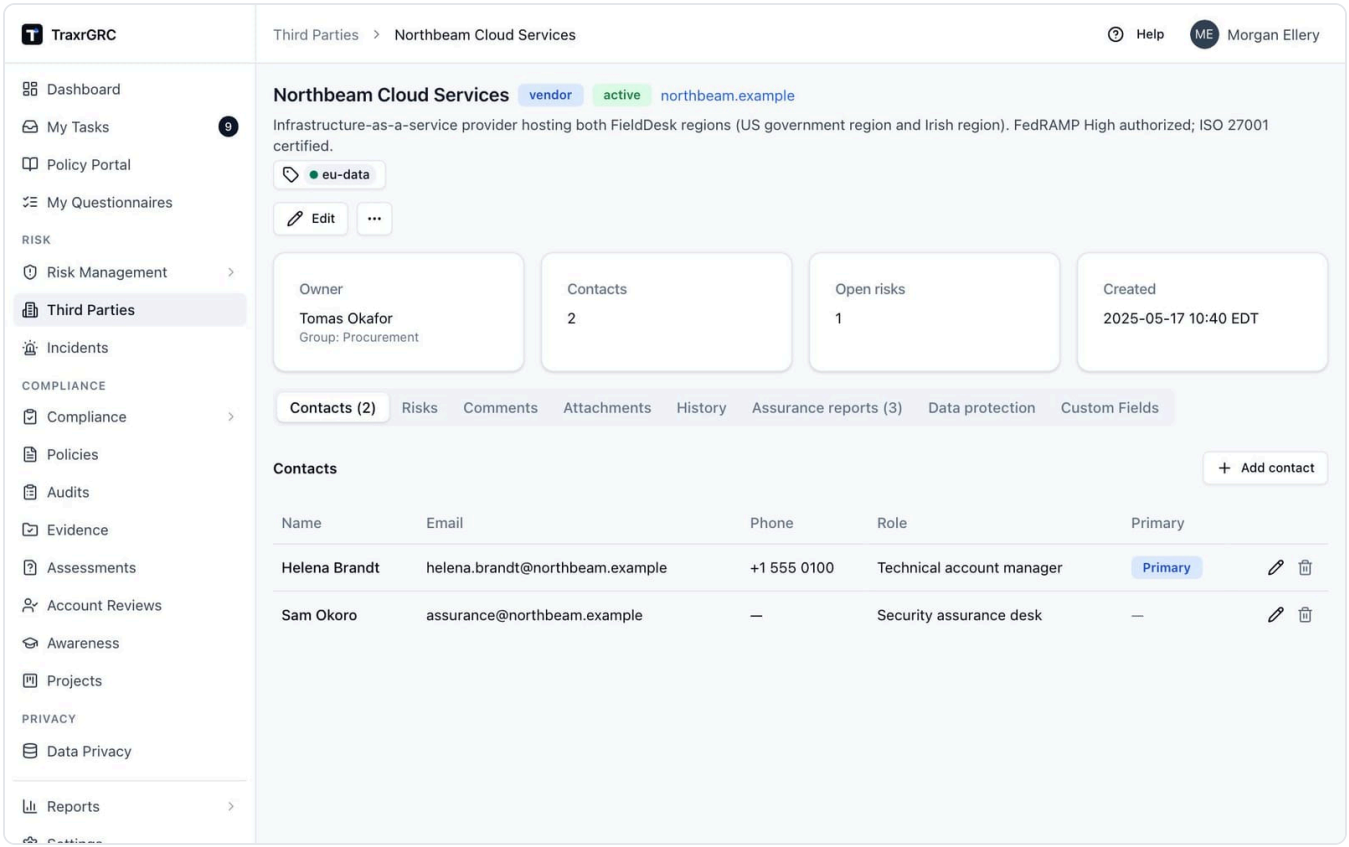


Figure 6.3. A vendor's detail page: the header with type and status badges and website, the summary cards, and the Contacts tab with the primary contact marked.

6.3.1 Contacts tab

The tab label shows the number of contacts. The contacts table lists each person's name, email, phone, role and whether they are the **Primary** contact (blue badge).

To add a contact:

1. Click **Add contact**.
2. Complete the dialog: **Name** (required), **Email**, **Phone**, **Role**, and the **Primary contact** checkbox.
3. Click **Add contact**.

To change a contact, click the pencil icon on its row, edit the fields and click **Save changes**. To remove one, click the bin icon and confirm with **Remove**.

Only one contact can be primary: marking a contact as primary automatically clears the flag on every other contact, both when adding and when editing. If you remove the primary contact, the record is left without one; nothing is promoted automatically. The **Remove contact?** confirmation reminds you of this. Adding, editing and removing contacts are written to the audit log.

6.3.2 Risks tab

This tab lists the risks from the **Third-party** classification that are linked to the organization: the vendor risks you track. The table shows each risk's title (linking to the risk register), inherent and residual score chips, status, owner and next review date.

Attaching a risk:

1. Click **Attach risk**. The button appears only when at least one risk in the Third-party classification is not yet linked to any organization. If it is missing, create the risk in the risk register first (see Risk Management).
2. In the dialog, choose the **Risk** and click **Attach risk**.

The risk stays in the risk register; attaching only creates the link. You need update access to the risk itself for the attach to succeed.

Detaching a risk: click **Detach** on the row and confirm with **Detach**. The link is removed; the risk itself stays in the register. Detaching is also audited.

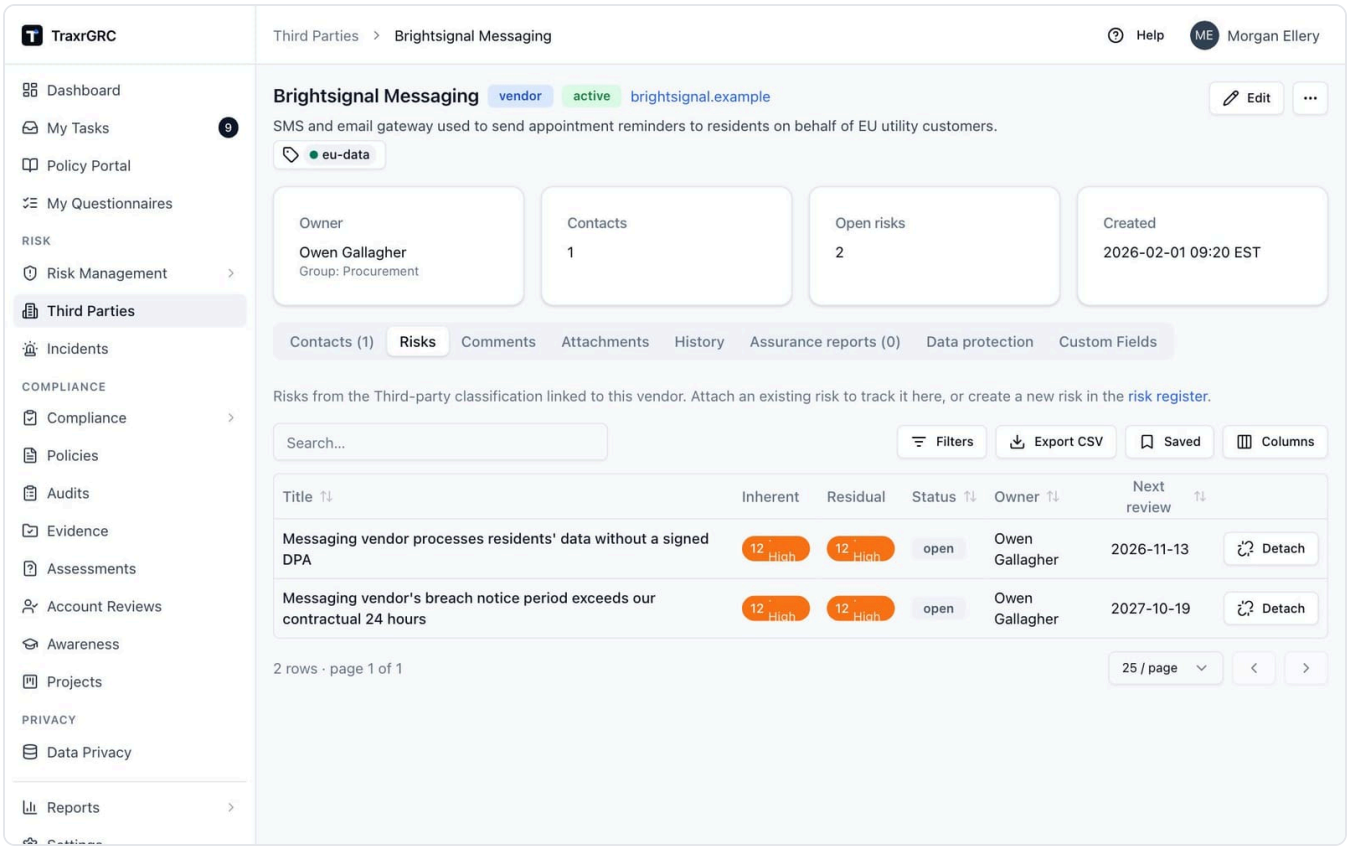


Figure 6.4. The Risks tab of a vendor, listing the linked third-party risks with their inherent and residual scores, status, owner, next review date and Detach button.

6.3.3 Comments, Attachments and History

- **Comments**: a discussion thread on the record, available to anyone who can update it.
- **Attachments**: supporting files such as contracts or due-diligence reports, available to anyone who can update the record.
- **History**: the last 100 audit entries for this record (create, update, attach-risk, detach-risk, delete), each showing the action, actor, time and changed fields. Status and DPA changes appear as update entries; adding or removing a sub-processor appears as a create or delete entry with "subProcessor" as the changed field. Contact changes and transfer impact assessment changes are recorded against their own records and appear in the administrator's audit log, not on this tab.

6.4 Data protection

The **Data protection** tab, after **History**, covers the vendor as a processor of personal data. It has up to four sections. Everything on it needs `thirdParties:update` on the vendor to change; without it the tab is read-only.

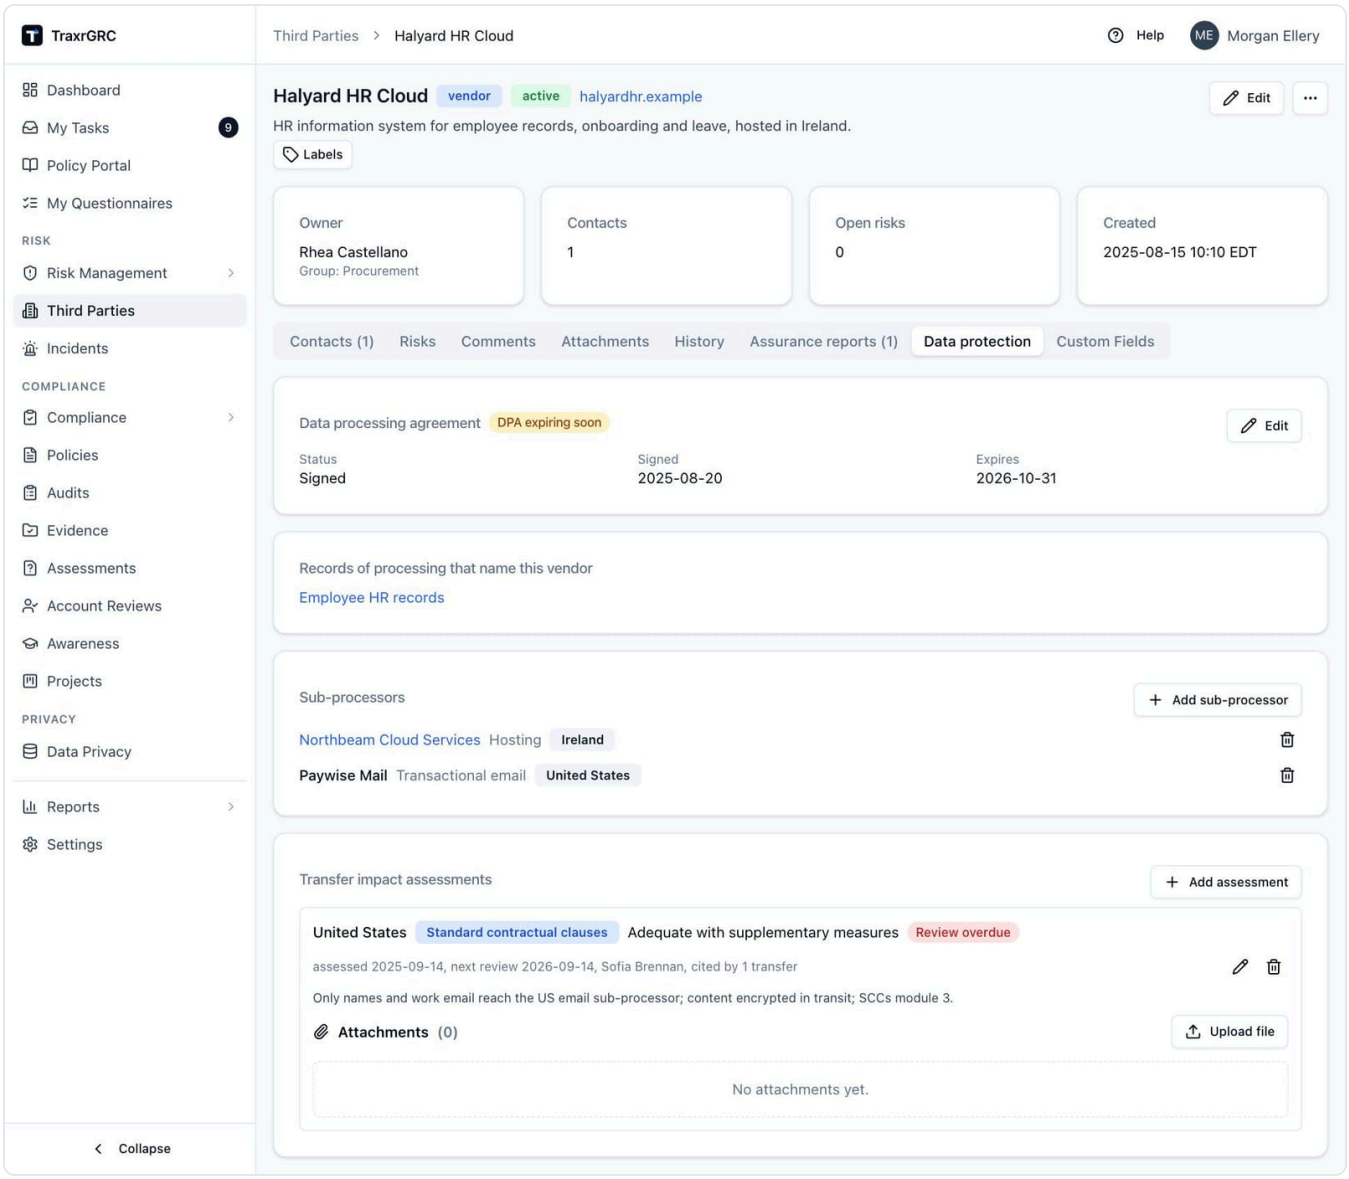


Figure 6.5. The Data protection tab of a vendor whose DPA expires soon and whose transfer impact assessment review is overdue, with the records of processing and sub-processors that name it.

6.4.1 Data processing agreement

The section heading carries a chip with the DPA's state, and the section shows the **Status**, **Signed** date and **Expires** date.

To record the agreement:

1. Click **Edit** in the section.
2. In the **Data processing agreement** dialog, choose the **Status**, and enter the **Signed** date and, optionally, **Expires (optional)**.
3. Click **Save**.

Status	Meaning
Not assessed	Nobody has decided whether a DPA is needed.
Not required	The vendor processes no personal data for you.
Required, not started	A DPA is needed and work has not begun.
In negotiation	Terms are being agreed.
Signed	The DPA is in force. Needs the signed date; the expiry cannot be before it.

The chip reads:

Chip	When
DPA not assessed (gray)	No status set.
DPA not required (gray)	Status Not required.
DPA missing (red)	Required, not started, or In negotiation.
DPA signed (green)	Signed, with no expiry or an expiry more than 60 days away.
DPA expiring soon (amber)	Signed, expiring within 60 days.
DPA expired (red)	Signed, and the expiry date has passed.

The vendor's owner gets one email when a signed DPA enters the 60 days and one when it expires; recording a new expiry date starts the cycle again. Terminated vendors are skipped.

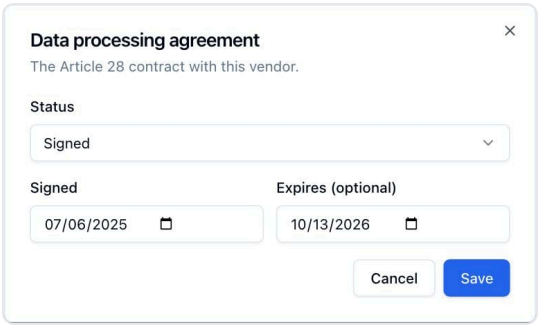


Figure 6.6. The Data processing agreement dialog, with the status, signed date and optional expiry date.

6.4.2 Records of processing that name this vendor

Lists the records of processing (see Data Privacy) that name this vendor as a processor, or "as sub-processor". Each links to the record. The section is shown only if you have `dataPrivacy:read`.

6.4.3 Sub-processors

Lists each company the vendor engages to process the data, with its service and country. A sub-processor that also has its own third-party record links to it.

1. Click **Add sub-processor**.
2. Enter the **Name** (required), **Service (optional)** and **Country (optional)**.
3. Optionally, under **Also a third party here (optional)**, choose the matching third-party record. A vendor cannot be its own sub-processor.
4. Click **Add sub-processor**.

Remove a sub-processor with the bin icon on its row.

6.4.4 Transfer impact assessments

Each assessment records whether a destination country's laws and practices let the transfer safeguard work as intended. Each entry shows the country, the safeguard, the conclusion, a review chip when needed, the assessed and next review dates, the owner, how many transfers in records of processing cite it, the supplementary measures, and its own attachments for the full assessment.

To add an assessment:

1. Click **Add assessment**.
2. Complete the dialog:

Field	What it means
Country	The destination country (required).
Safeguard	Adequacy decision, Standard contractual clauses, Binding corporate rules, EU-US Data Privacy Framework, Derogation (Art. 49) or Other. Defaults to Standard contractual clauses.
Conclusion	Adequate protection, Adequate with supplementary measures, or Not adequate.
Supplementary measures	The measures the conclusion relies on. Required when the conclusion is Adequate with supplementary measures.
Assessed	The assessment date; defaults to today.
Next review	When the assessment must be reviewed.
Owner	The person responsible for the assessment.

3. Click **Save assessment**.
4. Attach the full assessment document to the entry once it is saved.

Edit an assessment with its pencil icon, or delete it with its bin icon. Deleting an assessment leaves the transfers that cited it in place without the link.

Review chip	When
Review due soon (amber)	The next review is within 30 days.
Review overdue (red)	The next review date has passed.

The assessment's owner (or, without one, the vendor's owner) gets one email when the review comes within 30 days and one when it is overdue; a new review date starts the cycle again. With the daily digest on, these go in the digest and are still sent only once. Terminated vendors are skipped.

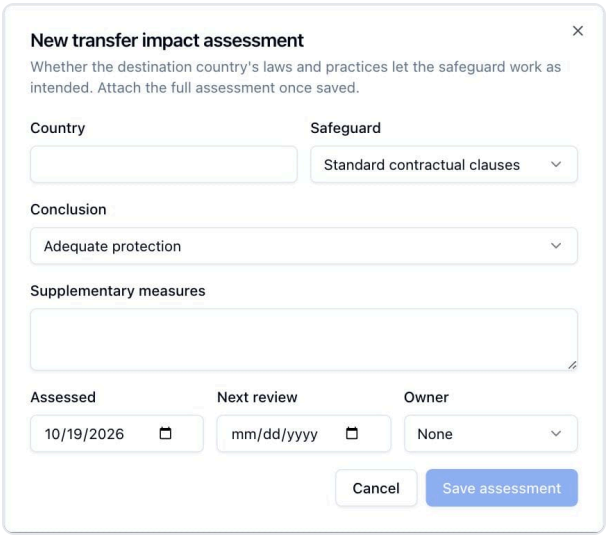


Figure 6.7. The New transfer impact assessment dialog, with the country, safeguard, conclusion, supplementary measures, dates and owner.

6.5 Deleting a third party

1. Open **More actions** (**...**) in the header and click **Delete**.
2. Confirm with **Delete** in the "Delete third party?" dialog.

Deleting removes the record together with its contacts, sub-processors, transfer impact assessments (with their files), comments, attachments and labels. A record of processing that named the vendor first moves on to its next processor. The delete is **blocked while any risks are attached**; the error tells you how many risks must be detached first. Deleting a record cannot be undone.

6.6 Who can do what

Access is controlled by the **thirdParties** module rights (see Administration & Settings):

Action	Permission needed
View the register and a record	`thirdParties:read`
Create a record	`thirdParties:create`
Edit a record (including status), add/edit/delete contacts, change the Data protection tab, comment, attach files	`thirdParties:update`; at the `own` level you must be the owner or in the owner group
Delete a record	`thirdParties:delete`; at the `own` level you must be the owner or in the owner group
See the records of processing that name a vendor	`dataPrivacy:read`

Attaching and detaching risks additionally require `risks:update` on the specific risk; at the `own` level you must own that risk or belong to its group. The attach dialog only lists risks for users who also have `risks:read`. Administrators bypass all checks.

CHAPTER 7

Incident Management

Incident Management is the record of security events that went wrong: a phishing campaign, a malware infection, a lost laptop, a service outage. You file each incident under a classification, give it an owner, link it to the affected asset or third party, and keep a full history from first report to resolution. Incidents that involve personal data can be assessed as personal data breaches, with the GDPR 72-hour notification clock and a log of who was told and when. Closed incidents keep their resolution, so the record stays useful for trend reporting and the next audit.

The security operations team typically files and works incidents; the Data Protection Officer assesses and tracks breaches; the CISO and auditors read the register.

7.1 Who can do what

Access to incidents is controlled by the **incidents** module rights. Administrators can do everything regardless of role.

Action	Permission
See the incidents list and detail pages	`incidents:read`
Create an incident (**New incident**)	`incidents:create`
Edit, close, reopen, assess a breach, log notifications, comment, upload attachments, edit custom fields	`incidents:update`, at the "all" level for every record or "own" for records you or your groups own
Link records of processing to a breach	`incidents:update` on the incident and `dataPrivacy:read`
Delete an incident	`incidents:delete` (all or own)
Manage the classification catalog	`incidents:update` at the "all" level
Import incidents from CSV	`incidents:create` or `incidents:update`

A record is "yours" when you are listed as its owner or your group is the owner group. The **Classifications** button appears only for users whose `incidents:update` right is at the "all" level, because the catalog is shared by every incident.

7.2 Managing the classification catalog

Classifications are the categories incidents are filed under. The demo catalog, for example, holds *Phishing*, *Malware*, *Unauthorized access*, *Personal data breach*, *Service outage*, *Lost or stolen device* and *Third-party incident*.

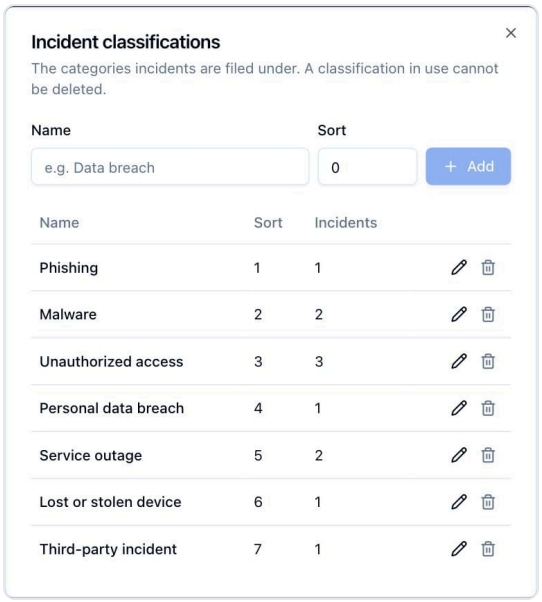


Figure 7.1. The Incident classifications dialog, listing each classification with its sort value and the number of incidents filed under it.

1. Open **Incidents** in the navigation.
2. Click **Classifications** (top right). The **Incident classifications** dialog lists the catalog with the **Sort** value and the number of **Incidents** filed under each classification.
3. To add one, type a **Name** and a **Sort** number and click **Add**. Classifications are ordered by sort value, then name, everywhere they are listed.
4. To rename or re-order one, click the pencil icon in its row, change the **Name** or **Sort** value in the **Edit classification** dialog, and click **Save changes**.
5. To remove one, click the trash icon and confirm with **Delete**.

A classification that is in use cannot be deleted. If any incident is filed under it, the delete fails with a message saying how many incidents are classified with it. Reclassify those incidents first.

7.3 The incidents list

The **Incidents** page lists every incident you can read. Click a column heading to sort; the **Classification** and **Status** columns have filter dropdowns.

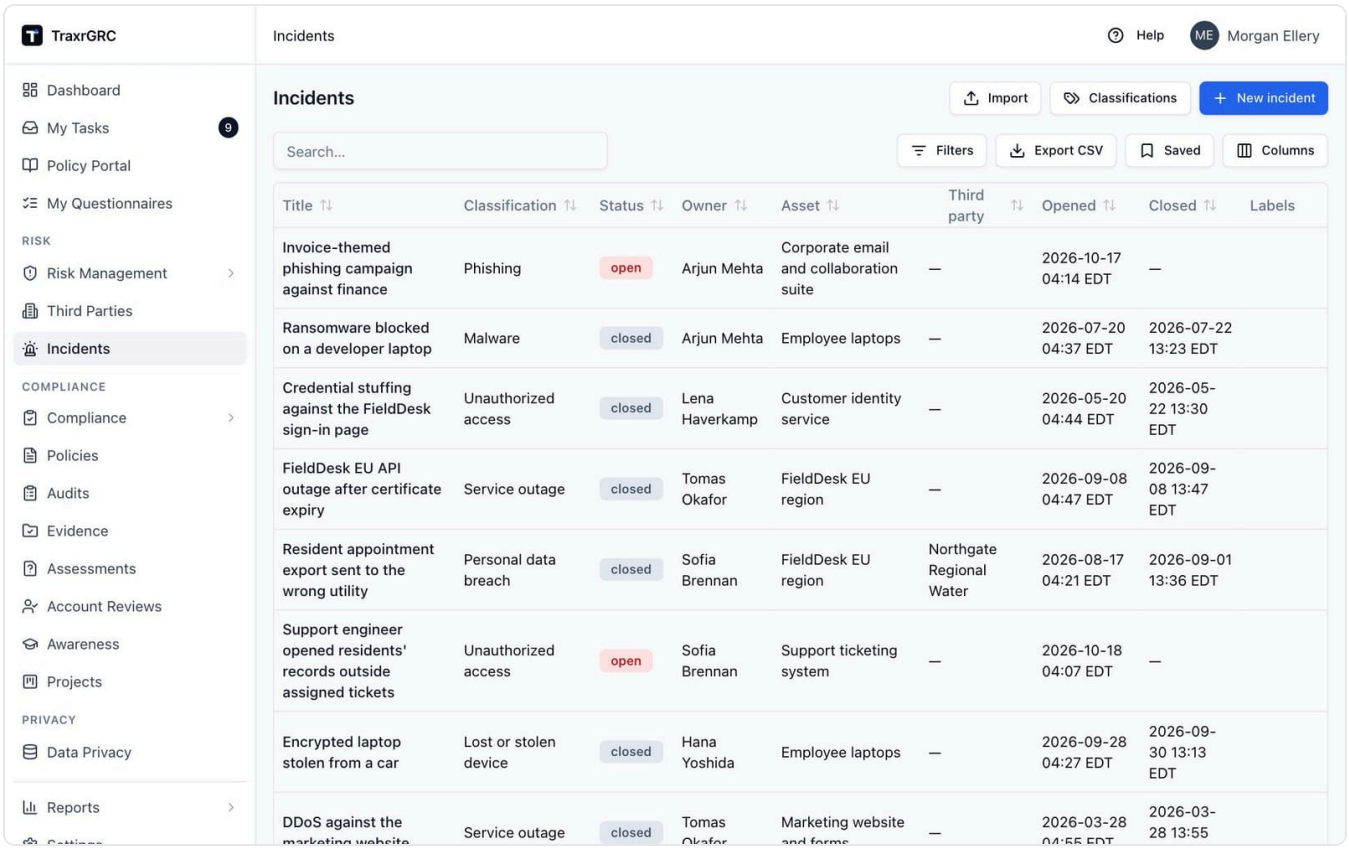


Figure 7.2. The incidents list, showing open and closed incidents with their classification, owner, linked asset or third party, and opened and closed dates.

Column	What it shows
Title	The incident's title. Click it to open the incident.
Classification	The incident's category, or a dash. Filter by one classification.
Status	**open** or **closed** (see below). Filter by status.
Owner	The owner user, or the owner group when no user is set.
Asset	The linked asset, or a dash.
Third party	The linked third party, or a dash.
Opened	When the incident was opened.
Closed	When the incident was closed, or a dash.
Labels	The incident's labels. Filter by typing part of a label name.

Incidents have two statuses, shown as colored badges:

Status	Badge	Meaning
open	Red	The incident is being worked. New incidents start here.
closed	Gray	The incident was resolved; a resolution and closed date are recorded.

The **Import** button beside **New incident** creates or updates incidents in bulk from a CSV file (see Importing from CSV under Getting Started). Personal data breach details are not imported.

7.4 Creating an incident

1. On the incidents list, click **New incident**.
2. Complete the dialog (only the title is required):

Field	What it means
Title	Short name for the incident, e.g. "Phishing campaign targeting finance".
Description	Free text: what happened, the impact, the current state.
Classification	Category from the catalog, or *None*.
Asset	The affected asset from your asset register, or *None*.
Third party	The third party involved, or *None*.
Owner	The user responsible for handling the incident.
Owner group	The group responsible; complements or replaces the owner user.

3. Click **Create incident**.

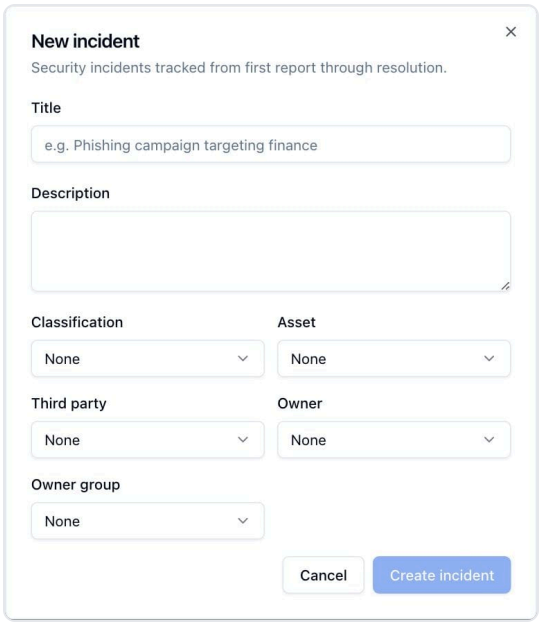


Figure 7.3. The New incident dialog with its title, description, classification, asset, third party, owner and owner group fields.

The incident is created with status **open**, and **Opened** is set to the current time. The creation is recorded in the audit log and appears on the incident's **History** tab.

7.5 The incident detail page

Click an incident's title to open it.

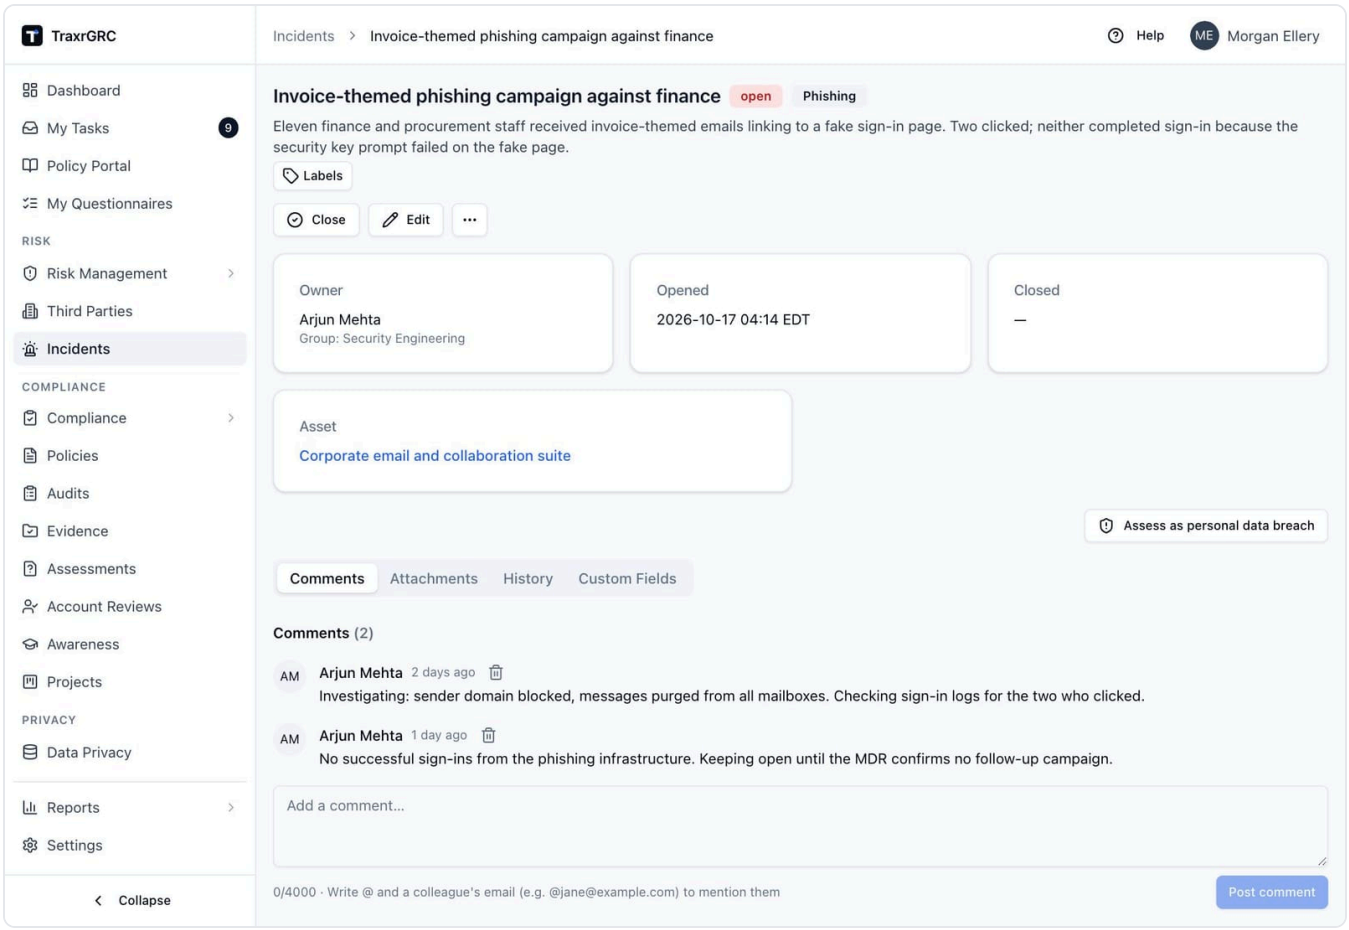


Figure 7.4. An open incident's detail page: status and classification badges, the Close, Edit and Delete buttons, the Owner, Opened and Closed cards, the linked asset, and the Assess as personal data breach button.

The page shows:

- **Header.** The title, the status badge, a red **Personal data breach** badge when the incident is marked as one, the classification badge, the description and the **Labels** picker (see Labels in Getting Started; attaching needs incidents:update on the incident). The **Close** (or **Reopen**) and **Edit** buttons and the **More actions** menu (**...**, with **Delete**) sit on the right; each appears only when you hold the right to use it.
- **Owner, Opened and Closed cards.** The owner card also names the owner group when one is set.
- **Asset and Third party cards**, when the incident is linked to either. The asset name opens the asset register; the third party name opens that third party's page.
- **The personal data breach section** (see Personal data breaches below).
- **Resolution card**, on a closed incident.
- **Comments**, **Attachments**, **History** and **Custom Fields** tabs.

7.6 Closing and reopening an incident

Closing records how the incident was resolved and stamps the closed date.

1. Open the incident and click **Close**.
2. In the **Close incident** dialog, write the **Resolution**: what fixed the incident, e.g. "Credentials reset, EDR coverage verified". A resolution is required; the button stays disabled until you have typed something.
3. Click **Close incident**.

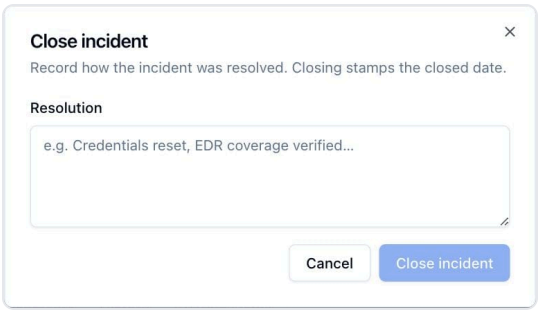


Figure 7.5. The Close incident dialog, asking for the resolution before the incident can be closed.

The incident moves to **closed**, the **Closed** card shows the closing time, and the resolution appears in a **Resolution** card. The close is recorded in the audit log.

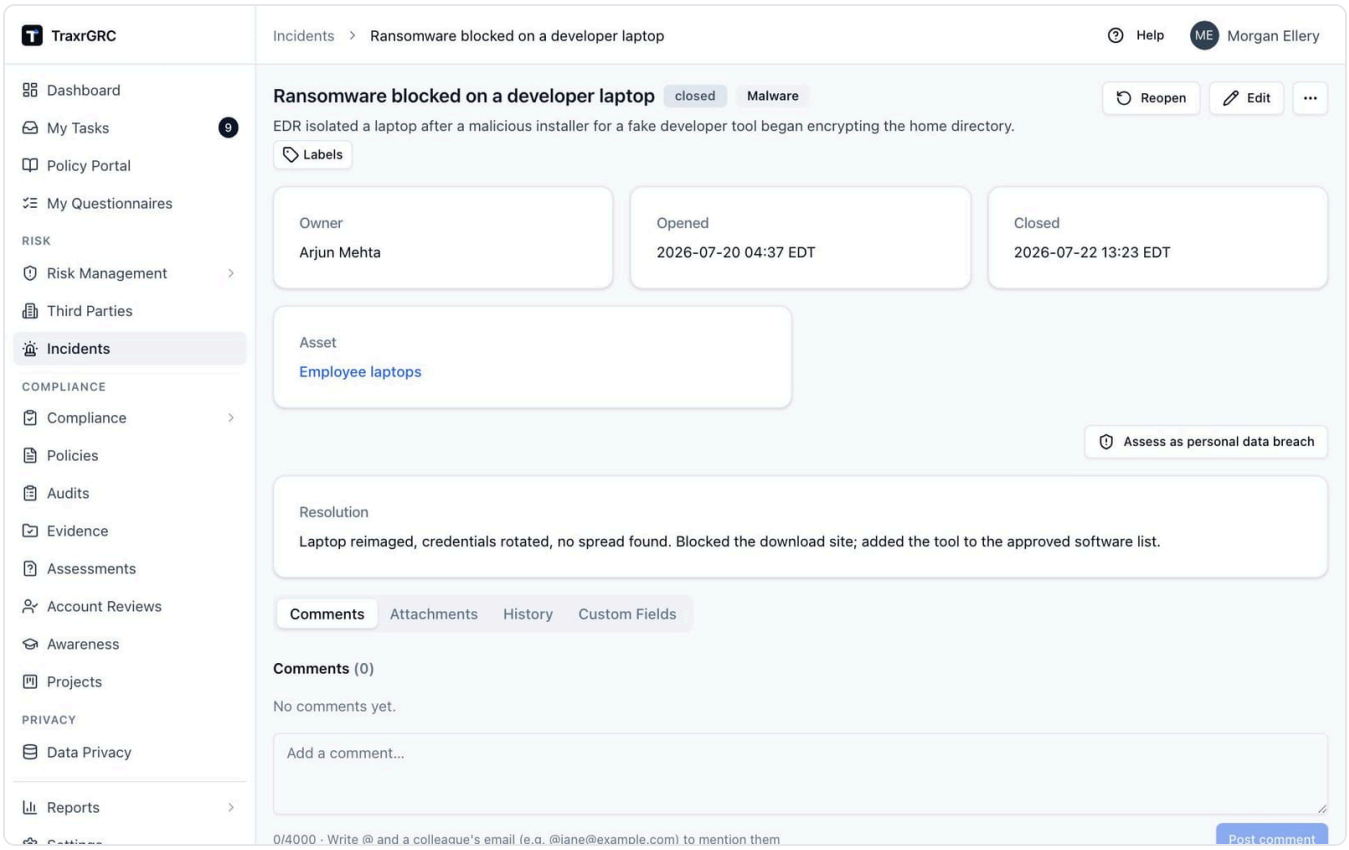


Figure 7.6. A closed incident, with the closing time on the Closed card and the Resolution card below.

To reopen a closed incident:

1. Click **Reopen**.
2. Confirm with **Reopen** in the **Reopen incident?** dialog.

The incident returns to **open** and the closed date is cleared. The recorded resolution is kept in the record: it is hidden while the incident is open, shown again if you close the incident a second time with a new resolution, and it stays visible in the History tab's change entries.

7.7 Personal data breaches

An incident that involves personal data is assessed as a personal data breach under GDPR Articles 33 and 34. The assessment records when the breach became known, the risk to individuals, and the decision to notify the supervisory authority and the people affected. TraxrGRC then tracks the 72-hour deadline for notifying the authority.

7.7.1 Assessing an incident as a breach

1. Open the incident and click **Assess as personal data breach** (shown on incidents not yet marked, to users with `incidents:update`).

2. The **Personal data breach assessment** dialog opens with **This incident is a personal data breach** ticked and **Became known** set to now. Times are entered and shown in your timezone, which the field labels name.
3. Complete the fields below and click **Save assessment**.

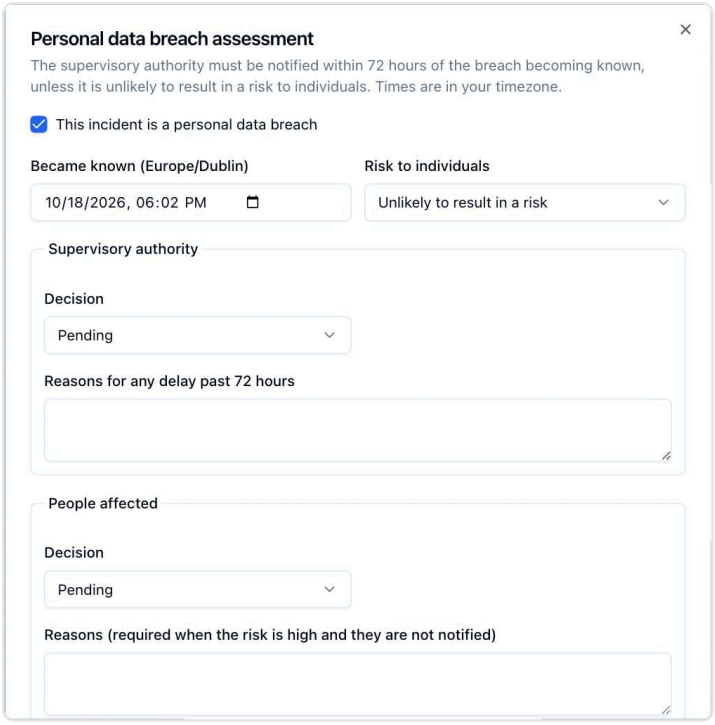


Figure 7.7. The Personal data breach assessment dialog, with the became-known time, risk to individuals, and the supervisory authority and people affected decisions.

Field	What it means
This incident is a personal data breach	The flag. Untick it if the assessment shows the incident did not involve personal data.
Became known	When your organization became aware of the breach. The 72-hour deadline runs from here. Required.
Risk to individuals	*Not assessed*, *Unlikely to result in a risk*, *Risk to individuals* or *High risk to individuals*.
Supervisory authority: Decision	*Pending*, *Notified* or *Not required*. *Notified* adds **Notified at** and an optional **Authority reference**.
Supervisory authority: reasons	Labeled **Why notification is not required** when the decision is *Not required*, otherwise **Reasons for any delay past 72 hours**. Required when notification is not required, or when the authority was notified after 72 hours.
People affected: Decision	*Pending*, *Notified* (adds **Notified at**) or *Not required*.
People affected: Reasons	Required when the risk is high and the decision is *Not required*: the Art. 34(3) ground, e.g. the data was encrypted.
People affected (approx.) / Records affected (approx.)	Approximate counts.
Categories of data affected	Tags, e.g. "Health data".

Saving refuses an incomplete assessment and tells you what is missing, for example "Enter when the authority was notified." To change an assessment later, click **Edit assessment** on the breach card.

7.7.2 The breach card

Once marked, the incident shows a red **Personal data breach** badge in its header and a **Personal data breach** card. While the authority notification is still open and **Became known** is set, the card opens with a deadline banner: the time left in hours and minutes ("41h 13m left", or "Overdue by 8h 3m" once the deadline has passed), a bar across the 72 hours marked at 24 and 48 hours, and the time the breach became known and the deadline in your own time zone. The banner is blue on the first day, amber on the second, and red on the last day and once overdue, and it updates every minute. If you may edit the incident, **Log authority notification** in the banner opens the log dialog with the supervisory authority already chosen.

Otherwise the card title carries a chip showing how the notification ended:

Chip	Meaning
Authority notified within 72 hours	The authority decision is *Notified* and the time is within the deadline. Green.
Authority notified after 72 hours	Notified, but after the deadline. Orange.
Authority notice not required	The authority decision is *Not required*. Gray.
Time of awareness not set	**Became known** is empty. Amber.

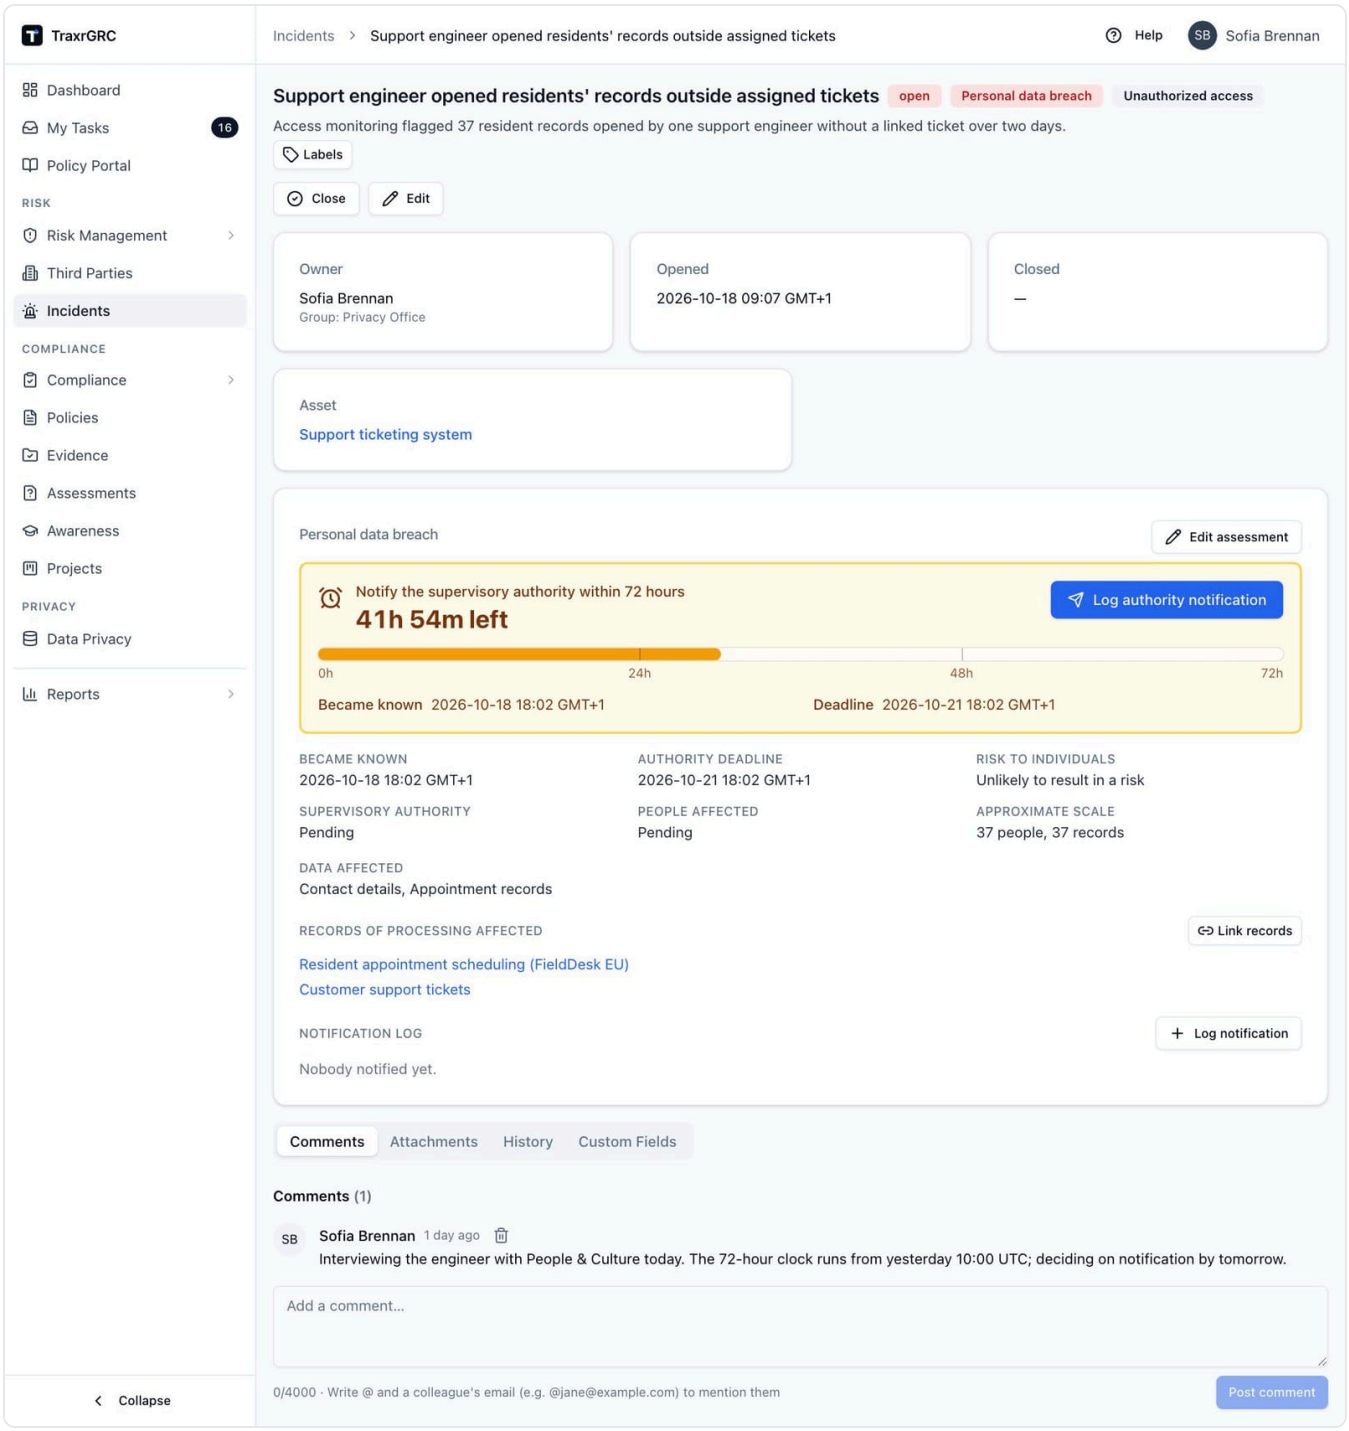


Figure 7.8. A breach under assessment: the deadline banner counting down the time left to notify the authority, with the decisions still pending.

The card lists **Became known**, the **Authority deadline** (72 hours later), **Risk to individuals**, the **Supervisory authority** and **People affected** decisions with their times and reference, the **Approximate scale**, any reasons, and **Data affected**. It has two further sections:

- **Records of processing affected.** The records of processing whose data was involved, each linking to its Data Privacy record. **Link records** opens a searchable checklist; ticking or unticking

a record saves at once (requires `dataPrivacy:read`). Records you cannot read are counted, not named.

- **Notification log.** A record of contact: who was told and when. Click **Log notification**, choose the **Recipient type** (*Supervisory authority*, *Data subjects*, *Controller* when you act as processor, or *Other*), enter the **Recipient**, the **Notified at** time and an optional **Reference**, and click **Log notification**. The trash icon removes an entry. Logging a notification to the supervisory authority while the authority decision is pending records the decision as notified at that time, which stops the countdown; a notification after the 72 hours does so only once the reasons for the delay are in the assessment. Logging one to data subjects does the same for the data subjects decision.

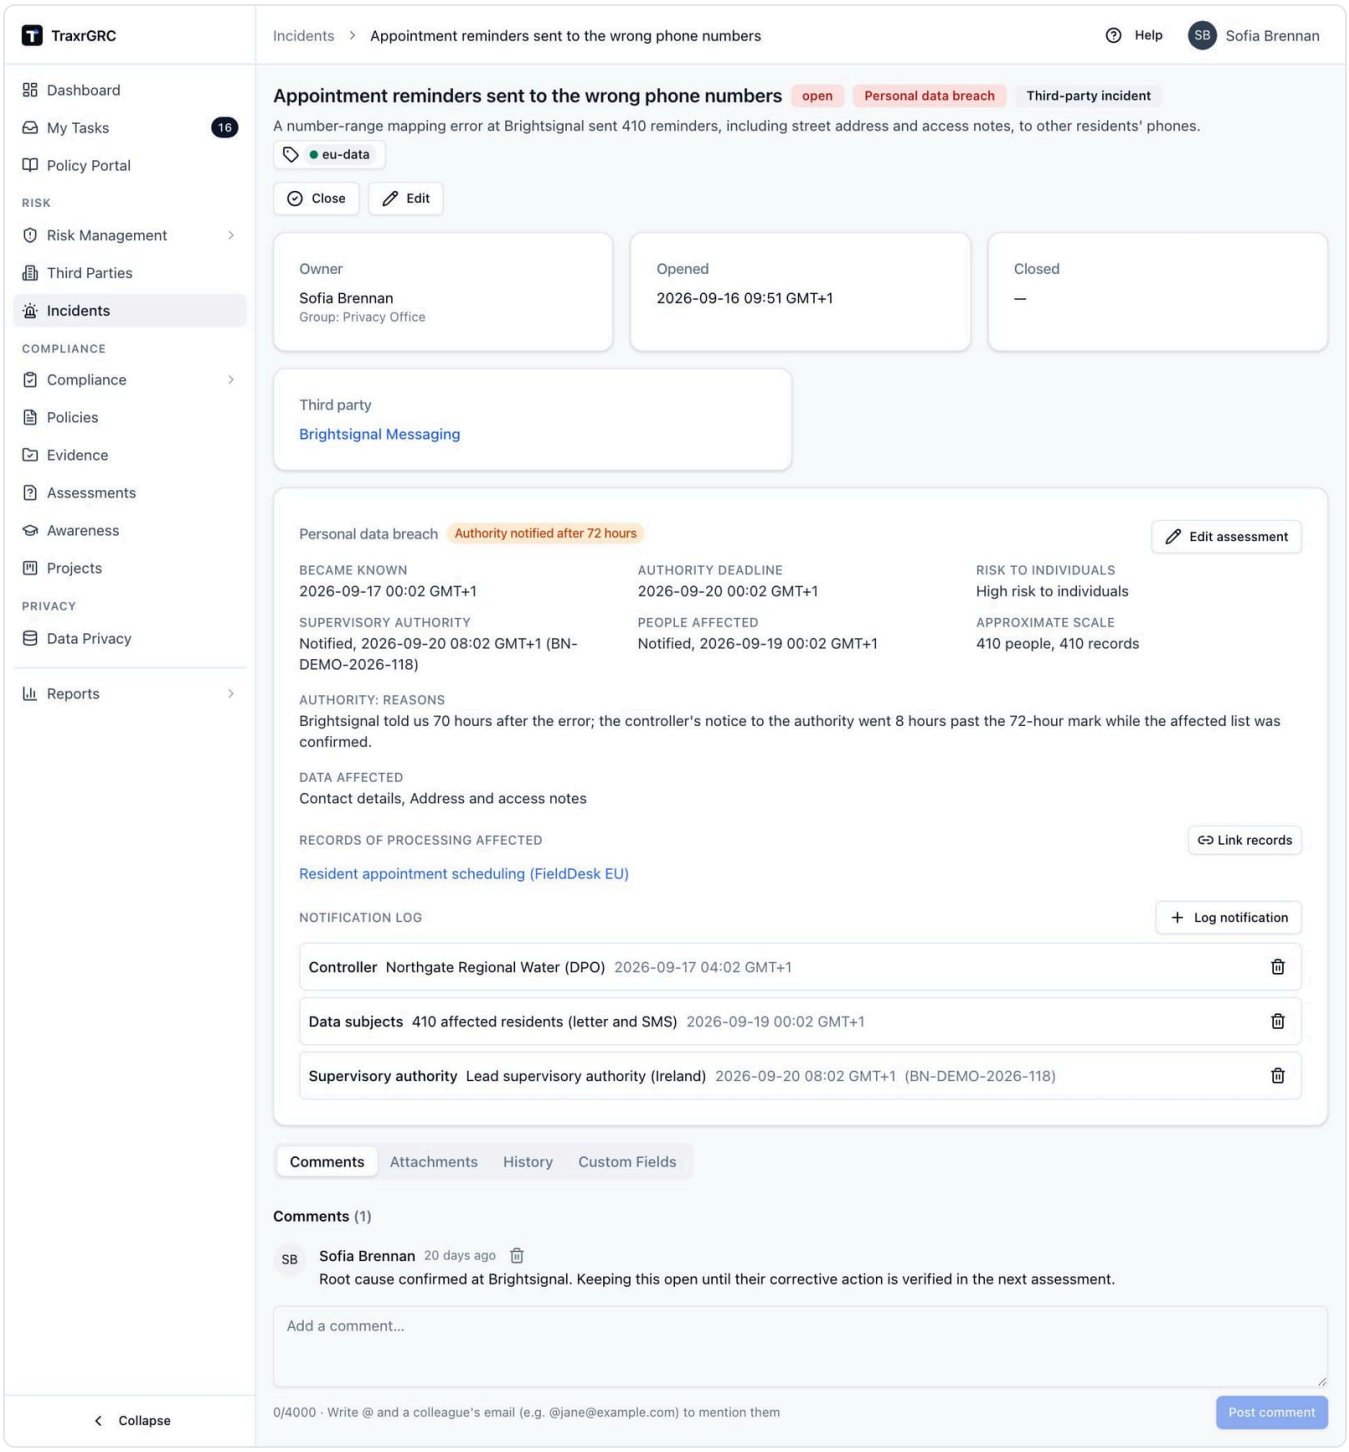


Figure 7.9. A breach notified to the authority after the 72-hour deadline, with the reasons for the delay and a notification log of the controller, authority and residents.

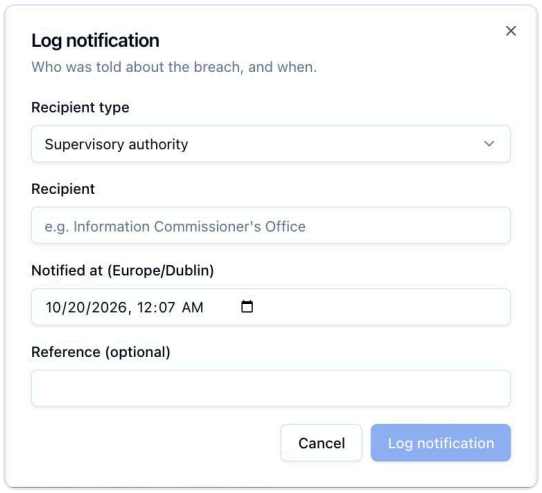


Figure 7.10. The Log notification dialog, recording the recipient type, recipient, time and reference of a notification.

Open breaches also appear on the **Personal data breaches** dashboard widget.

7.8 Editing and deleting

- Click **Edit** to change the title, description, classification, asset, third party, owner or owner group, then click **Save changes**. Editing does not change the status; only **Close** and **Reopen** do.
- Open **More actions** (**...**), click **Delete** and confirm with **Delete** to remove the incident permanently. Deleting also removes its comments, attachments and labels and cannot be undone. Deleting requires `incidents:delete`.

7.9 Comments, attachments and history

The **Comments**, **Attachments** and **History** tabs sit at the bottom of the page, followed by **Custom Fields**.

- **Comments.** Write in the box (up to 4,000 characters) and click **Post comment**. Write @ followed by a colleague's email address to mention them. You can delete your own comments; administrators can delete any.
- **Attachments.** Click **Upload file** to add evidence such as screenshots or logs; files up to 25 MB are accepted. Click a file name to download it. You can delete a file you uploaded; administrators can delete any.
- **History.** Every change to the incident (create, update, close, reopen, breach assessment and notification entries) as an audit log entry: the action, who made it, when, and which fields changed. This mirrors what administrators see under the audit log.

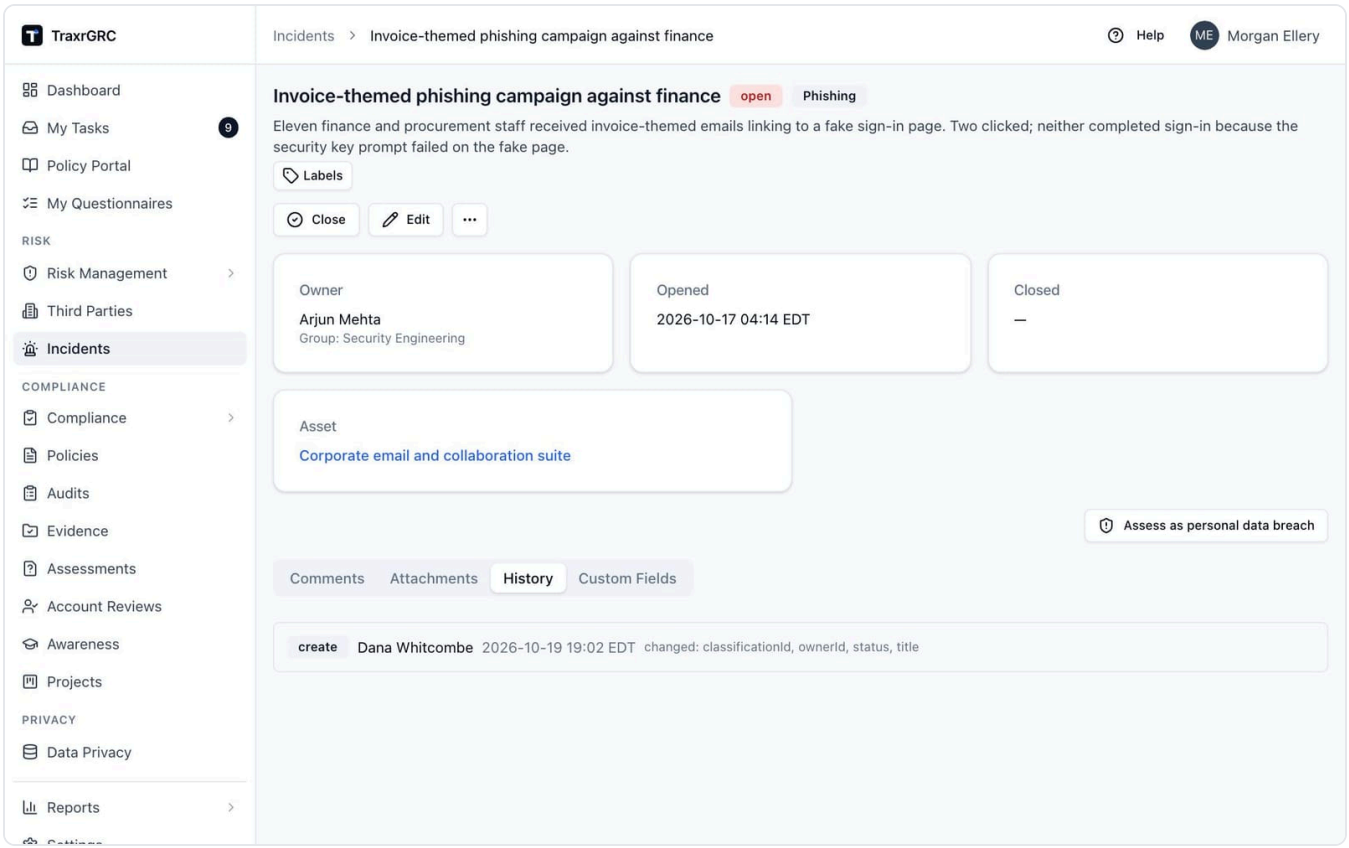


Figure 7.11. The History tab of an incident, listing each audit entry with the action, the person, the time and the fields changed.

Commenting and uploading require `incidents:update` on the incident.

The **Custom Fields** tab, last in the row, shows any custom fields your organization has defined for incidents, and lets you fill them in when you have `incidents:update` on the record (see Custom Fields).

CHAPTER 8

Internal Audits

Internal Audits tracks your audit program: what is planned, what is under way, and what each audit found. Each audit plan carries its schedule and a list of **findings**, the gaps the audit identified, and every finding can be linked to a remediation project so that follow-up work is visible in one place. Plans move through planned, in progress and completed; findings move between open, in remediation and closed.

Internal auditors and the GRC team create plans and record findings; finding owners move their findings through remediation; the CISO and audit committee read the results.

8.1 Who can do what

Access is controlled by the **audits** module rights. Administrators bypass all of them.

Action	Permission
See the audits list and plan pages	`audits:read`
Create an audit plan	`audits:create`
Edit a plan, change its status, add, edit and transition findings, comment, upload, edit custom fields	`audits:update`, "all" for every plan or "own" for plans you or your groups own
Delete an audit plan	`audits:delete` (all or own)
Create a remediation project from a finding	`audits:update` on the finding **and** `projects:create`
Import audit findings from CSV	`audits:create` or `audits:update`

8.2 The audit plans list

The **Audits** page lists every plan you can read. Click a column heading to sort; the **Status** column has a filter dropdown.

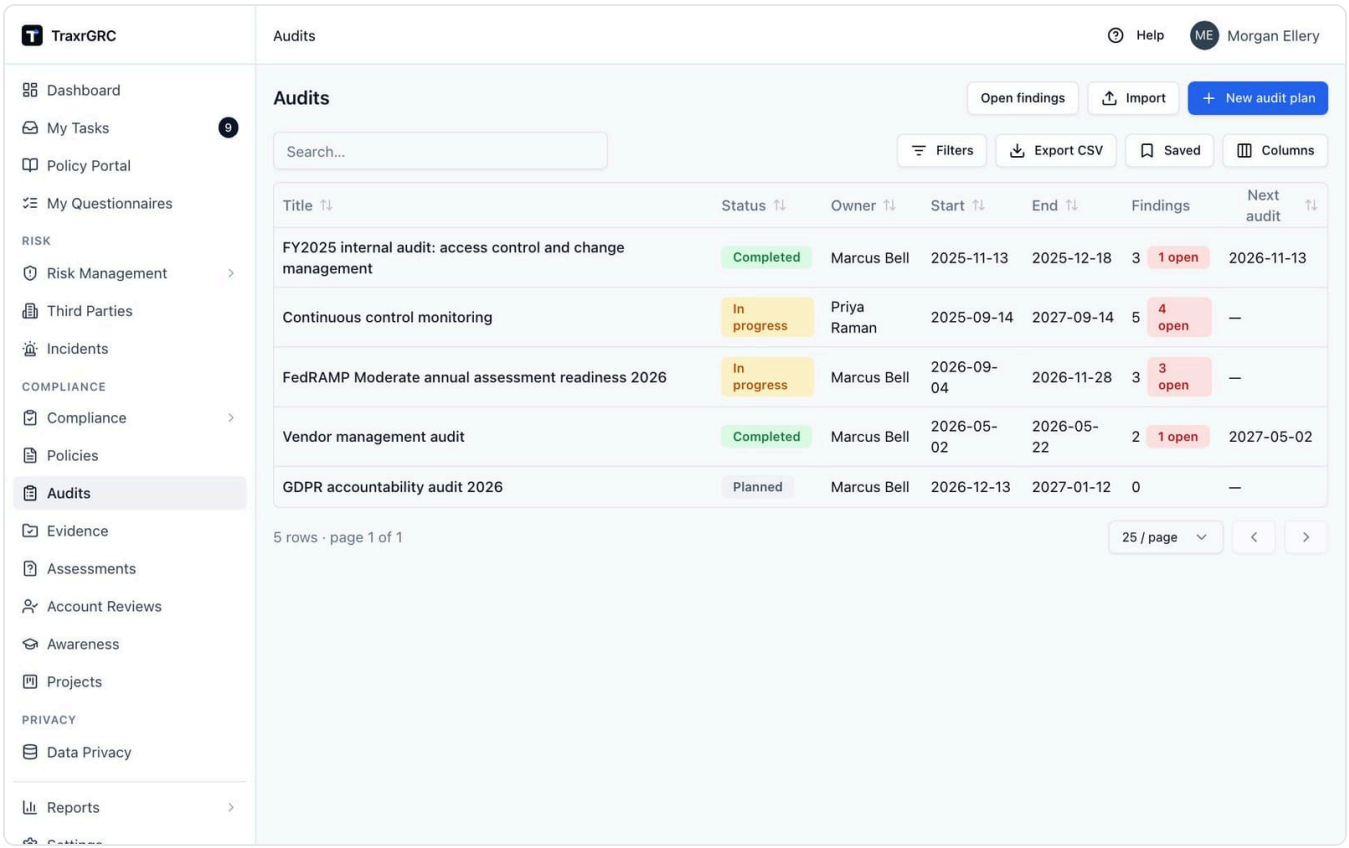


Figure 8.1. The audit plans list, showing each plan's status, owner, audit window, finding counts with open findings in red, and next audit date.

Column	What it shows
Title	The plan's title. Click it to open the plan.
Status	Planned, In progress or Completed (see below).
Owner	The owner user, or the owner group when no user is set.
Start / End	The audit window.
Findings	Total number of findings. When some are not yet closed, a red badge shows how many are open (findings in *Open* or *In remediation*).
Next audit	The scheduled date of the next audit. Shown in red when it lies in the past and the plan is **Completed**, a signal that the next cycle is overdue.

Plan statuses:

Status	Badge	Meaning
Planned	Gray	Scheduled but not started. New plans start here.
In progress	Amber	The audit is under way.
Completed	Green	The audit is done.

The **Import** button beside **New audit plan** creates or updates audit findings in bulk from a CSV file, each row naming its audit plan (see Importing from CSV under Getting Started). Plans themselves are created in the app.

8.3 Creating an audit plan

1. On the Audits page, click **New audit plan**.
2. Complete the dialog (only the title is required):

Field	What it means
Title	Name of the audit, e.g. "ISO 27001 internal audit, 2026".
Description	Scope, domains covered, reference standards.
Owner	The user running the audit.
Owner group	The team responsible for the audit.
Start date	When fieldwork begins.
End date	When fieldwork ends.
Next audit date	When the next cycle of this audit is due; drives the red overdue highlight on the list.

3. Click **Create audit plan**.

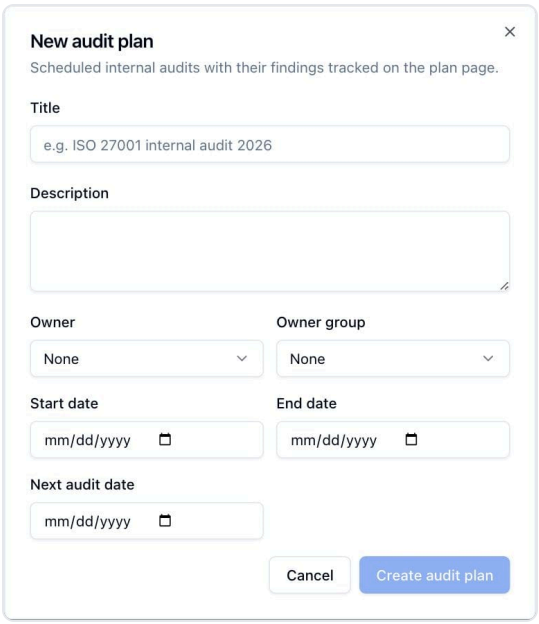


Figure 8.2. The New audit plan dialog with title, description, owner, owner group and the start, end and next audit dates.

The plan is created with status **Planned**, and the creation is recorded in the audit log.

8.4 The plan detail page

Click a plan's title to open it.

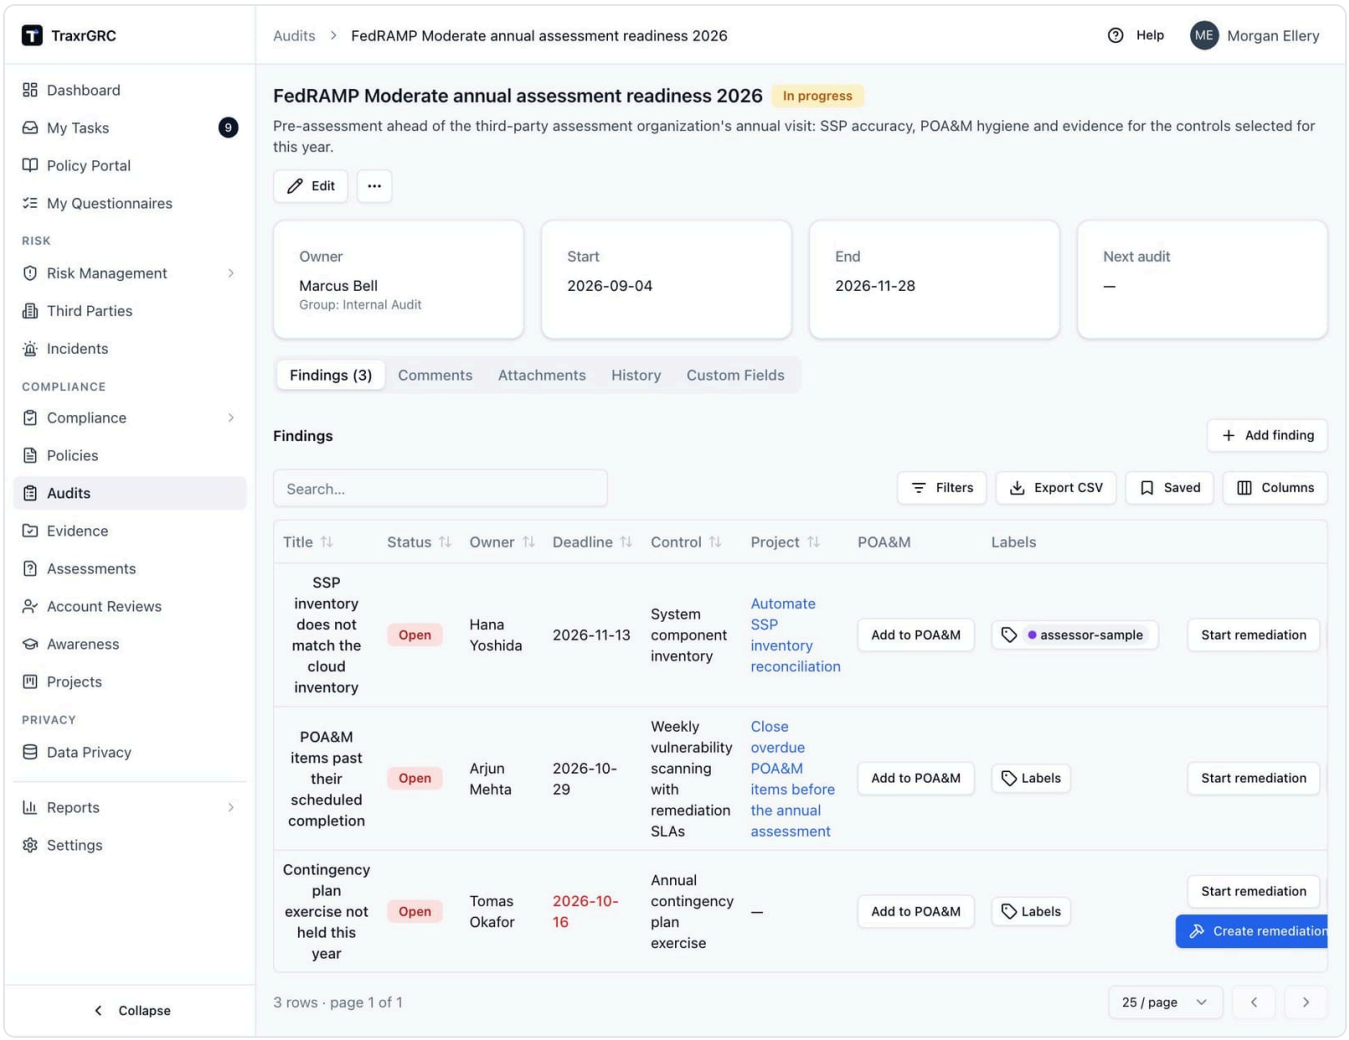


Figure 8.3. An in-progress audit plan: the status badge, the Owner, Start, End and Next audit cards, and the Findings tab listing each finding with its status, deadline, control, project and action buttons.

The page shows the title, status badge and description, the **Edit** button and the **More actions** menu (**…**, with **Delete**), and cards for **Owner**, **Start**, **End** and **Next audit**. Below them are the **Findings (*n*)**, **Comments**, **Attachments**, **History** and **Custom Fields** tabs.

8.4.1 Editing a plan and changing its status

1. Click **Edit**.
2. Change the title, description, owner, owner group or dates, or pick a new **Status**. The status field appears only when editing.
3. Click **Save changes**.

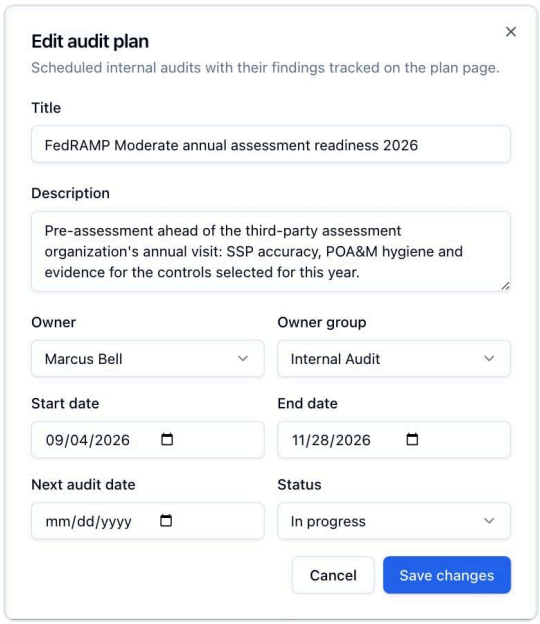


Figure 8.4. The Edit audit plan dialog, with the Status field that moves the plan forward.

Status moves strictly forward. The allowed transitions are:

1. **Planned → In progress**: start the audit.
2. **In progress → Completed**: finish it.

There is no way back: a completed plan cannot be reopened, and a planned plan cannot be marked complete without first moving it to in progress. If you try, the save fails with a message such as *Cannot move from "planned" to "completed"*. Dates and other details can be edited at any stage.

8.4.2 Deleting a plan

Open **More actions** (**...**), click **Delete** and confirm with **Delete**. This deletes the plan **together with all its findings** and their comments, attachments and labels. Remediation projects linked to findings are kept: they stay in Projects but lose the link to the deleted finding. Deleting requires `audits:delete` and cannot be undone.

8.5 Working with findings

Open the **Findings** tab on a plan (it opens by default). The tab label shows the total count, and users with `audits:update` see an **Add finding** button.

8.5.1 Adding or editing a finding

1. Click **Add finding**, or click a finding's title in the table to edit it.
2. Complete the dialog (only the title is required):

Field	What it means
Title	Short name for the gap found, e.g. "Privileged access review gaps".
Description	Details: what was found, the evidence, the risk.
Owner	The user responsible for remediation.
Owner group	The team responsible for remediation.
Deadline	Target date for closing the finding.
Linked control	The control from your control catalog that the finding relates to, or *None*.

3. Click **Add finding** (for a new finding) or **Save changes**.

Editing a finding's title, description, owner, linked control or status is for whoever can update its audit plan. A finding's owner (or a member of its owner group) who cannot update the plan sees the finding's details as text, with one change open to them: **Deadline**. Moving it asks for a reason, and the move and the reason are posted to the finding's comments.

Comments. An existing finding's dialog ends with its comment thread. Whoever can update the plan and the finding's owner can comment; anyone who can read the finding reads the thread. Comments the app records (a repeat audit result, a deadline move, the POA&M item it is tracked on) are marked **Recorded by the app** and cannot be deleted.

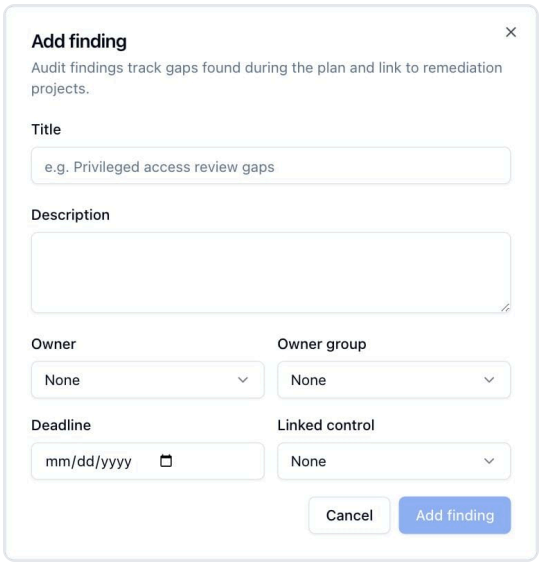


Figure 8.5. The Add finding dialog with title, description, owner, owner group, deadline and linked control.

New findings start with status **Open**.

The findings table shows **Title**, **Status** (filterable), **Owner**, **Deadline**, the linked **Control**, the linked remediation **Project** (a link, when there is one), **POA&M**, **Labels** and the action buttons. A

Requirements column is hidden by default; turn it on from the table's column choices. It lists the compliance requirements the finding's control is mapped to: the first three codes as links to the requirement, then "+N more"; requirements you cannot read are counted ("1 you cannot see"). The list is worked out each time from the control's current mappings. A deadline in the past is shown in red while the finding is not closed. Findings have no page of their own, so their labels are attached in the table: with `audits:update` on the finding, its **Labels** cell is the label picker (see Labels in Getting Started); otherwise it shows the labels read-only. Type part of a label name in the column's filter to list only findings carrying it.

8.5.2 Finding status workflow

Each finding row carries buttons for the moves allowed from its current status. You see them on a finding you can update: anyone who can update the plan sees them on every finding, and a finding's owner (or a member of its owner group, with `audits:update` at the own level) sees them on their own finding, except **Close**: closing verifies the fix, so it is for those who can update the plan. **Create remediation project** is also theirs. Clicking a finding's title opens it for editing for those who can update the plan; its owner sees the details with the deadline to move, and anyone else sees them read-only. Closing or reopening a finding also closes or reopens its POA&M item (see The POA&M register).

Current status (badge)	Button	Moves to
Open (red)	Start remediation	In remediation
Open (red)	Close	Closed
In remediation (amber)	Reopen	Open
In remediation (amber)	Close	Closed
Closed (gray)	Reopen	Open

Moving back is always allowed: a finding closed too early can be reopened, and one in remediation can be returned to open. Every transition is written to the audit log against the finding.

The app has no button for deleting a finding; findings are removed with their plan, or individually through the REST API. A finding cannot be deleted while a remediation project is linked to it; unlink the project first (see Projects).

8.5.3 Creating a remediation project

Findings with no linked project offer a **Create remediation project** button. Clicking it:

1. Creates a new project in the Projects module, pre-filled with the finding's title, description, owner and owner group, with status **In progress** and a link back to the finding.

2. Links the finding to that project, so the finding row shows the project and the project shows the finding.
3. Takes you straight to the new project's page.

This requires `audits:update` on the finding **and** `projects:create`; if you lack either, the action fails with a permission message. A finding can have at most one linked project, so the button is hidden once a link exists.

8.5.4 Findings raised by failed control audits

A user with `settings:update` can name an audit plan that receives a finding every time a control audit fails (**Controls → Settings**, see Risk Management). Each such finding is titled *Failed control audit: <control name>*, is owned by the control owner, is due by the re-audit date and is linked to the control. It is worked like any other finding. A control has at most one open finding of this kind: later failures and a passing re-audit are noted on it (see Control testing under Risk Management).

Because the finding names its control, it reaches the requirements that control is mapped to: the requirement page lists it under **Open audit findings on mapped controls**, and **Add to POA&M** fills in the requirement codes (see The POA&M register). Nothing opens a POA&M item for it automatically; a person decides, and until then the finding is flagged.

8.6 Open findings

Open findings (the button at the top of the Audits page) lists every open or in-remediation finding across the audit plans you can read, soonest deadline first. Each row shows the finding (a link to its plan), the **Audit plan**, **Status**, **Needs attention**, **Owner**, **Deadline**, **Control**, **Requirements** and the POA&M item or remediation project answering it, with **Add to POA&M** when there is neither and you hold `compliance:create`.

A finding that needs attention shows one badge, for its most urgent reason:

Reason	When
Overdue	Its deadline has passed.
Not tracked	It is not overdue, and neither a POA&M item nor a remediation project answers it.

Filter **Needs attention** on one reason, or on **Needs attention (any reason)** for the list to work through. The **Requirements** column filters on **Affects requirements** or **Affects requirements, no POA&M item** (its control is mapped to at least one requirement, counted even when you cannot read it). The **Findings to resolve** dashboard widget links here with the Needs attention filter set, and the notice on the POA&M page with the Requirements filter set.

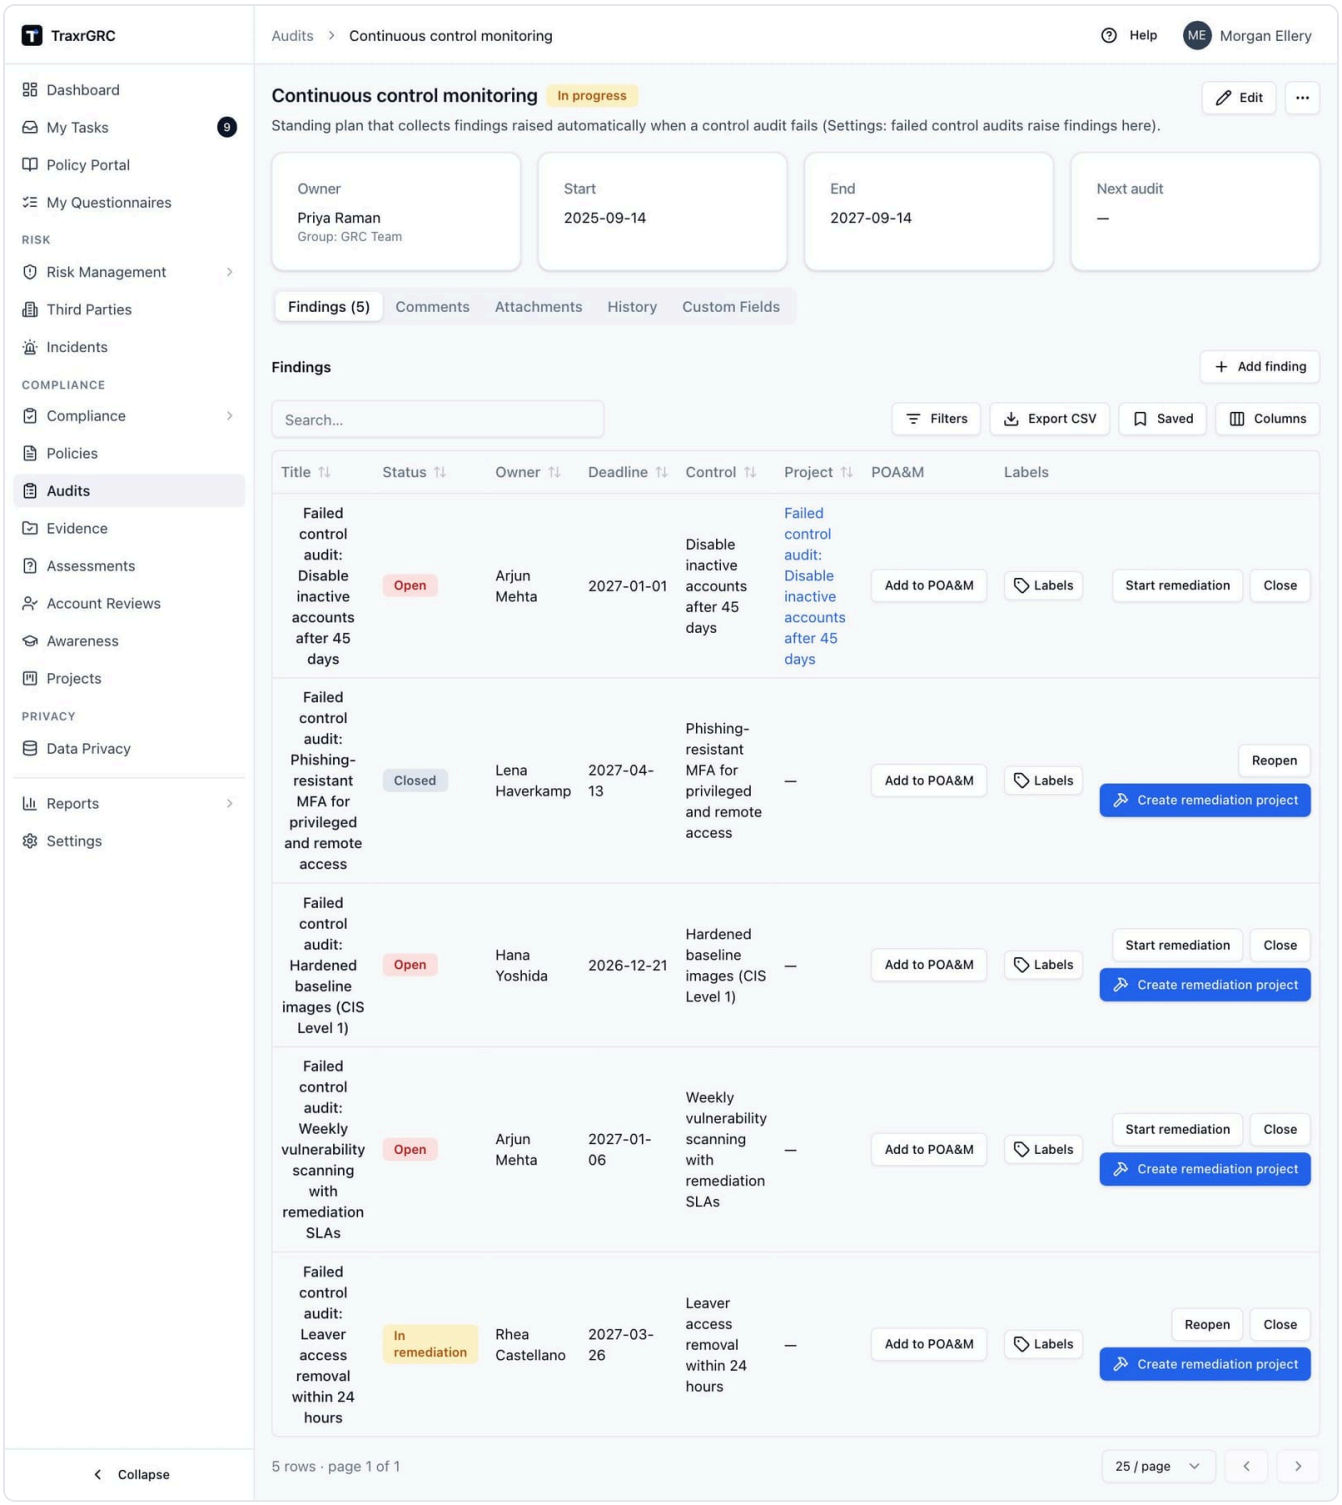


Figure 8.6. A standing monitoring plan holding findings raised by failed control audits, one of them already linked to a remediation project.

8.7 Comments, attachments and history

The plan's **Comments**, **Attachments** and **History** tabs work like the incident ones: comments up to 4,000 characters, file uploads up to 25 MB with download links, and the audit trail of changes to the plan (create, edits and status changes, delete). Finding changes (creation, edits, status transitions, remediation projects) are recorded in the audit log against each finding, where administrators see them under Settings → Audit Log. Commenting and uploading require `audits:update` on the plan.

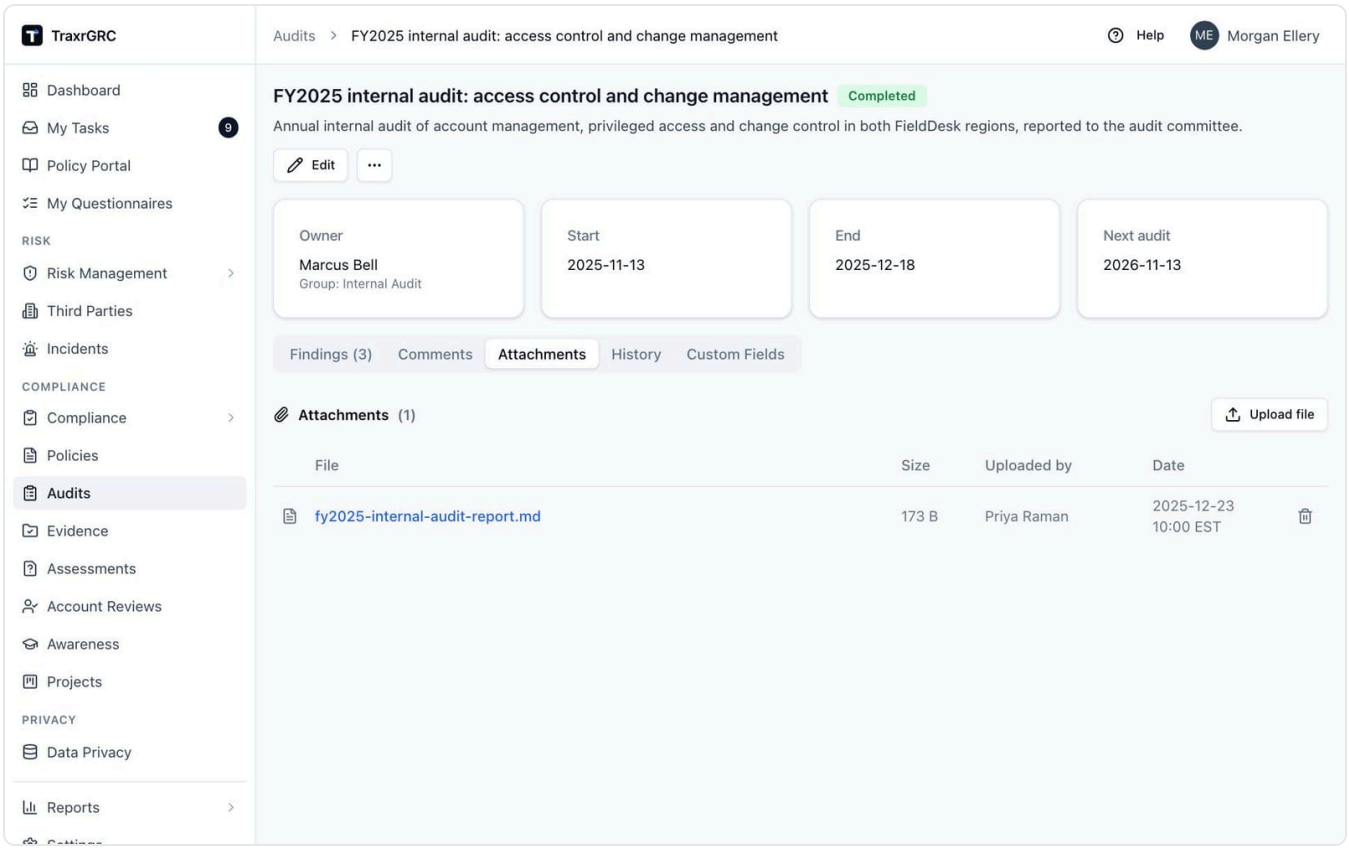


Figure 8.7. The Attachments tab of a completed audit plan, holding the audit report.

The **Custom Fields** tab, last in the row, holds organization-defined fields for audit plans (see Custom Fields).

CHAPTER 9

Projects

Projects tracks improvement and remediation work, from rolling out phishing-resistant MFA to fixing what an audit found. A project can stand alone, be linked to an audit finding, and be linked to the compliance requirements whose gaps it closes (see The gap view under Compliance Management), so remediation stays connected to the evidence. Completing a project stamps its completion date, giving you a clean record of when the work landed.

Project owners keep their projects current; the GRC team and the CISO use the list to follow remediation against deadlines.

9.1 Who can do what

Access is controlled by the **projects** module rights. Administrators bypass all of them.

Action	Permission
See the projects list and detail pages	`projects:read`
Create a project (including from a finding)	`projects:create`
Edit, change status, unlink a finding, comment, upload, edit custom fields	`projects:update`, "all" or "own" (you or your groups own the project)
See linked compliance requirements	`compliance:read`
Link or unlink compliance requirements	`projects:update` on the project and `compliance:read` on each requirement newly linked
Delete a project	`projects:delete` (all or own)
Import projects from CSV	`projects:create` or `projects:update`

9.2 The projects list

The **Projects** page lists every project you can read. Click a column heading to sort; the **Status** column has a filter dropdown.

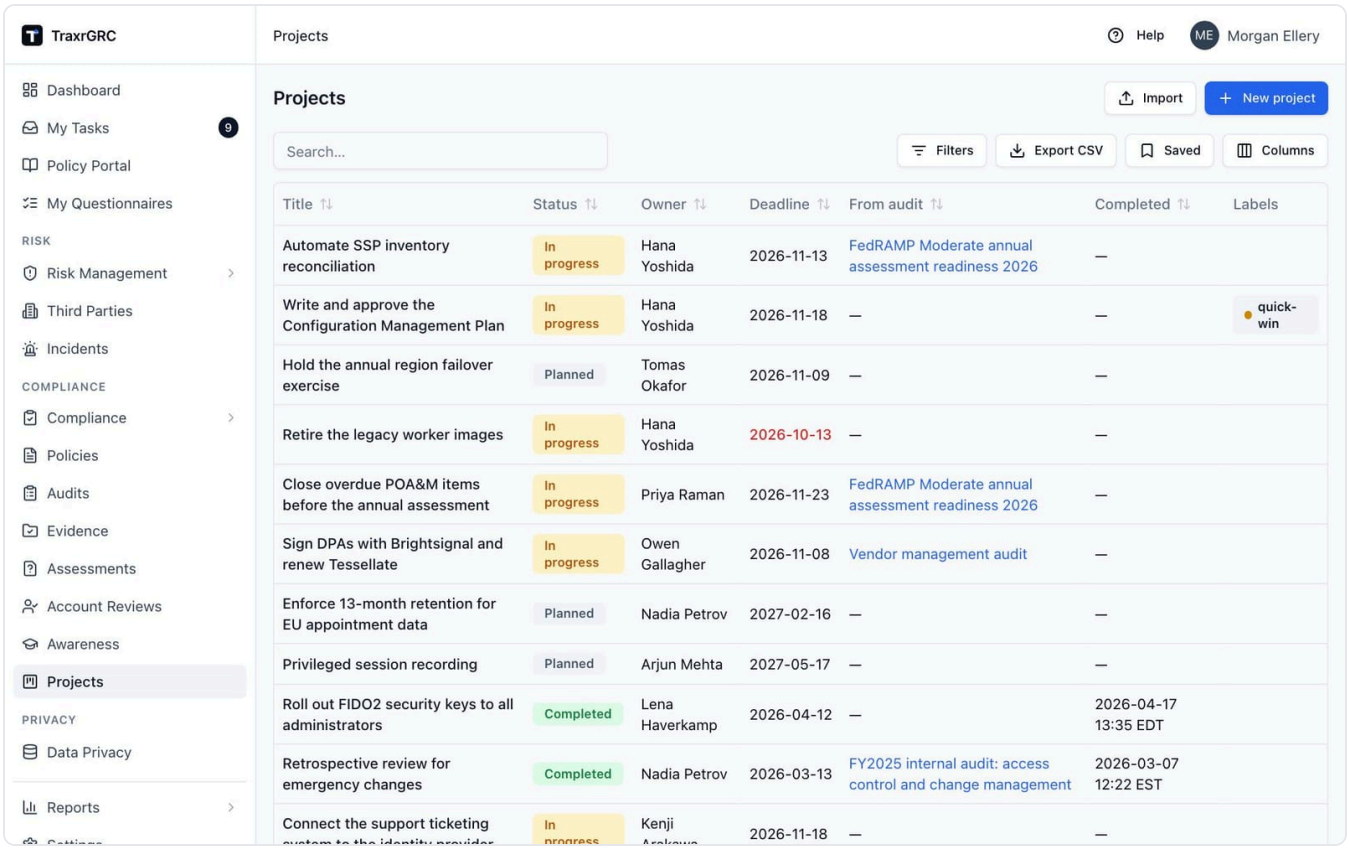


Figure 9.1. The projects list, showing each project's status, owner, deadline (overdue ones in red), the audit plan it came from and its completion time.

Column	What it shows
Title	The project's title. Click it to open the project.
Status	planned, in_progress or completed (see below).
Owner	The owner user, or the owner group when no user is set.
Deadline	Target date. Shown in red when it lies in the past and the project is not completed.
From audit	For a project linked to an audit finding: the title of the audit plan the finding belongs to. Click it to open that plan. Without read access to the plan the cell says only "Audit finding".
Completed	When the project was completed, or a dash.
Labels	The project's labels. Filter by typing part of a label name.

Project statuses:

Status	Badge	Meaning
planned	Gray	The work is agreed but not started. New projects start here.
in_progress	Amber	The work is under way. Projects created from a finding start here.
completed	Green	The work is done; the completion time is stamped.

The **Import** button beside **New project** creates or updates projects in bulk from a CSV file (see Importing from CSV under Getting Started).

9.3 Creating a project

1. On the Projects page, click **New project**.
2. Complete the dialog (only the title is required):

Field	What it means
Title	Name of the work, e.g. "Deploy phishing-resistant MFA".
Description	What the project covers and why.
Owner	The user responsible for delivery.
Owner group	The team responsible for delivery.
Deadline	Target completion date.

3. Click **Create project**.

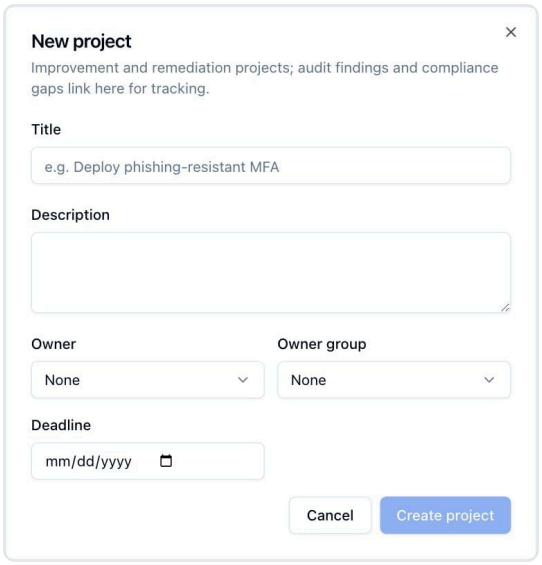


Figure 9.2. The New project dialog with title, description, owner, owner group and deadline.

The project starts as **planned**. The create dialog does not offer a finding link. To track remediation for an audit finding, create the project from the finding instead (see Internal Audits): it arrives pre-filled and linked in both directions. Requirements are linked from the project page once it exists.

9.4 The project detail page

Click a project's title to open it.

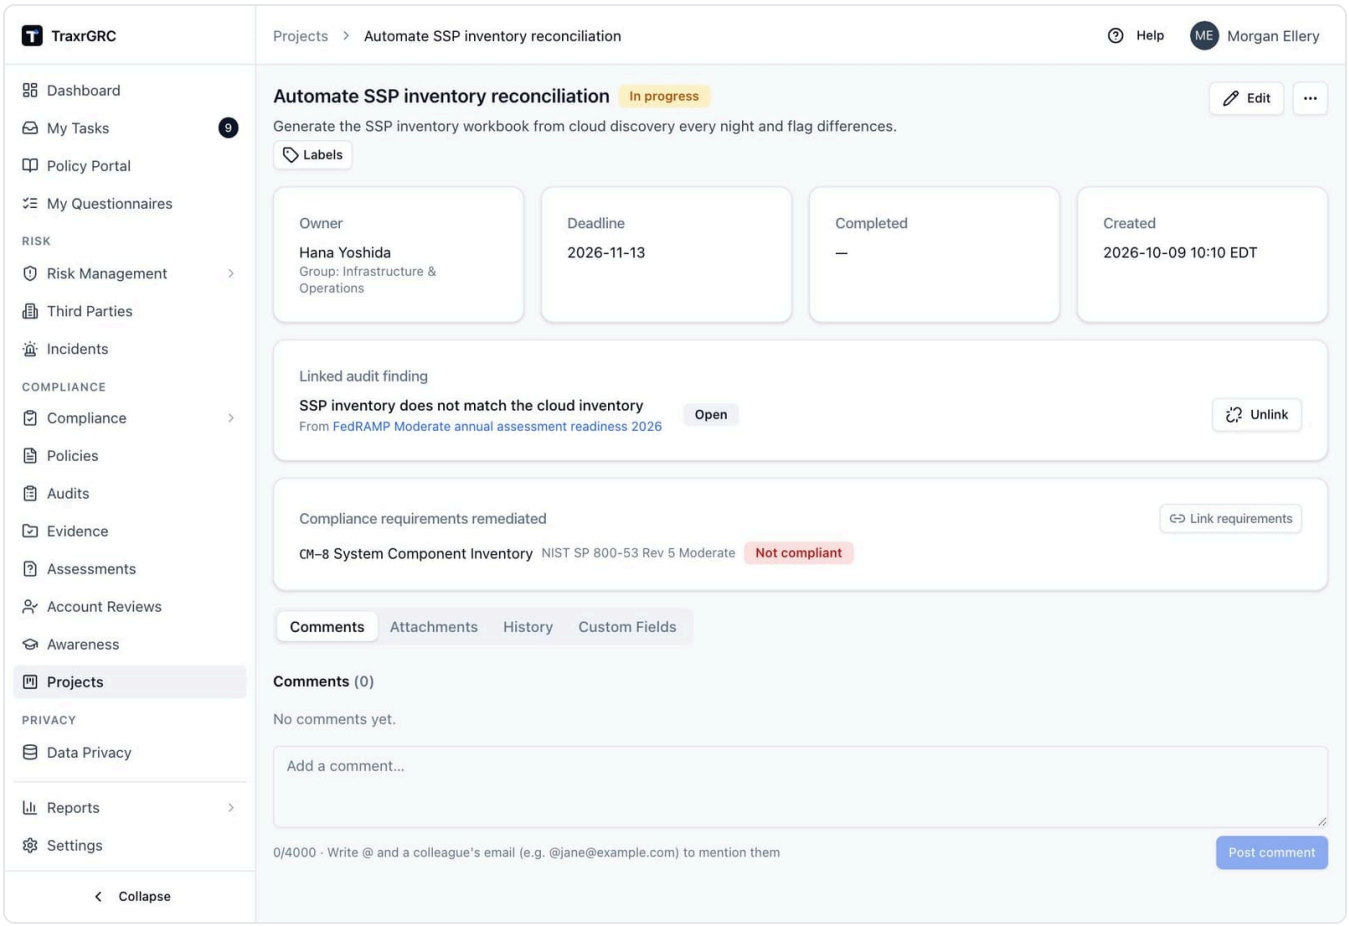


Figure 9.3. A remediation project linked to an audit finding and to the compliance requirement it closes, with the Owner, Deadline, Completed and Created cards.

The page shows the title, status badge and description, the **Labels** picker (see Labels in Getting Started; attaching needs projects:update on the project), the **Edit** button and the **More actions** menu (**...**, with **Delete**), cards for **Owner**, **Deadline**, **Completed** and **Created**, the **Linked audit finding** and **Compliance requirements remediated** cards when they apply, and the **Comments**, **Attachments**, **History** and **Custom Fields** tabs.

9.4.1 Changing status

1. Click **Edit**.
2. Pick a new **Status** (the field appears only when editing). Any of the three values is reachable from any other.
3. Click **Save changes**.

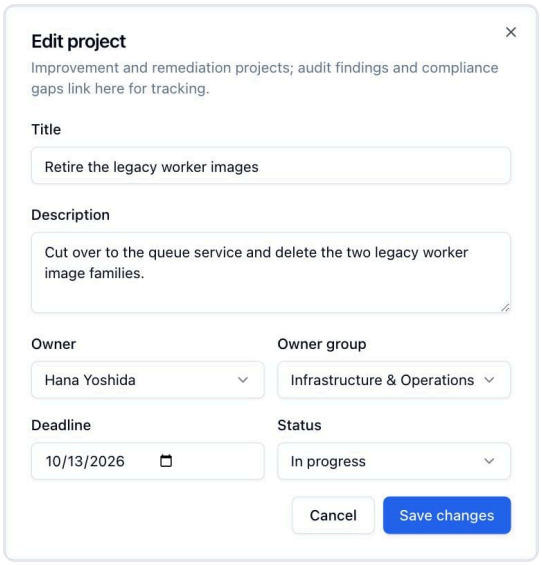


Figure 9.4. The Edit project dialog with its Status field.

Two things happen automatically:

1. Moving a project **to completed** stamps the **Completed** card with the current time.
2. Moving it **out of completed** (back to planned or in progress) clears that time.

The Completed time therefore always reflects the most recent time the project entered completed.

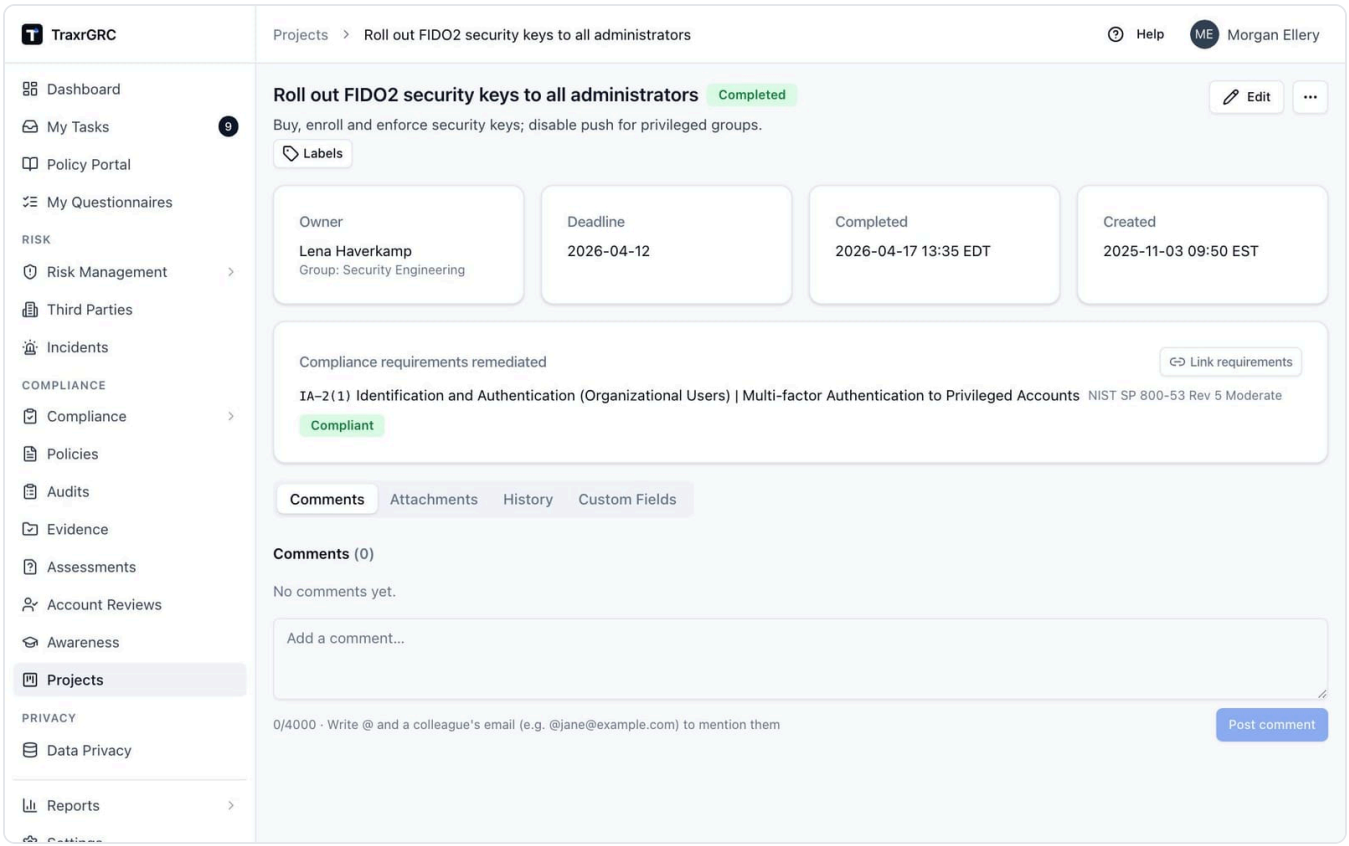


Figure 9.5. A completed project, with its completion time on the Completed card.

9.4.2 The linked audit finding

When the project was created from a finding (or linked to one through the REST API), a **Linked audit finding** card shows the finding's title and status, and the plan it came from ("From *plan*"; click the plan name to open it). The plan is shown only when you have read access to it.

To remove the link:

1. Click **Unlink** on the card (requires `projects:update`).
2. Confirm with **Unlink** in the **Unlink finding?** dialog.

Unlinking removes the connection only. The finding itself is unchanged, and it can later get a new remediation project. The same link is visible from the other side, in the **Project** column of the finding's row in its audit plan.

9.4.3 Compliance requirements remediated

With `compliance:read`, the **Compliance requirements remediated** card lists the requirements this project is linked to, each with its code, name, package and analysis status; click one to open it. Requirements you cannot read are counted ("*n* more you do not have access to"), not named, and stay linked.

To link or unlink requirements:

1. Click **Link requirements** (requires `projects:update`).
2. Search the checklist of requirements you can read (headings are left out).
3. Tick or untick a requirement. Each change saves straight away and is recorded on the history of both the project and the requirement.

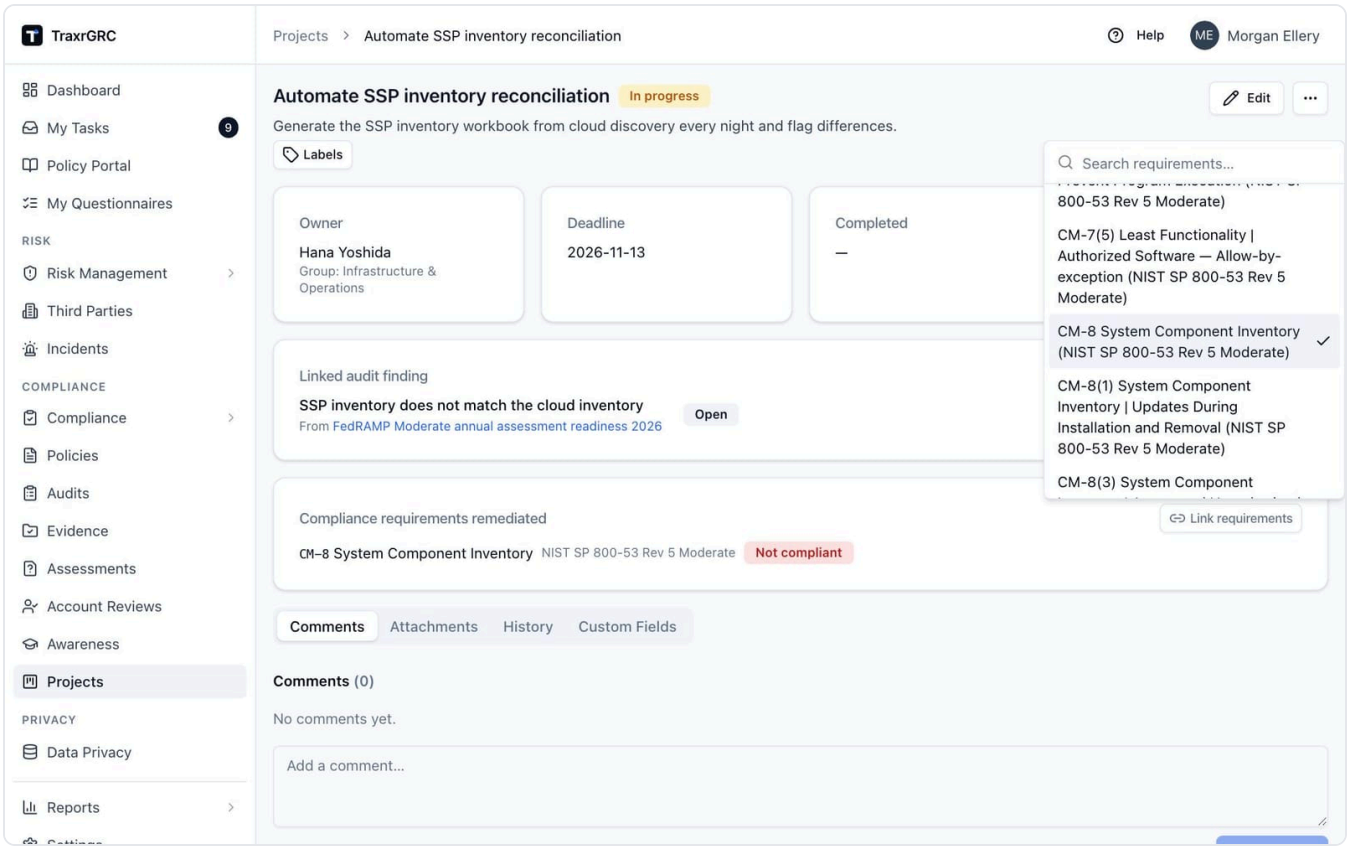


Figure 9.6. The Link requirements checklist, open on a project, for searching and ticking the requirements the project remediates.

The same link shows on the requirement's **Remediation projects** card and in the gap view.

9.4.4 Comments, attachments and history

These tabs work as described for incidents: comments up to 4,000 characters, uploads up to 25 MB with download links, and an audit trail of edits, status changes, finding unlinks and requirement links. Commenting and uploading require `projects:update` on the project.

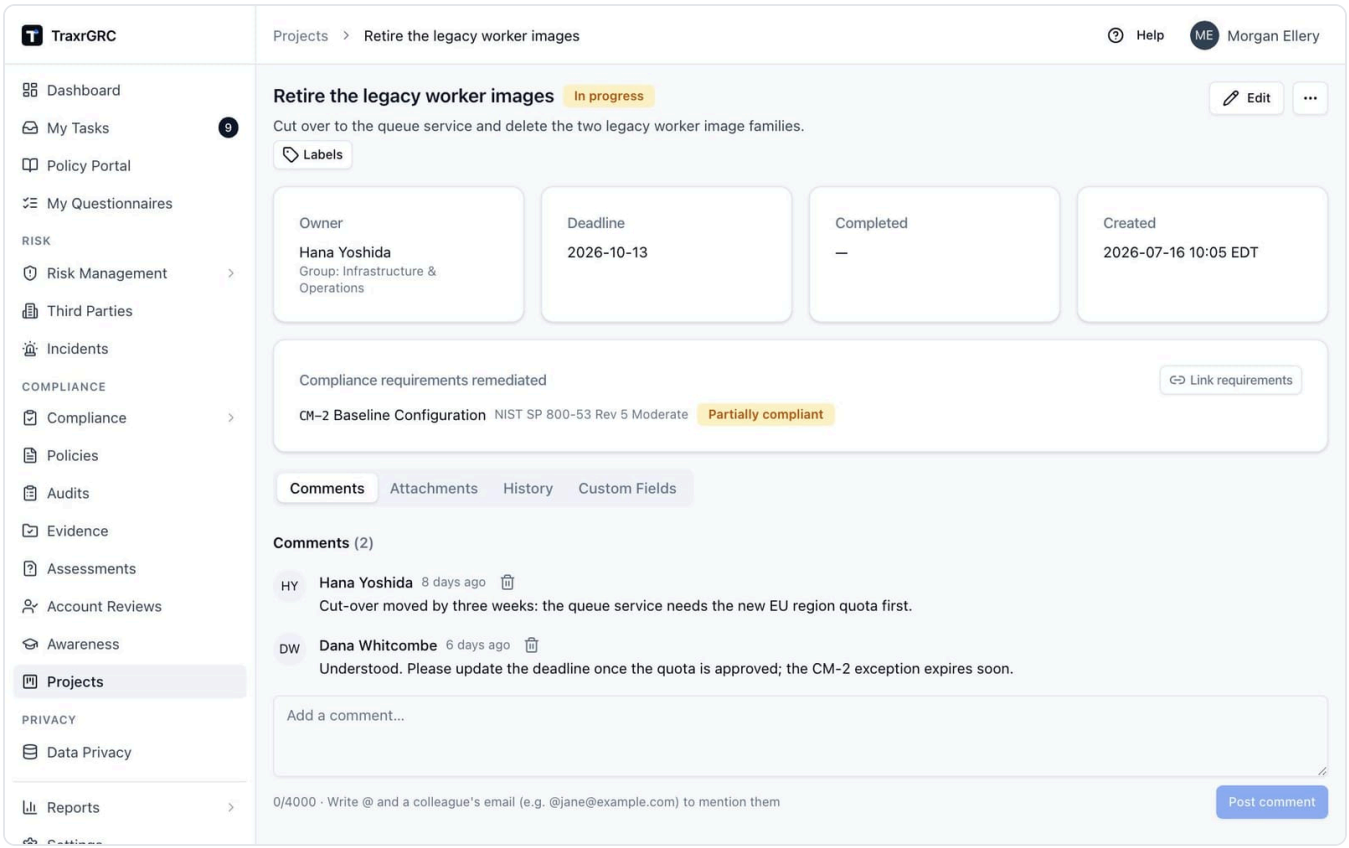


Figure 9.7. The Comments tab of an overdue project, with the owner and the CISO discussing a revised deadline.

The **Custom Fields** tab, last in the row, holds organization-defined fields for projects (see Custom Fields).

9.4.5 Deleting a project

Open **More actions** (**...**), click **Delete** and confirm with **Delete**. The project is removed together with its comments, attachments and labels. A linked audit finding is **not** deleted: it loses the link, so its remediation history stays in the audit record. Linked compliance requirements lose the link too, and each requirement's history records that the project was deleted. Deleting requires `projects:delete` and cannot be undone.

CHAPTER 10

Security Awareness

Security Awareness is where you build, deliver and track staff training programs, for example an annual security awareness course with a scored test at the end. A program pairs training material (written in Markdown) with a multiple-choice test. Everyone who takes the test gets a score, and the program reports who has passed, failed or not completed.

Two groups of people use this module. Program owners (typically the security or GRC team) create programs, write the material and the test, publish them and chase completion. Every other member of staff opens the module to read the material and take the test. This chapter covers both: administering programs first, then taking a test. Who can do what is controlled by the awareness module permissions, summarized at the end.

10.1 The programs list

Open **Awareness** in the navigation (under Compliance). The list shows one row per program:

Column	What it shows
Title	Click it to open the program.
Status	Draft, Active or Archived (see Program status below).
Owner	The named owner, or the owner group if no individual owner is set.
Due date	The completion deadline. Shown in **red with "(overdue)"** when the program is active and the date has passed.
Pass score	The percentage a taker must reach to pass.
Completion	How many active users have passed, out of all active users, for example `12/40 (30%)`.

The table can be searched and filtered by status, and **Export CSV** downloads the current view.

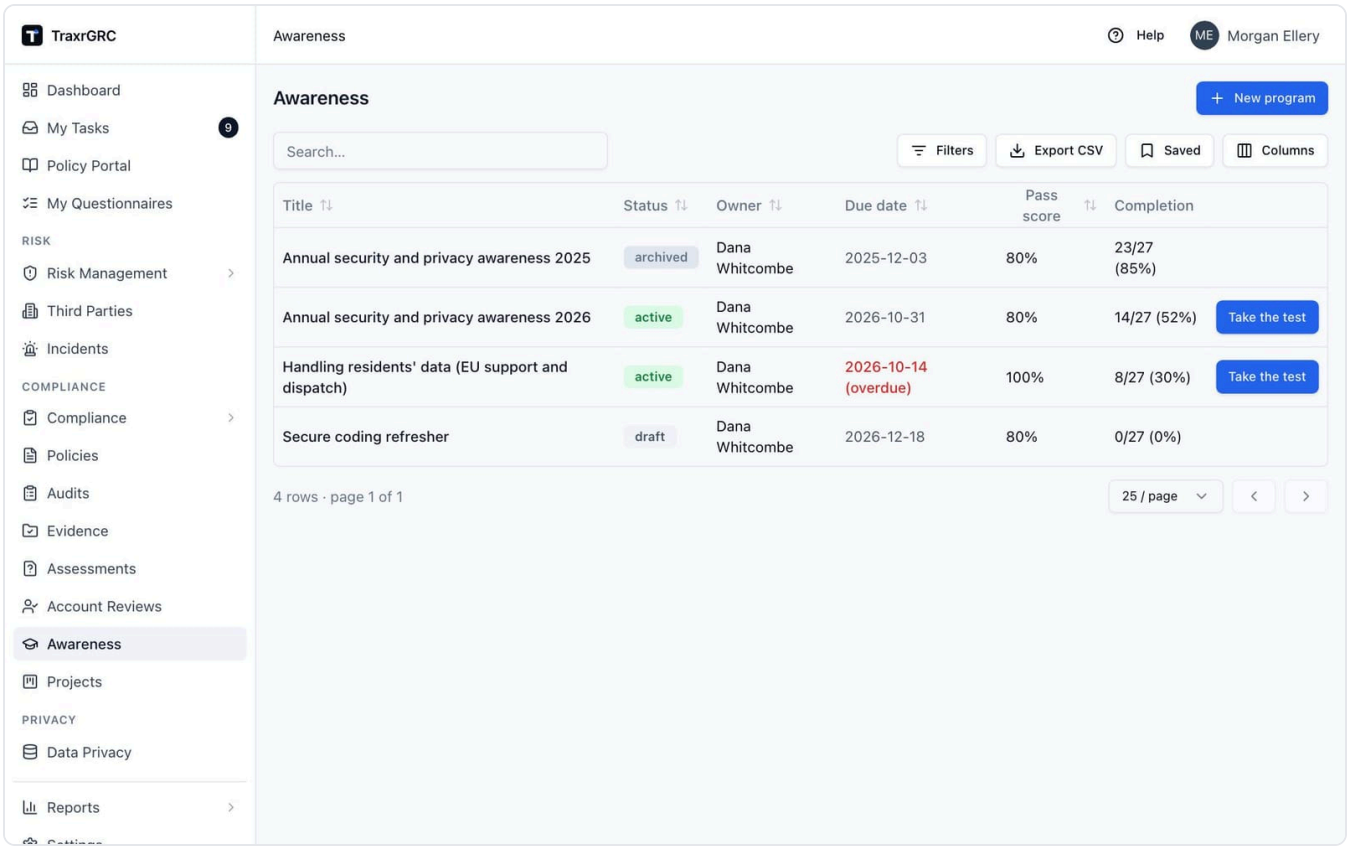


Figure 10.1. The Awareness programs list, with completion figures for each program and an overdue due date shown in red.

If a program is **active** and has a test, the last column shows a shortcut for you personally:

Shown	Meaning
Take the test	Opens the test (see Taking a test below). You need awareness:read.
Passed	You have already passed; no shortcut is offered.
Nothing	The program is a draft or archived, or it is active but has no test yet.

10.2 Creating a program

1. Click **New program** (requires awareness:create).
2. Complete the **New awareness program** dialog:

Field	What it means
Title	The program name, e.g. "Annual Security Awareness 2026". Required.
Description	Free text shown to takers under the title.
Owner	An individual user responsible for the program (optional).
Owner group	A group responsible for the program (optional). Used for "own"-level permission scoping.
Due date	The deadline by which staff should complete the program. Optional.
Pass score (%)	The score needed to pass, from 1 to 100. Defaults to 80.

3. Click **Create program**. The button is enabled once the title is filled in and the pass score is between 1 and 100. The program is created as a **Draft**.
4. Open the program from the list to add the training content and the test. The dialog only captures the basics.

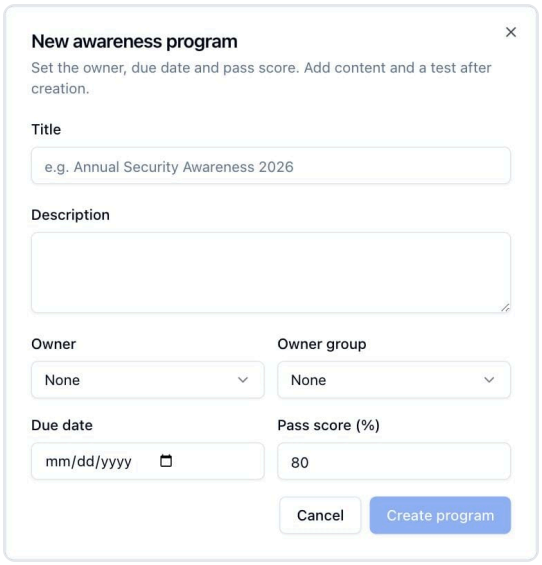


Figure 10.2. The New awareness program dialog with its title, owner, due date and pass score fields.

Every creation, change and deletion is recorded in the audit log and appears on the program's History tab.

10.3 Program status: draft, active, archived

A program moves through three statuses:

Status	Badge	Meaning
Draft	Gray	Being prepared. Not open for completion.
Active	Green	Published. Takers can open and complete the test; due-date reminders are sent.
Archived	Gray	Closed. Hidden from takers: the test can no longer be opened or submitted.

The allowed moves are:

From	To	Typical use
Draft	Active	Publish the program.
Draft	Archived	Shelve it before publishing.
Active	Archived	Close it, for example after the campaign ends.
Archived	Active or Draft	Bring it back.

To change status:

1. Open the program and click **Edit**.
2. Pick the new value in **Status**. The list offers only the moves allowed from the current status.
3. Click **Save changes**.

Archiving an active program immediately blocks any further test submissions, even mid-campaign.

10.4 The program detail page

Click a program title to open it. The header shows the title, the status badge and, if you have passed, a green **You passed** badge with your score. Four summary cards sit underneath:

Card	What it shows
Owner	The owner, with the owner group below it.
Due date	The deadline, in red with "(overdue)" when the program is active and the date has passed.
Pass score	The percentage needed to pass.
Completion	"*passed*/*total* passed (%)", counted over all active users.

The header buttons depend on the program and your permissions:

Button	When it appears
Take the test / **Retake the test**	The program is active and has a test. Shown to everyone with awareness:read; it reads **Retake the test** once you have a result.
Edit	You have awareness:update on the program. Edits the settings, including status.
Delete	You have awareness:delete on the program.

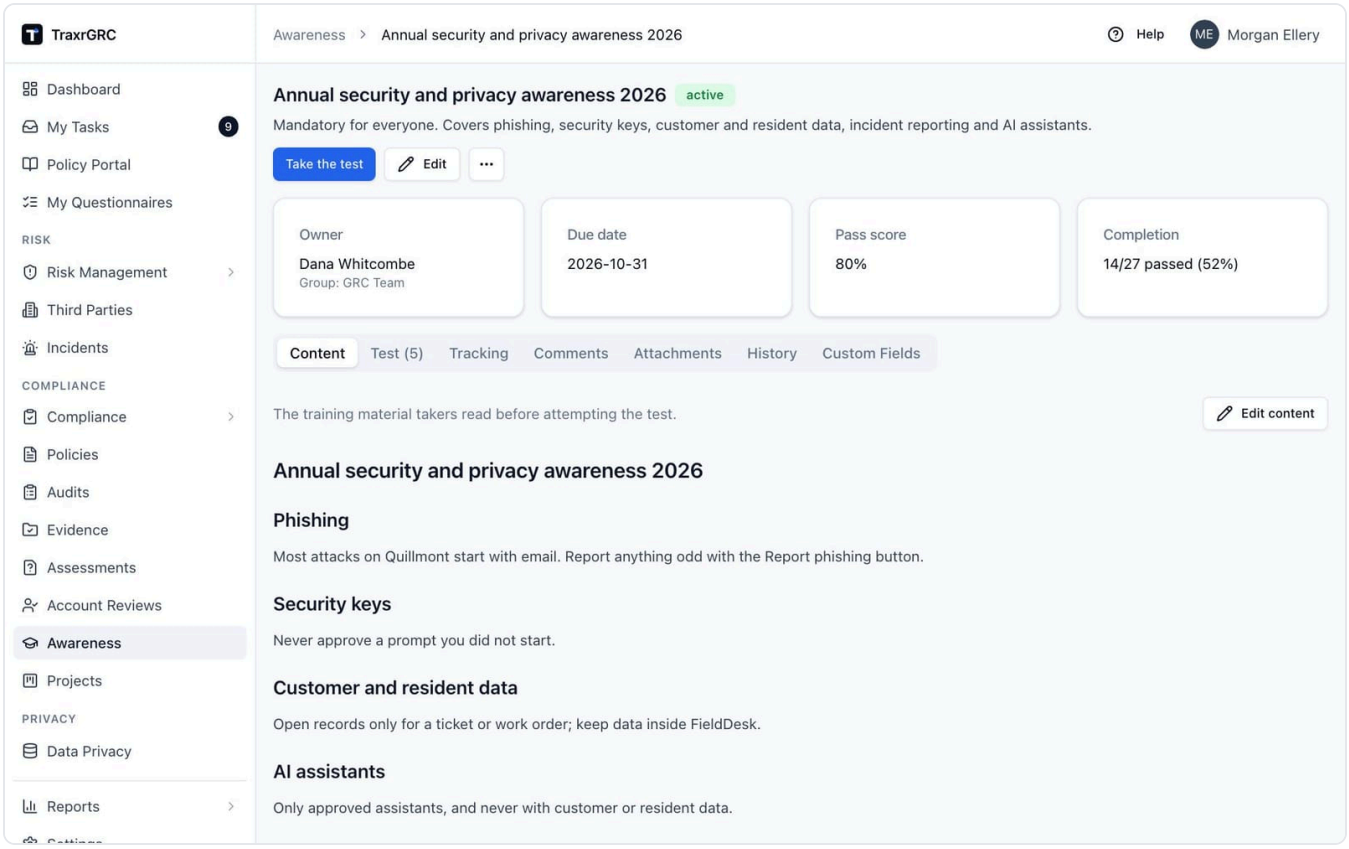


Figure 10.3. The detail page of an active awareness program, showing the summary cards and the training material on the Content tab.

The program is organized into these tabs:

Tab	Contents
Content	The training material takers read before the test.
Test (*N*)	The question editor; *N* is the number of questions.
Tracking	Per-person completion status and reminders.
Comments	Discussion about the program.
Attachments	Supporting files.
History	The audit trail for the program.
Custom Fields	Organization-specific fields (see Comments, attachments, history and custom fields below).

10.5 Content tab: the training material

The Content tab displays the program's training material as formatted text. A program with no material shows "No content yet."

To write or change the material (requires awareness:update):

1. Click **Edit content**.
2. Write the material in the **Write** view. Markdown formatting is supported, including tables and checklists, and the toolbar beside the **Write** and **Preview** buttons inserts headings, bold, italic, code, lists and a table (the same toolbar as the policy editor, see Writing content).
3. Switch to **Preview** at any time to check how it renders for takers.
4. Click **Save content**.

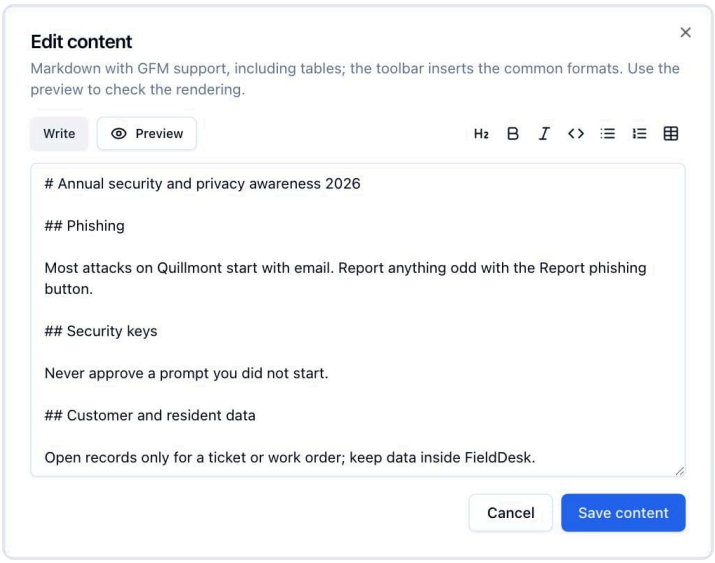


Figure 10.4. The Edit content dialog, with Write and Preview views of the Markdown training material.

Takers read this material before attempting the test, so write it first. The test should only cover what the material teaches.

10.6 Test tab: building the question set

The Test tab is the question editor. Each program has one test, and questions are shown to every taker in the order you define.

If you can update the program you see the full editor, including which choices are correct. Anyone else sees a read-only list of the questions and their choices without the answer key; the key is withheld through the API as well.

To build or change the test (requires awareness:update):

1. Click **Add question**. Repeat for each question (up to 100).
2. For each question, type the question text.
3. Each new question starts with two choices. Edit the choice texts, and click **Add choice** for more (up to 20 per question).
4. Tick the checkbox next to **every** correct choice. A question may have more than one correct choice; the editor reminds you that more than one is allowed.
5. Reorder questions with the up and down arrows, or remove one with its trash icon. A choice can be removed with its trash icon, but a question always keeps at least two choices.
6. Click **Save test**. The button is enabled only once you have made changes and the test has at least one question.

A test only saves if it meets all of these rules:

- At least one question.

- Each question has text.
- Each question has between 2 and 20 non-empty choices.
- Each question has at least one choice ticked as correct.

If results already exist, an amber warning above the editor states how many. Editing the test **does not affect stored results**: existing scores and passes are kept as they were, even if questions are reworded or removed afterward. People who take the edited test are scored against the new version.

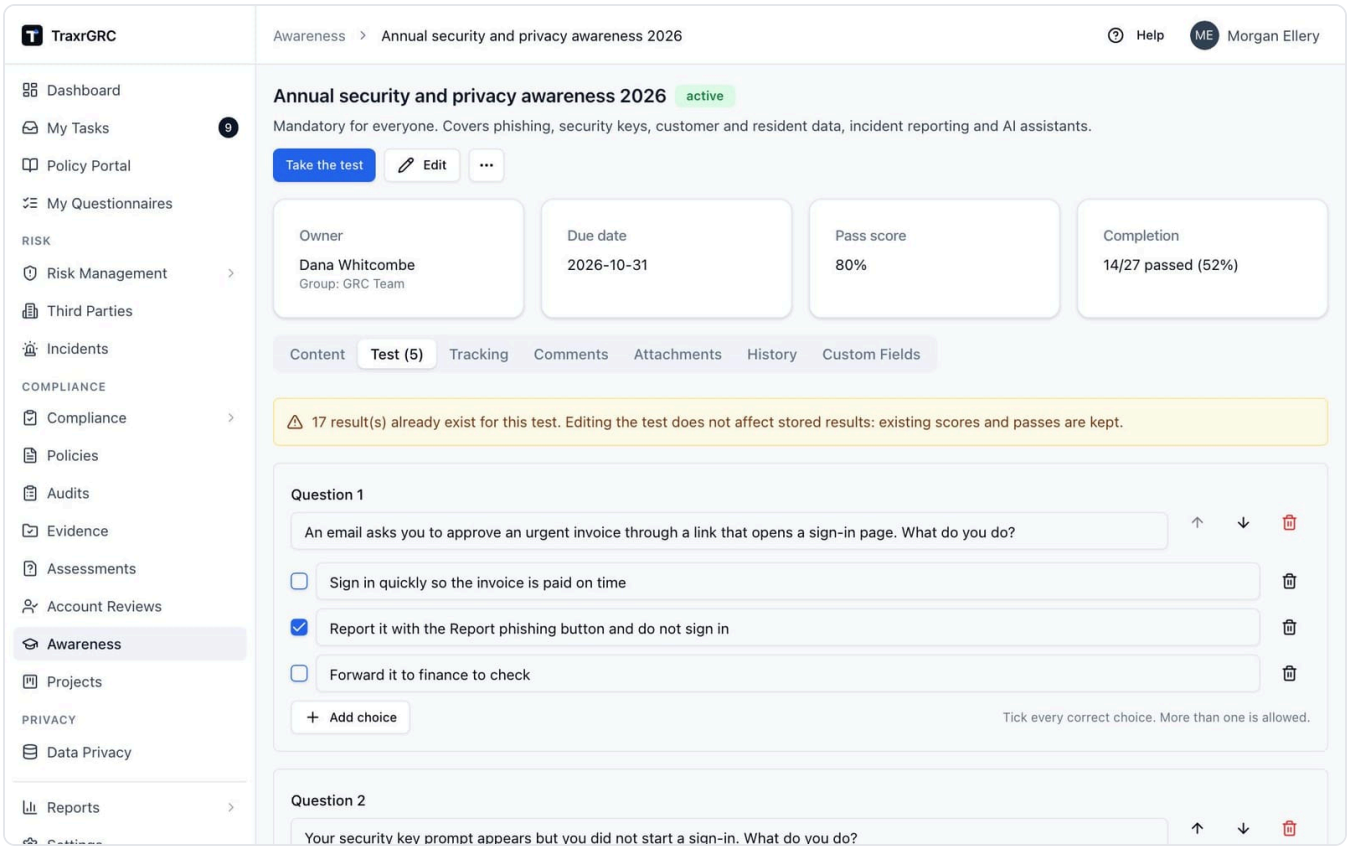


Figure 10.5. The test editor for an active program, with the amber warning about existing results and the correct choices ticked.

10.7 Tracking tab: who has completed the program

The Tracking tab lists **every active user**, one row each, with their latest state. It is shown to those who can update the program and to those with `auditLog:read` (the built-in GRC Manager and Auditor roles); reading the program is enough to take it, and someone with read access alone sees only their own result. Draft and archived programs are likewise listed only for them.

Column	What it shows
User	Name and email address.
Status	Passed (green), Failed (amber) or Not completed (gray).
Score	The percentage from the person's most recent attempt, or a dash if they have never submitted.
Completed	The date and time of the most recent submission.

Use the status filter or the search box to narrow the list, and **Export CSV** to download it, for example as evidence or for reporting.

If the program has no test yet, a notice tells you to add one in the Test tab before asking people to complete it, and the list shows everyone as Not completed.

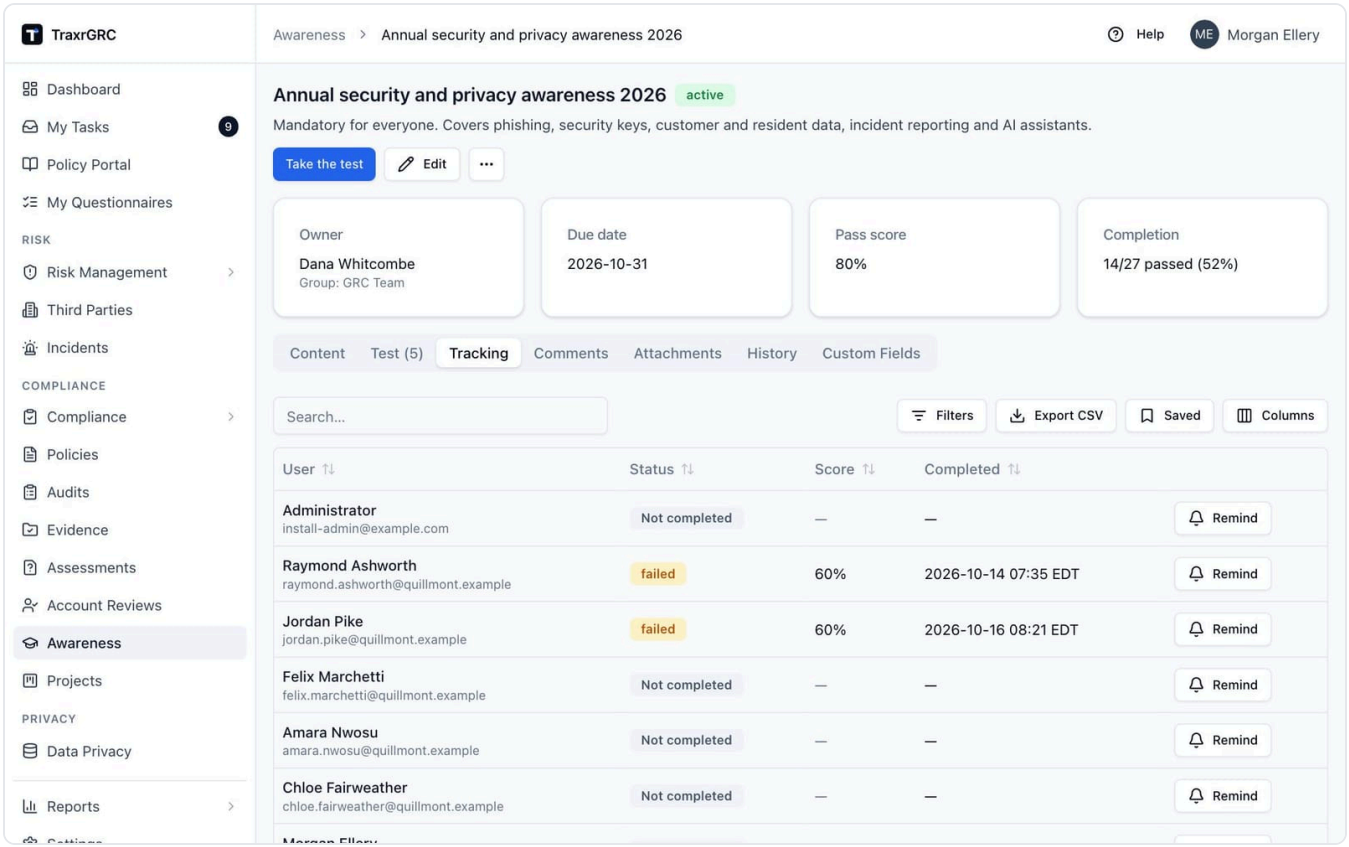


Figure 10.6. The Tracking tab, listing each user's status, score and completion date, with Remind buttons for people who have not passed.

10.7.1 Sending a reminder

For anyone who has not passed, you can send a reminder by hand (requires awareness:update):

1. Find the person in the Tracking list.
2. Click **Remind** on their row.

The person is emailed the "Awareness program due" notification at once, naming the program and its due date. The button then reads **Reminder sent** for the rest of your session. Each person gets at most one reminder per program a day, counting the daily reminder emails; a second **Remind** within 24 hours is refused with a message saying so.

10.7.2 Automatic reminders

Independently of the Remind button, the system's daily sweep emails every active user who has not passed an active program whose due date is within the notification lead window (default 7 days) or already past. The same email template is used, and each person is reminded at most once per sweep period. Lead windows and templates are configured in Settings → Notifications (see Administration & Settings).

10.8 Taking a test

This section is for everyone, not just program owners. You need awareness:read, the same permission that lets you see the Awareness module.

1. Open **Awareness** and click **Take the test** on a program, or open the program and use the button in its header. If the program is not active or has no test, the test cannot be opened.
2. Read the training material on the program's Content tab if you have not already.
3. On the test page, check the pass score and due date shown under the title, then work through each question.
4. Tick your answers. **More than one choice can be correct per question**: tick every choice you believe is right. A question only counts as correct when your selection exactly matches the intended set. There is no partial credit, and an unanswered question counts as incorrect.
5. Click **Submit answers**. The button stays disabled until at least one question is answered; the counter next to it shows how many questions you have answered out of the total.

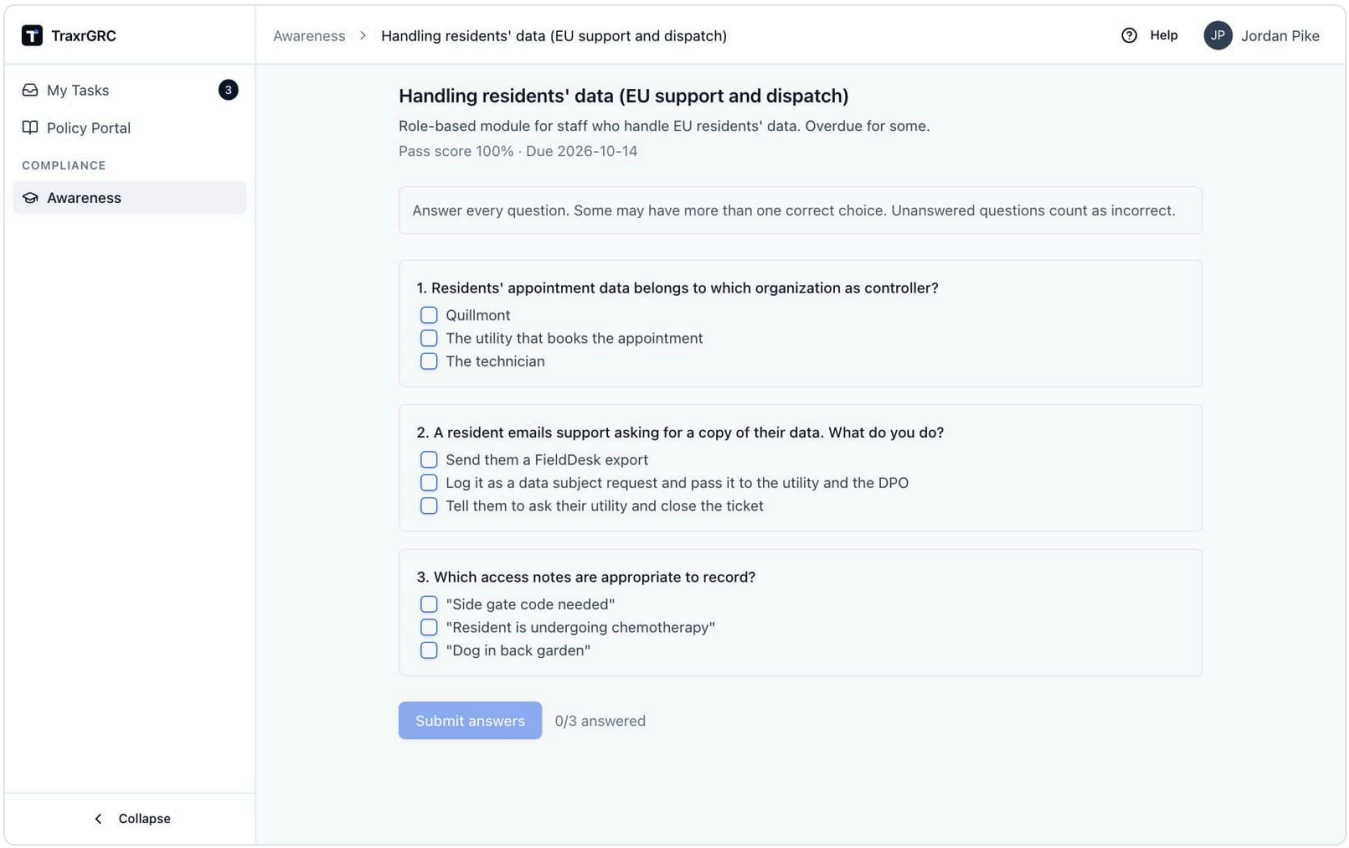


Figure 10.7. The test page as a staff member sees it, with checkbox choices for each question and the Submit answers button.

You immediately see a result screen:

- Your **score** as a percentage, and a **Passed** (green) or **Failed** (amber) badge next to the title, measured against the program's pass score.
- Per-question feedback: questions you got right are marked with a green tick and wrong ones with a red cross, with your own picks labeled "(your answer)". The right choices are never shown, on a pass or a fail, so the test cannot be passed by reading the key from an earlier attempt.
- **Back to program** returns you to the program page.

10.8.1 Retakes

You may retake a test as often as you like. When you open a test you have taken before, a banner above the questions shows your previous result: green if you passed, amber with **Not yet passed**, your latest score and the pass score if you did not. Two rules keep the record fair:

- A pass is never lost: if you passed earlier, a lower score on a retake does not remove your pass.
- The recorded score is always your **latest** attempt, so a retake that improves your score is reflected in Tracking.

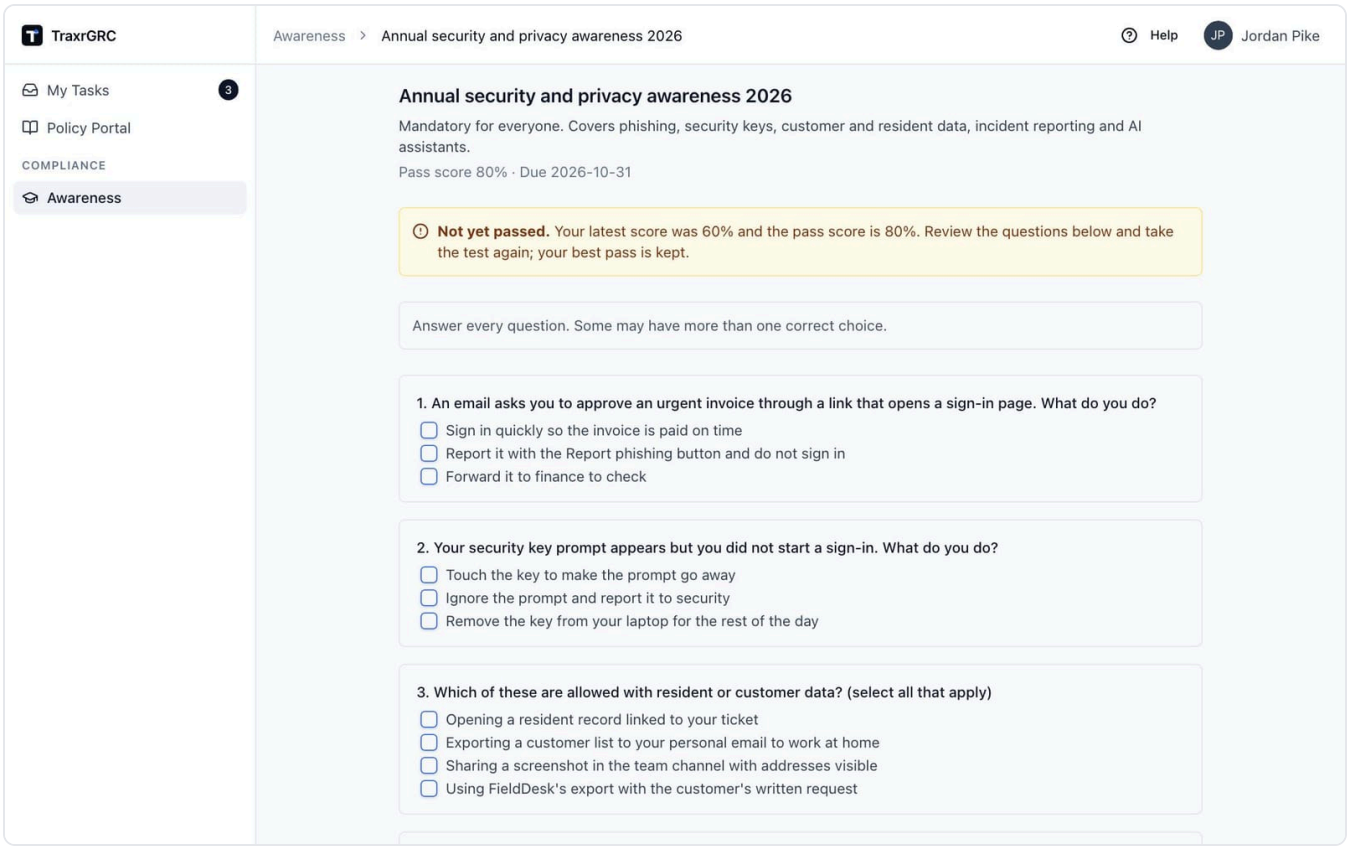


Figure 10.8. A retake of a test the staff member previously failed, with the amber Not yet passed banner showing their latest score and the pass score.

Once you have passed, the program list shows "Passed" for you, the detail page shows the "You passed" badge, and automatic due-date reminders stop for you.

10.9 Comments, attachments, history and custom fields

- **Comments.** Discuss the program in the Comments tab. Posting requires awareness:update on the program.
- **Attachments.** Upload reference files (policies, slide decks) in the Attachments tab. Uploading likewise requires awareness:update; anyone who can read the program can download them.
- **History.** The History tab lists the audit-log entries for this program: creation, updates, content changes, test saves, test submissions, reminders and deletions, each with the action, the person who made it, the date and the fields that changed (last 100 entries).
- **Custom Fields.** The Custom Fields tab, last in the row, shows any extra fields your organization has defined for awareness programs, or a note that none are defined. Editing values requires awareness:update on the program (field definitions are managed under the customFields module; see Administration & Settings).

10.10 Deleting a program

1. Open the program, open **More actions** (**...**) and click **Delete** (requires awareness:delete).
2. Read the confirmation. It states exactly what goes away: the program **with its test, all results, comments, attachments and labels**.
3. Click **Delete** to confirm, or **Cancel**.

Deletion cannot be undone and is recorded in the audit log. Consider archiving instead if you need to keep the completion evidence.

10.11 Awareness permissions

Action	Permission needed
See the list and open a program	awareness:read
Take or retake a test	awareness:read (program must be active, with a test)
See the answer key on the Test tab	awareness:update on the program
Create a program	awareness:create
Edit settings, content or the test; send reminders	awareness:update on the record; at "own" level you must own it (or its owner group)
Post comments, upload attachments	awareness:update on the record
Delete a program	awareness:delete on the record

Roles are assigned to groups and inherited through group membership; administrators bypass all of the above (see Managing Users, Groups & Roles).

CHAPTER 11

Data Privacy

Data Privacy is your register of records of processing: one entry for each activity in which your organization handles personal data (for example "CRM customer data processing"), documenting the legal basis, the data involved, where it is stored, and who is responsible for keeping the entry up to date. Each record carries a review schedule so the register stays current over time. Around the register the module provides data protection impact assessments (DPIAs), the Article 30 register export, a data map, and data mapping questionnaires that collect the facts from the business people who know each activity.

The privacy team and the Data Protection Officer (DPO) use this module day to day. Data Liaisons, the business contacts for each activity, see only their own questionnaires (see Data mapping questionnaires below). Auditors and colleagues with read access use the register and its exports as evidence. Who can do what is controlled by the dataPrivacy module permissions, summarized at the end.

11.1 The records list

Open **Data Privacy** in the navigation (under Privacy). The page has three tabs:

Tab	Contents
Records of processing	The register itself, described below.
DPIAs (*N*)	Every DPIA on records you can read (see DPIAs below).
Questionnaires	The questionnaires sent to Data Liaisons and the questionnaires available to send (see Data mapping questionnaires below).

On **Records of processing**, each row is one record of processing:

Column	What it shows
Name	Click it to open the record.
Status	Draft, Active or Archived.
Legal basis	The GDPR lawful basis as a **blue badge**: Consent, Contract, Legal obligation, Legitimate interest, Vital interests or Public task (or a dash if unset).
Data categories	Up to two category chips, with "+N more" when there are more.
Processor	The first processor on the record, if any. Click to open its profile (see Third Parties).
Owner	The named owner, or the owner group.
Next review	When the next review is due. Shown in **red with "(overdue)"** once the date has passed, unless the record is archived.

The table can be searched and filtered by status and legal basis, and **Export CSV** downloads the current view. The buttons above the table are:

Button	What it does
Data map	Opens the data map (see The data map below).
Article 30 register (CSV)	Downloads the full register as a spreadsheet (see The Article 30 register below).
Article 30 register (PDF)	Downloads the full register as a formatted document.
New record	Creates a record (requires dataPrivacy:create).

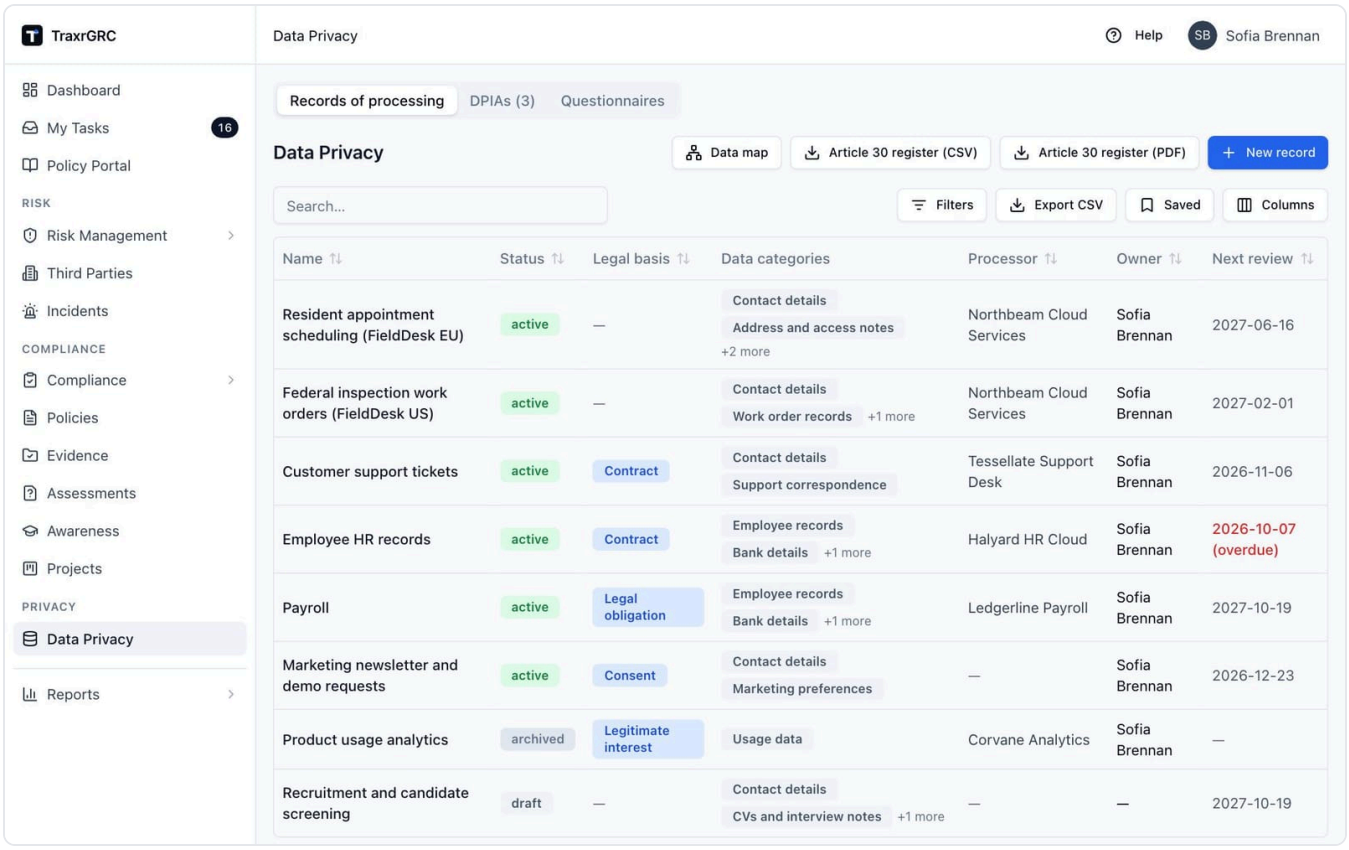


Figure 11.1. The Data Privacy records list, with legal basis badges and an overdue next review shown in red.

11.2 Creating a record

1. Click **New record** (requires dataPrivacy:create).
2. Complete the **New record of processing** dialog:

Field	What it means
Name	A short name for the processing activity, e.g. "CRM customer data processing". Required.
Description	What the activity covers.
Owner	An individual user responsible for keeping the entry accurate (optional).
Owner group	A responsible group (optional). Used for "own"-level permission scoping.
Review frequency (months)	How often the record must be reviewed, 1 to 120. Defaults to 12.
Processor (third party)	The third party that processes the data on your behalf, if any (only non-terminated third parties are offered; see Third Parties).
Schedule the first review *N* months from now	Tickbox, on by default. When ticked, the record is created with its first review already scheduled and the Next review date set accordingly. Untick to create the record with no scheduled review.

3. Click **Create record**. The record is created as a **Draft**.
4. Open the record and click **Edit** to fill in the full documentation (below).

11.3 The record detail page

Click a record name to open it. The header shows the name, the status badge, the legal basis badge and the description. The header buttons depend on the record's status and your permissions:

Button	When it appears
Activate, **Archive**, **Reactivate**	You have dataPrivacy:update; which ones depend on the status (see Record status below).
Edit	You have dataPrivacy:update on the record.
Delete	You have dataPrivacy:delete on the record.
Send questionnaire	You have dataPrivacy:update on the record and it is not archived (see Data mapping questionnaires below).

Four summary cards sit under the header:

Card	What it shows
Owner	The owner, the owner group and the record's **Data Liaison** (or "none").
Processors	Each processor, linked to its third-party profile, with sub-processors marked as such.
Next review	The next review date, marked overdue in red once it has passed, and the review frequency.
Retention	The retention period.

Below the cards come the Article 30 completeness notice and the **Linked assets, risks and controls** card, both described below.

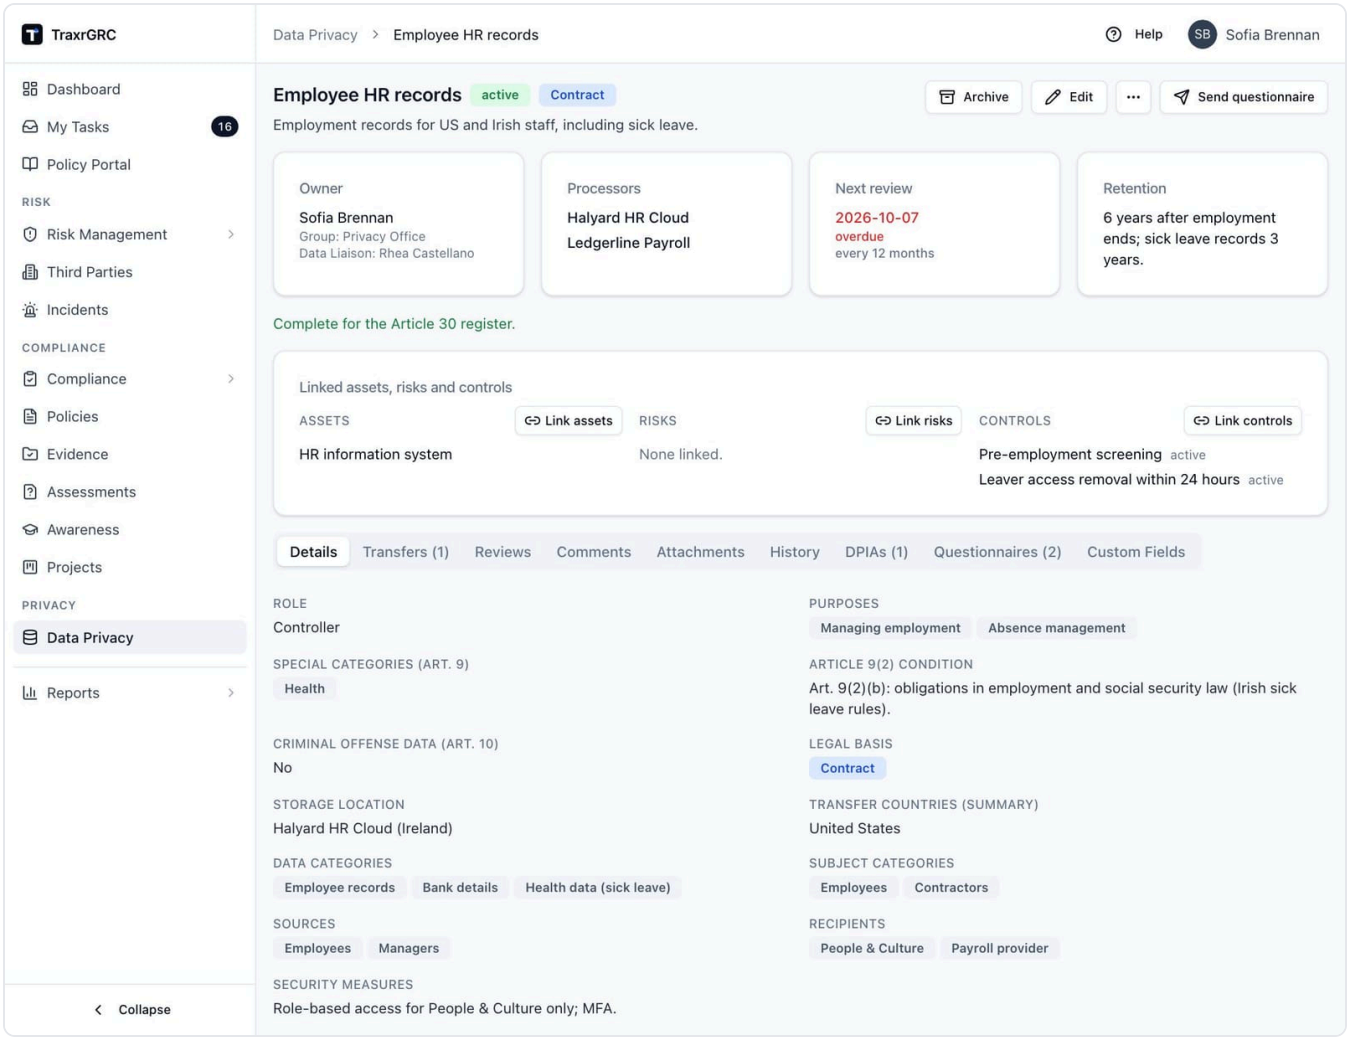


Figure 11.2. A record of processing with an overdue review, showing the header buttons, the summary cards, the Article 30 notice and the linked assets, risks and controls.

The page has one row of tabs:

Tab	Contents
Details, Transfers (*N*), Reviews, Comments, Attachments, History	Everything about the record itself, described in the sections below.
DPIAs (*N*)	The record's data protection impact assessments (see DPIAs below).
Questionnaires (*N*)	Questionnaires sent about this record, with the questionnaire, Data Liaison, status, sent date, answer-by date and the date answered.
Custom Fields	Organization-specific fields (see Comments, attachments, history and custom fields below).

11.4 Documenting the record

The creation dialog captures only the essentials. To document the record fully:

1. Open the record and click **Edit** (requires dataPrivacy:update).
2. Complete the **Edit record of processing** dialog. It contains everything from the creation dialog plus the fields below.
3. Click **Save changes**.

Field	What it means
Data Liaison	The business contact who answers questionnaires about this activity and receives its annual checks (see Data mapping questionnaires below). Only people with the Data Liaison role are offered.
Role	Whether your organization acts as **Controller**, **Joint controller** or **Processor** for this activity. A processor's record needs less for Article 30 (see below).
Purposes of processing	Why the data is processed, e.g. "Payroll administration". A tag input like the category fields.
Joint controllers	Shown for the joint-controller role: the name and contact details of each joint controller.
Processors and sub-processors	The third parties that process the data for you. Pick one from **Add a third party** and set it to **Processor** or **Sub-processor**; remove it with the bin icon. The first one listed is the record's main processor.
Special categories of data (Art. 9)	Tick each special category involved: health, genetic, biometric (for identification), racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, sex life or sexual orientation.
Article 9(2) condition	Shown once a special category is ticked: the condition that permits processing it, e.g. "9(2)(i) public interest in public health".
Criminal convictions and offenses data (Art. 10)	Tick when the activity involves such data.
Legal basis	The lawful ground for the processing: **Consent, Contract, Legal obligation, Legitimate interest, Vital interests** or **Public task** (or "None"). Shown as a blue badge on the record.
Retention	How long the personal data is kept, e.g. "5 years after account closure".
Storage location	Where the data is stored, e.g. "EU data center".
Transfer countries (summary)	A free-text summary kept from earlier versions. Record each transfer and its safeguard on the Transfers tab instead.
Security measures	Free text describing the safeguards in place, e.g. encryption and access control.

Field	What it means
Data categories	The kinds of data involved, e.g. "Contact details".
Subject categories	Whose data it is, e.g. "Customers", "Employees".
Sources	Where the data comes from, e.g. "Web sign-up forms".
Recipients	Who receives the data, e.g. "Sales team".

The purposes field and the four category fields (data categories, subject categories, sources, recipients) are tag inputs: type a value and press **Enter** to add it as a chip (you can also separate values with a comma). Duplicate values are ignored. Remove a value by clicking the **×** on its chip.

Everything you document is displayed on the record's **Details** tab, ready for an auditor or a colleague to read.

11.5 Article 30 completeness

Below the summary cards the record either says it is "Complete for the Article 30 register" or shows an amber **Missing for the Article 30 register** notice listing what it still lacks.

Record's role	What it needs besides its role
Controller, joint controller, or no role set	Purposes, categories of data subjects, categories of personal data, recipients, retention period, security measures and legal basis; an Article 9(2) condition when special categories are ticked; the joint controllers for the joint-controller role. A record with no role set is also asked for its role.
Processor (Article 30(2))	Only the categories of personal data and the security measures.

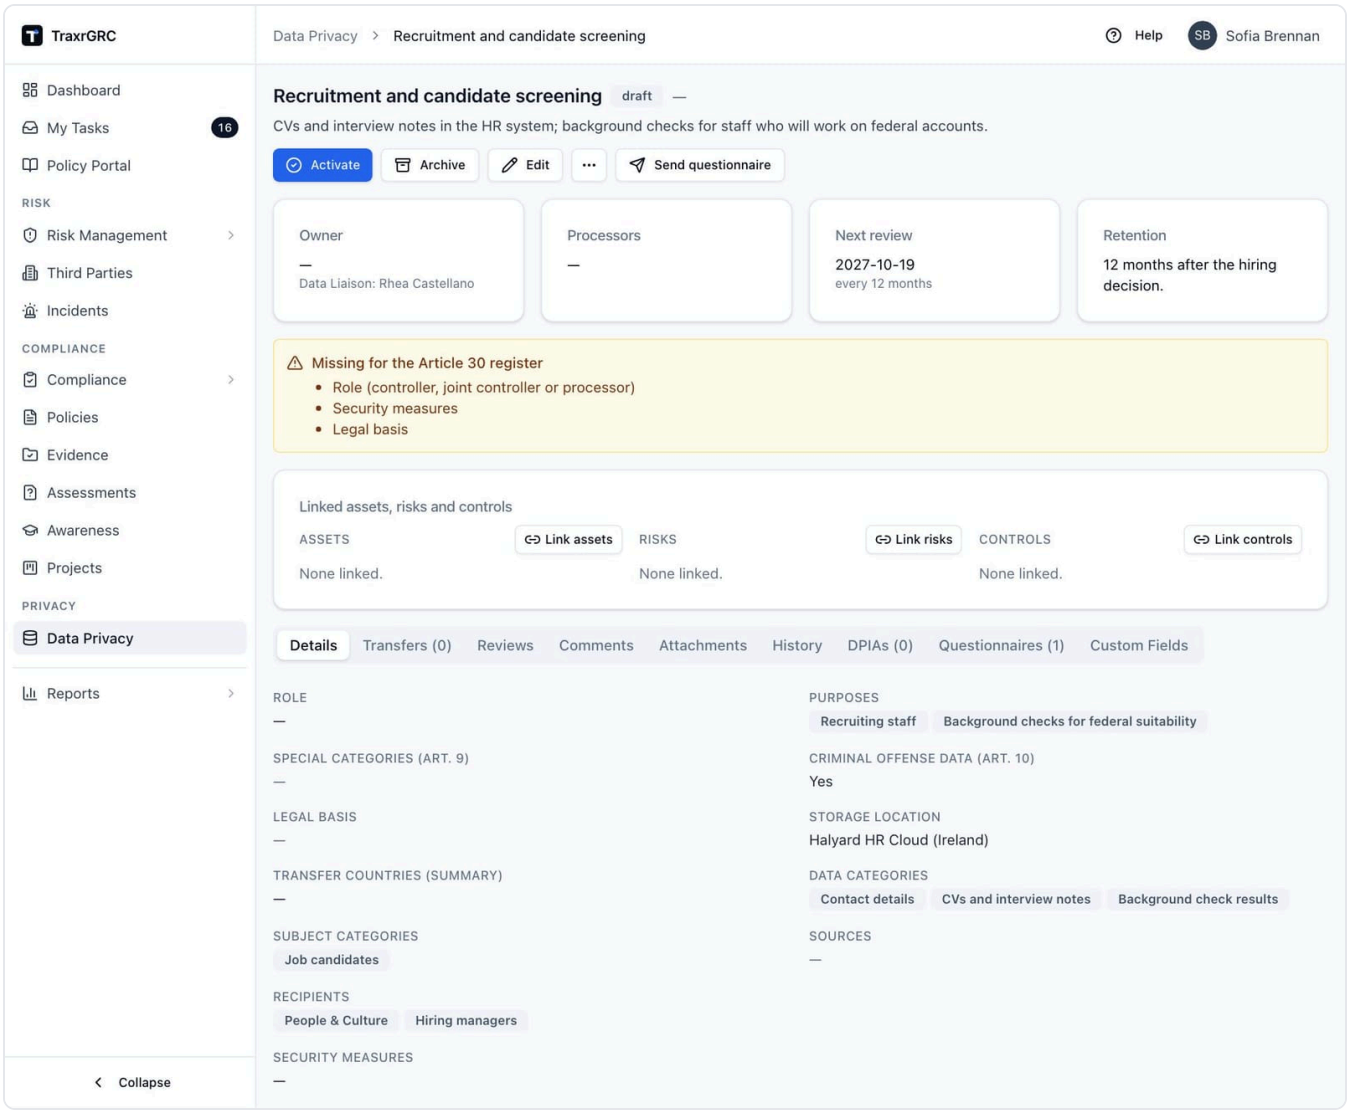


Figure 11.3. A draft record created from a Data Liaison's answers, with the amber notice listing what is missing for the Article 30 register.

11.6 Transfers

The **Transfers** tab lists each transfer of the record's data to a country outside the EEA, under the heading **Transfers to third countries**. To add one (requires dataPrivacy:update):

1. Click **Add transfer**.
2. Complete the dialog:

Field	What it means
Country	The destination country. Required.
Recipient (optional)	Who receives the data there.
Safeguard	The transfer mechanism relied on: Adequacy decision, Standard contractual clauses, Binding corporate rules, EU-US Data Privacy Framework, Derogation (Art. 49) or Other.
Transfer impact assessment	The vendor's transfer impact assessment the transfer rests on (see Third Parties: Data protection). Shown only when such assessments exist.
Notes (optional)	Anything else worth recording.

3. Click **Add transfer**.

A transfer under standard contractual clauses or binding corporate rules with no assessment cited shows an amber **No transfer impact assessment** badge. Remove a transfer with the bin icon.

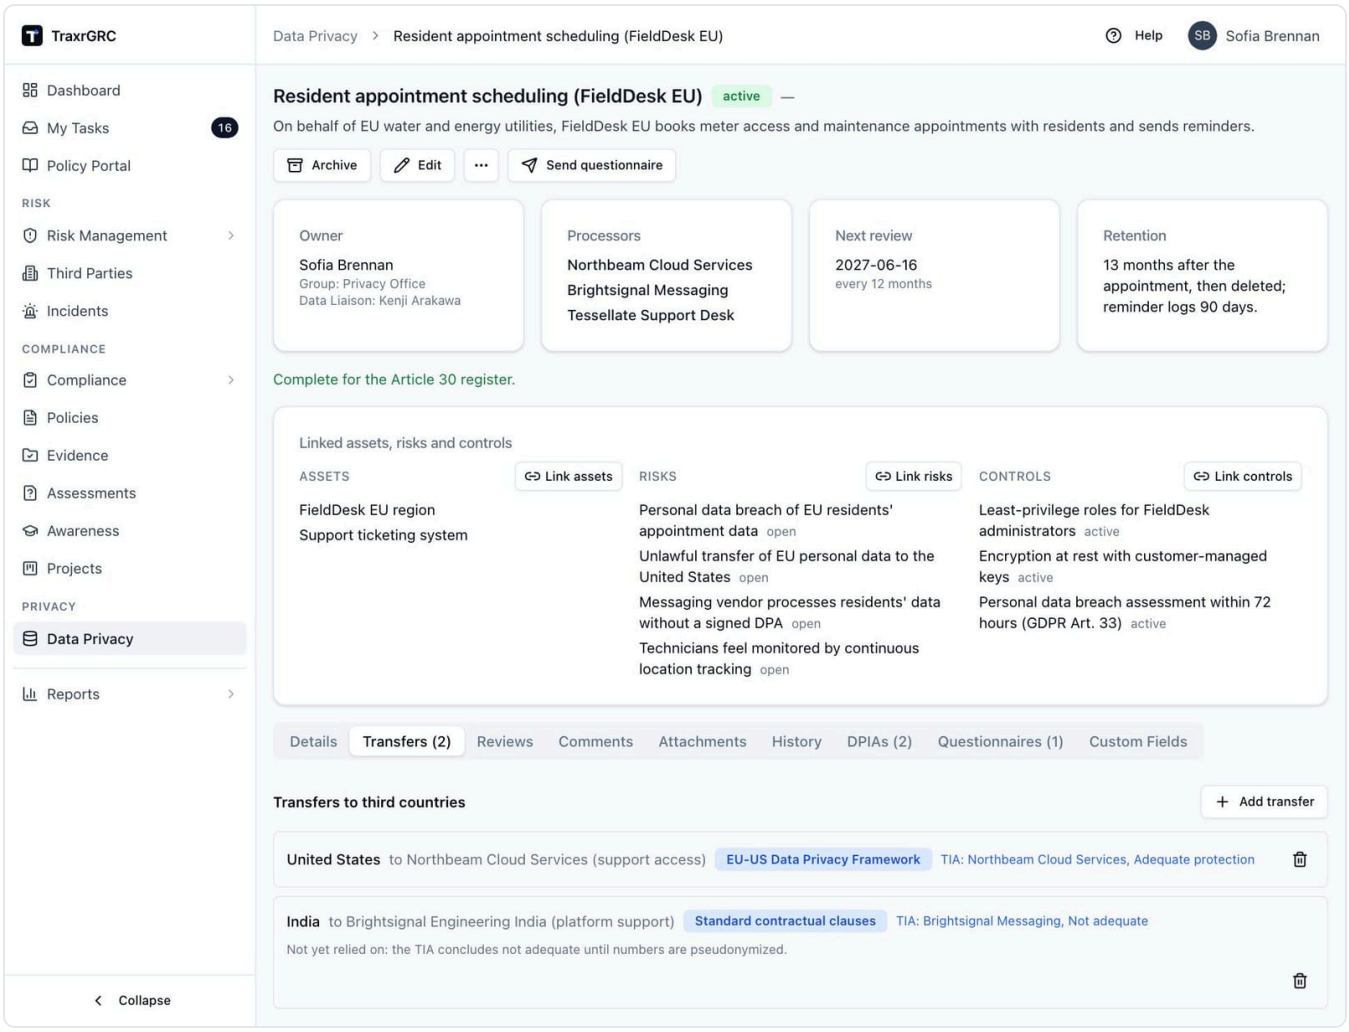


Figure 11.4. The Transfers tab of a record, listing two transfers with their safeguards and transfer impact assessments.

11.7 Linked assets, risks and controls

The **Linked assets, risks and controls** card ties the record to the rest of the register. To link an entry:

1. Click **Link assets**, **Link risks** or **Link controls**.
2. Search, and pick an entry to link it. Pick it again to unlink it.

Linking needs dataPrivacy:update on the record and read on the linked module; entries in modules you cannot read are counted, not shown. Deleting a linked asset, risk or control removes the link and records that on the record's history, naming what was deleted; a deleted risk is recorded the same way on any DPIA it was linked to.

11.8 The Article 30 register

Article 30 register (CSV) and **Article 30 register (PDF)** on the Data Privacy page download the records of processing activities.

Both contain a header naming the controller, its contact details, representative and DPO (from Settings → Privacy) and the date, then one entry per active record that you can read, with (drafts and archived records are left out, since a draft is not yet a live processing activity):

- role and joint controllers;
- purposes and legal basis (printed with its Article 6(1) letter);
- categories of data subjects and of personal data, special categories and the Art. 9(2) condition, and criminal offense data;
- recipients and processors;
- each transfer with its safeguard (a record with no transfers listed falls back to its transfer summary);
- retention, security measures and owner.

The PDF lays out each activity as a block of labeled fields in Article 30(1) order, with a page footer naming the controller and date, and prints any empty required field as **Not recorded** so gaps are visible before the register goes to a supervisory authority.

11.9 DPIAs

A data protection impact assessment (Art. 35) belongs to one record of processing. The **DPIAs** tab on the Data Privacy page lists every DPIA on records you can read; the **DPIAs** tab on a record lists that record's DPIAs. Both tables show:

Column	What it shows
DPIA	The title. Click it to open the DPIA.
Record of processing	The record it belongs to (on the Data Privacy page only).
Status	See the table below.
Screening	Not screened, Likely required or Likely not required.
Decided	When it was approved or rejected.
Next review	The DPIA's next review date.

A DPIA moves through these statuses:

Status	Meaning
Draft	Created; screening not yet saved.
Screened	Screening saved. Either kept as the evidence that no DPIA was needed, or ready for the full assessment.
In progress	The full assessment is being written.
Awaiting approval	Submitted for the DPO's advice, where required, and a decision. Frozen apart from prior consultation.
Approved	Approved by someone with the Approve DPIAs right. Locked.
Rejected	Rejected by someone with the Approve DPIAs right. Locked until reopened.

11.9.1 Starting and screening a DPIA

1. Open the record's **DPIAs** tab and click **New DPIA** (requires dataPrivacy:update on the record). DPIAs are always started from their record; the Data Privacy page's DPIAs tab has no **New DPIA** button.
2. Give it a **Title** and click **Create DPIA**. You arrive on the DPIA page.
3. Under **Screening (WP248 criteria)**, tick each of the nine criteria from the WP29 guidelines endorsed by the EDPB that applies: evaluation or scoring; automated decisions with legal or similarly significant effect; systematic monitoring; sensitive or highly personal data; large scale; matching or combining datasets; vulnerable data subjects; innovative use or new technology; and processing that prevents people exercising a right or using a service or contract. Two or more met makes a DPIA likely required; the chip next to the heading says so as you tick.
4. Click **Save screening**. If fewer than two are met, the screened DPIA stays as your evidence that none was needed, or you can still carry on.

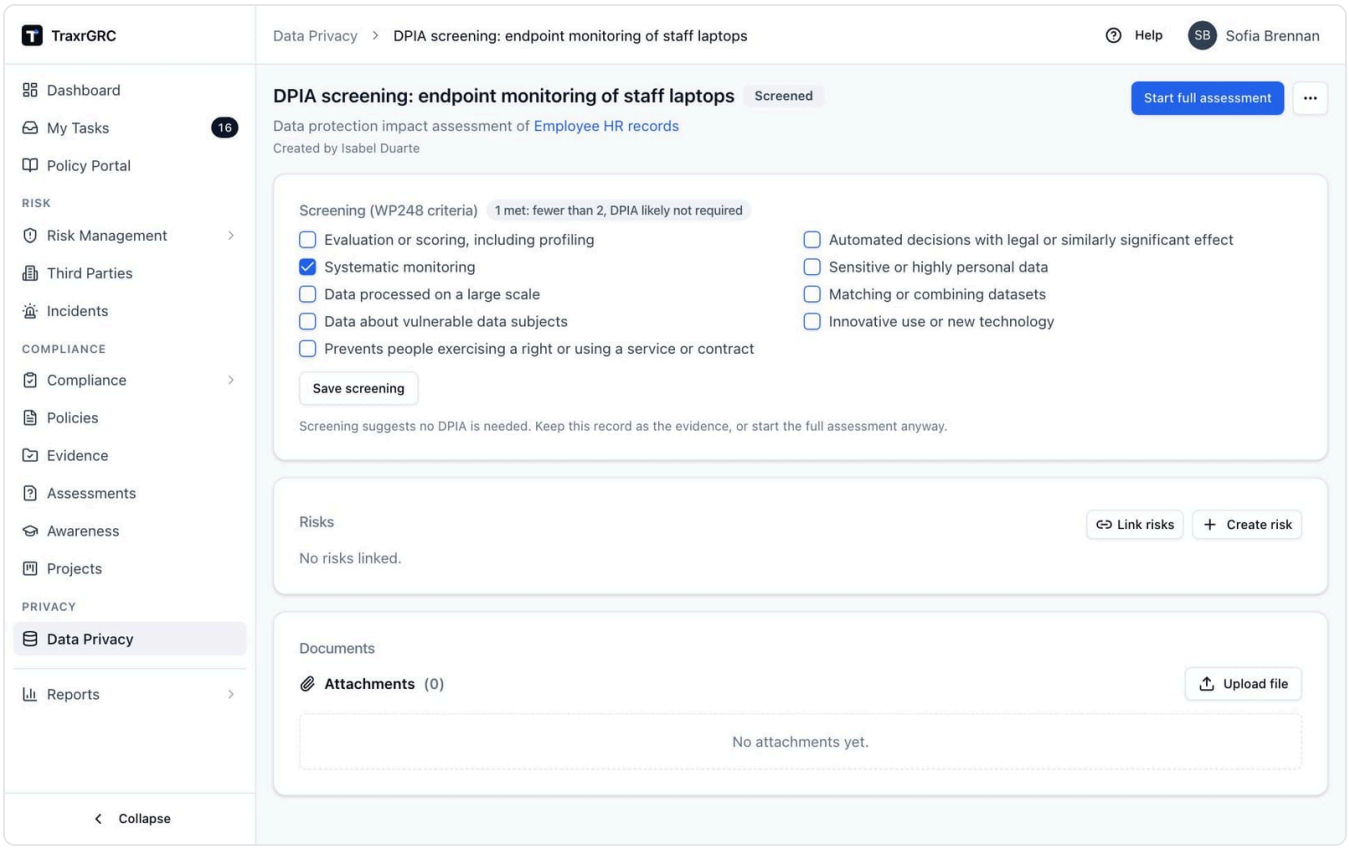


Figure 11.5. A screened DPIA with one criterion met, so the chip reports that a DPIA is likely not required.

11.9.2 The full assessment

1. Click **Start full assessment**. The **Assessment (Art. 35(7))** section opens with three texts: the systematic description of the processing and its purposes, necessity and proportionality, and the risks to individuals with the measures that address them.
2. Tick **A high risk remains after the measures** when it does, and record **Prior consultation** with the supervisory authority (Art. 36): **Not required**, **Required, not yet made**, or **Authority consulted**, with its date (**Consulted on**) and **Reference**.
3. Set an **Owner** and a **Next review** date if you like, then click **Save assessment**.
4. Under **Risks**, **Link risks** connects existing risks. **Create risk** makes a new one from the DPIA: give it a **Title** and **Classification**; it is linked to the DPIA and to the record, starts at the lowest likelihood and impact for you to rate on the risk page, and needs risks:create.
5. Click **Submit for approval**. This is refused until the description, necessity and measures are written. From then until it is decided, the DPIA is frozen: nobody can change its text, screening, residual-risk answer, linked risks, owner or review date. The one exception is prior consultation, which an approver or the DPO (with dataPrivacy:update on the record) may still set with **Save prior consultation**, since approval can depend on it. Changing it clears any DPO advice already given, since the advice covered the earlier text.

The DPIA page names who created it and who submitted it for approval.

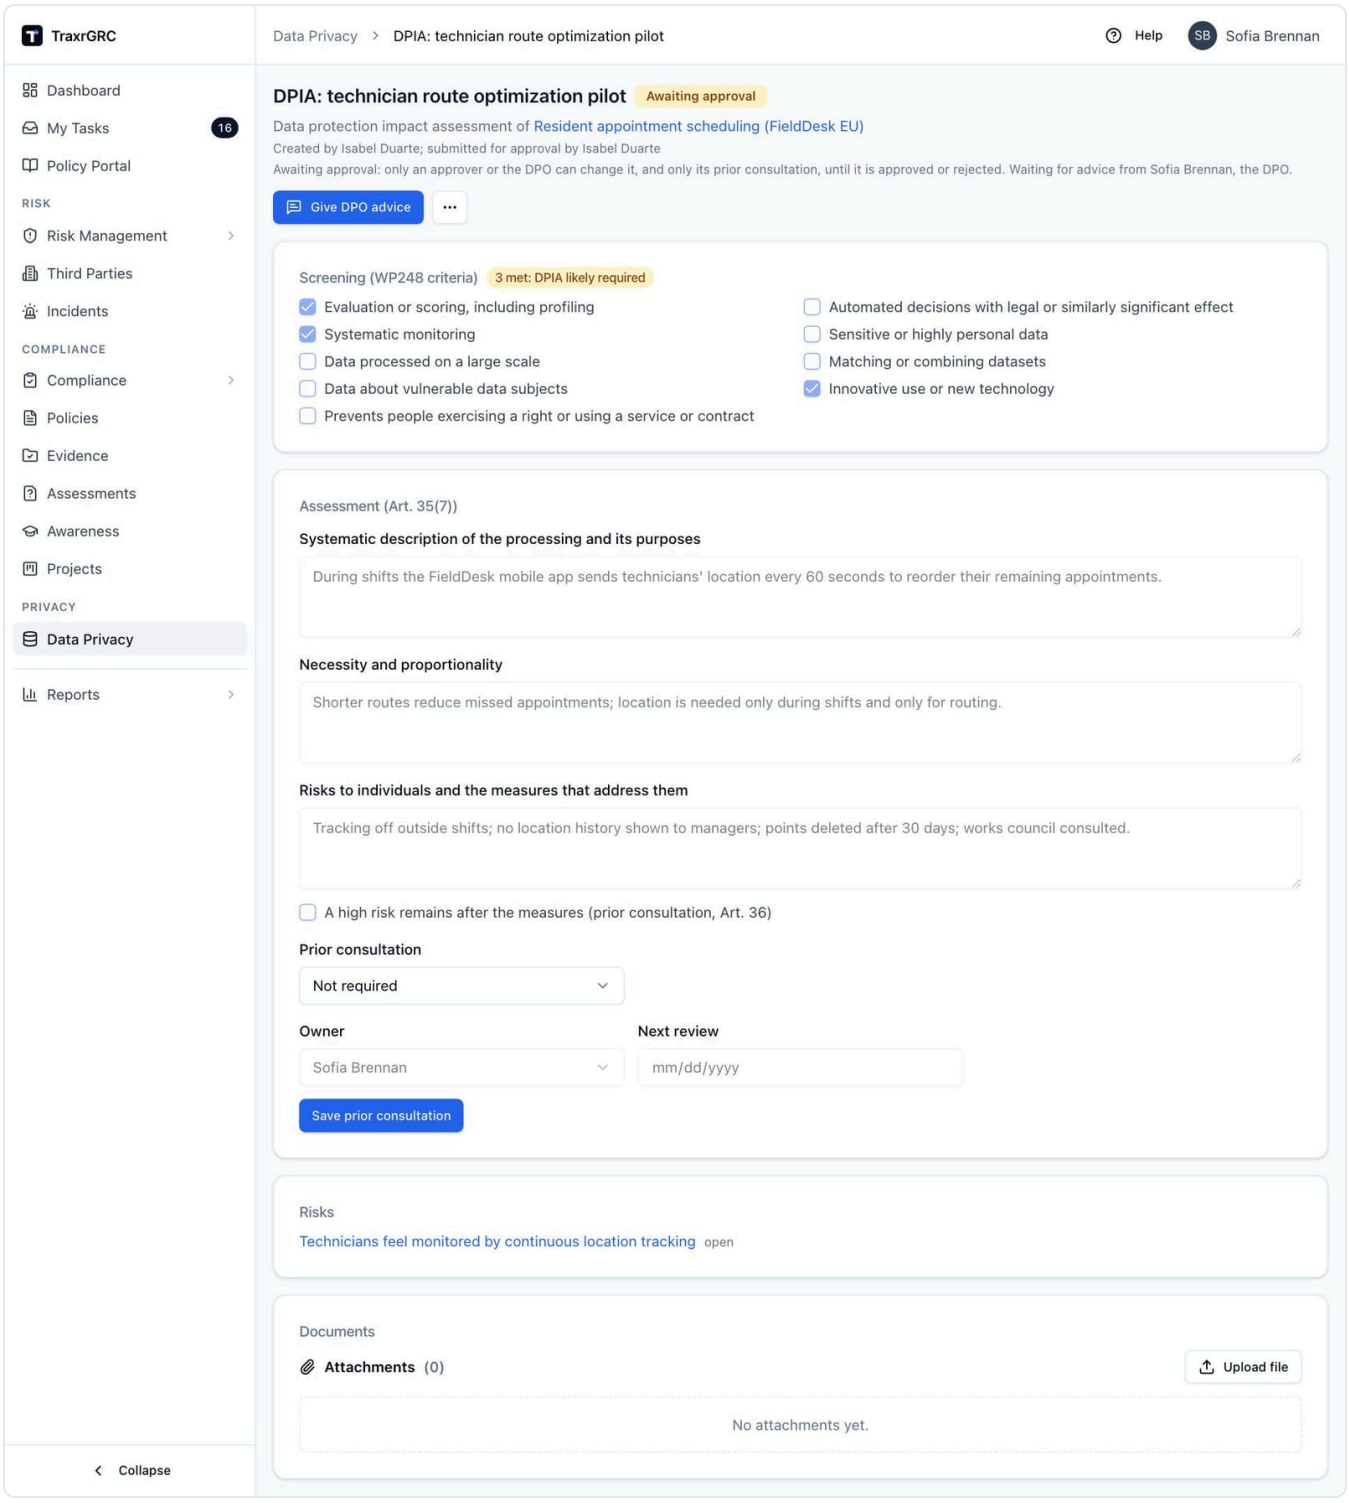


Figure 11.6. A DPIA awaiting approval, showing the frozen-state notice, the completed assessment and the linked risk.

11.9.3 DPO advice

Article 35(2) requires the organization to seek the DPO's advice when carrying out a DPIA. The DPO gives advice; deciding is a separate step for someone with the Approve DPIAs right, who may or may

not be the DPO.

An administrator names the **Data Protection Officer** in Settings → Privacy (see the Administration chapter). Only that person can give DPO advice, and their DPIAs awaiting advice appear in their My Tasks as **Advise on …**. Whether advice is required before a decision is also set there; by default it is required whenever a DPO is named.

1. As the DPO, open a DPIA awaiting approval and click **Give DPO advice**.
2. Choose a **Recommendation**: **Proceed**, **Proceed with changes** or **Do not proceed**.
3. Write your **Advice** and click **Save advice**.

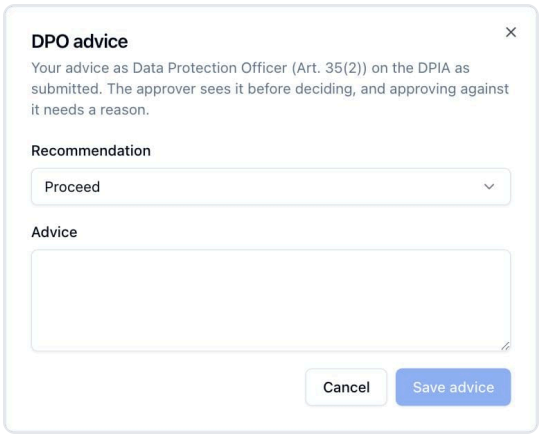


Figure 11.7. The DPO advice dialog, with the recommendation, the advice text and Save advice.

The advice appears on the DPIA under **DPO advice**, with your name and the time. Until the DPIA is decided you can change it with **Revise DPO advice**. Like a decision, it covers exactly the DPIA you have open and is refused if the DPIA changed meanwhile. Resubmitting or reopening the DPIA clears it, so the DPO advises on each submission.

While advice is required and missing, approvers see "Waiting for advice from" the DPO on the DPIA, no **Approve or reject** button, and no task in My Tasks.

11.9.4 Approving or rejecting

1. Open the DPIA and click **Approve or reject** (requires the Approve DPIAs right, dataPrivacy:approve, on the record). The dialog repeats the DPO's recommendation, if any.
2. Write a **Decision note**.
3. If the DPO recommended changes or not proceeding, approving also needs a **Reason for approving against the DPO's advice**.
4. Click **Approve** or **Reject**.

The decision is recorded in your name: the DPIA shows **Decision** with your note and "Approved by" or "Rejected by" you, with the time. An approval against the DPO's advice is marked **Approved against DPO advice** with your reason, recorded in the audit log, and emailed to the DPO. Rejecting never needs a reason beyond the note.

The rules for deciding are:

- When DPO advice is required, the decision waits for it.
- The author may decide their own DPIA unless an administrator has turned on **Require a second person to approve DPIAs** (Settings → Privacy). Then whoever created the DPIA or submitted it sees why the button is missing, and another approver decides.
- The decision covers exactly the DPIA you have open. It is refused if the DPIA changed since the page was loaded or since it was submitted (a linked risk deleted meanwhile, for example; reject it so it can be resubmitted), if the DPO's advice changed meanwhile, or if someone else has already decided it.
- Approval is refused while a high residual risk has prior consultation set to not required, and when consultation is marked as made without its date.
- Approving sets the next review a year out unless a date is already set.

11.9.5 After the decision

An approved or rejected DPIA is locked; **Reopen** returns it to In progress and clears the decision and the DPO's advice. **Documents** holds the full assessment and any supporting files. **Delete**, in the **More actions** menu (**...**), asks for confirmation ("Delete DPIA?") and then removes the DPIA (requires dataPrivacy:delete on the record); deleting the record deletes its DPIAs.

11.10 The Reviews tab

Records are kept honest through scheduled reviews. Open the record's **Reviews** tab.

Upcoming reviews lists pending reviews by scheduled date; a date in the past is shown in red. To complete a pending review:

1. Click **Complete review** (requires dataPrivacy:update).
2. In the dialog, choose the **Result**:

Result	Meaning
Approved	The record is accurate as documented.
Changes required	The record needs corrections (describe them in Notes).

3. Add optional **Notes**.
4. Click **Complete review**.

Completing a review records the outcome, your name and the date in the review history, and **automatically schedules the next review** at the record's review frequency. The Next review date moves forward to match, leaving one pending review ready.

To schedule an additional, ad-hoc review without completing one:

1. Click **Add review**.
2. Pick a **Scheduled date**.
3. Click **Schedule review**.

The Next review date always reflects the earliest pending review.

History lists completed reviews newest first: the result badge (**Approved** in gray, **Changes required** in amber), the completion date, the reviewer's name and any notes.

Changing the **Review frequency** on the record re-anchors pending reviews: they are rescheduled from the most recent completed review (or the record's creation, if none) at the new frequency.

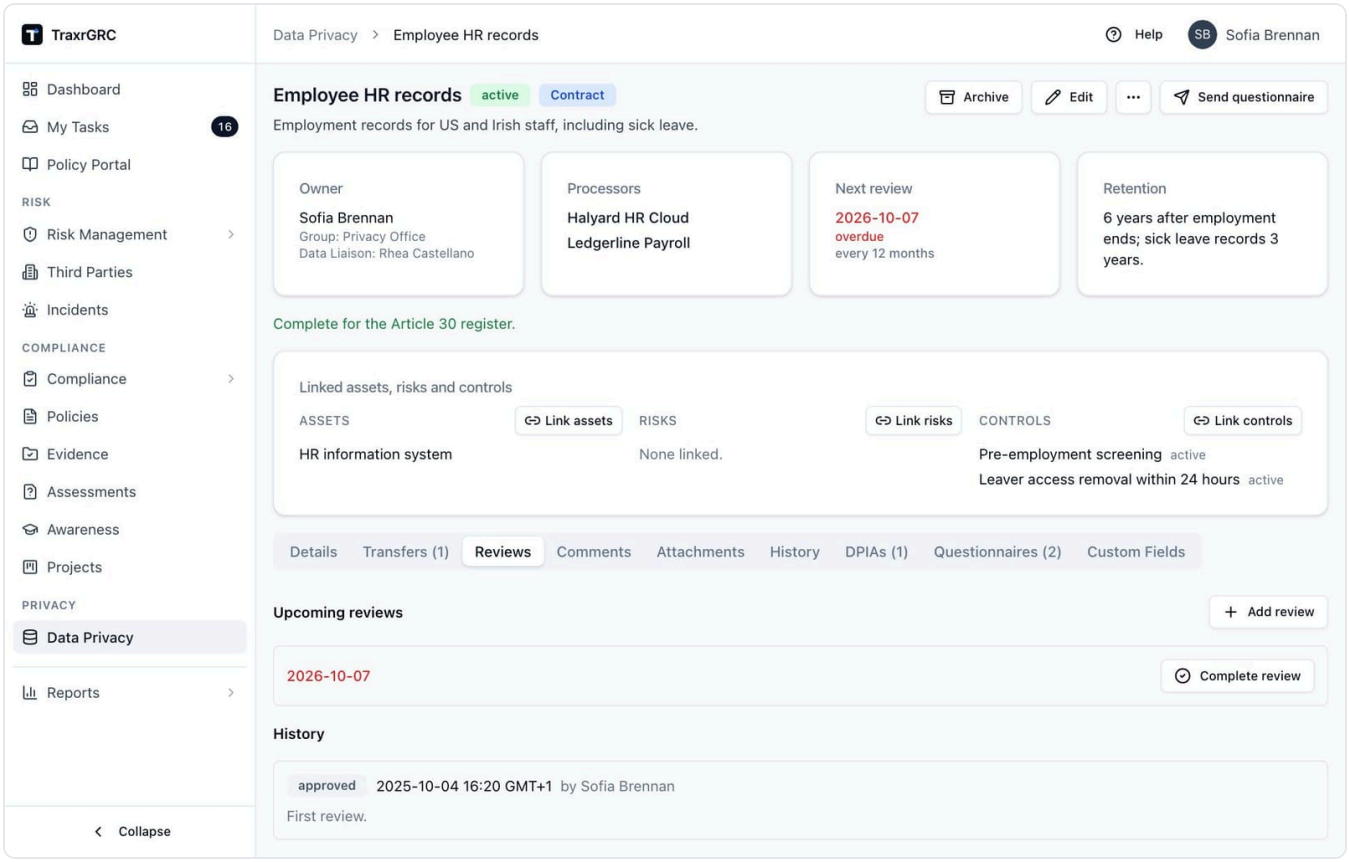


Figure 11.8. The Reviews tab of a record with an overdue pending review in red and the completed review history below.

11.11 Review notifications

The daily notification sweep watches every record that is not archived:

When	What the owner receives
A review comes due within the lead window (default 7 days)	A "Data privacy review due soon" email naming the record and its due date.
The review date has passed	An "OVERDUE data privacy review" email.

Owners who are deactivated receive nothing, and archived records are excluded entirely. A record with a Data Liaison gets an annual check instead while that check is open (see Annual checks below). Templates and lead windows are configured in Settings → Notifications (see Administration & Settings). Reminders also appear in the notification center (see Notifications).

11.12 Record status: draft, active, archived

Status	Badge	Meaning
Draft	Gray	Being documented. Not yet part of the live register.
Active	Green	The live register entry. Review scheduling and due-date notifications apply.
Archived	Gray	Retired. Excluded from review notifications; its next-review date no longer flags red.

The detail page offers dedicated buttons (requires dataPrivacy:update):

Button	Move
Activate	Draft to Active.
Archive	Draft or Active to Archived.
Reactivate	Archived to Draft; activate it again afterward to return it to the live register.

11.13 Comments, attachments, history and custom fields

- **Comments.** Discuss the record in the Comments tab. Posting requires dataPrivacy:update on the record.
- **Attachments.** Upload supporting evidence (DPIAs, processor agreements, data maps) in the Attachments tab. Uploading requires dataPrivacy:update; anyone who can read the record can download.
- **History.** The History tab lists the audit-log entries for the record and its reviews: creation, edits, status changes, scheduled and completed reviews (last 100 entries), each with the action, actor,

date and changed fields. A change applied from a Data Liaison's answers links to the questionnaire it came from.
- **Custom Fields.** The Custom Fields tab, last in the row, shows any extra fields your organization has defined for records of processing, or a note that none are defined. Editing values requires dataPrivacy:update on the record (see Administration & Settings).

11.14 Deleting a record

1. Open the record, open **More actions** (**...**) and click **Delete** (requires dataPrivacy:delete).
2. Read the confirmation, which warns that the record is deleted **together with its reviews**, along with its comments, attachments and labels. Its transfers, links, DPIAs and their documents go too.
3. Click **Delete** to confirm, or **Cancel**.

Deletion cannot be undone and is recorded in the audit log. If the processing activity has simply ended, archiving is usually the safer choice.

11.15 Data mapping questionnaires

Data mapping questionnaires ask the business people who know a processing activity, called **Data Liaisons**, to describe it, so the privacy team does not have to chase the facts by hand. Each answer can fill a field of the record of processing, and nothing changes on the record until the privacy team has reviewed the answers.

11.15.1 Who can be a Data Liaison

Anyone with the **Data Liaison** role, given through a group (Settings → Groups), plus administrators. A Data Liaison needs an account: with SAML single sign-on, the account is created the first time they sign in through SSO, after which an administrator adds them to a group carrying the role; without SSO, an administrator creates the account. The role grants nothing else, so a Data Liaison sees only **My Questionnaires** and lands there after signing in; give them other roles as well where they need more.

11.15.2 The Questionnaires tab

The **Questionnaires** tab on the Data Privacy page has two sections.

Sent questionnaires lists every questionnaire sent, with the number overdue next to the heading:

Column	What it shows
About	The record, or the proposed name of a new activity (marked **New activity**). Click it to open the questionnaire.
Questionnaire	The questionnaire used, marked "annual check" for annual checks.
Data Liaison	Who was asked.
Status	See the table below.
Answer by	The answer-by date, in red with "(overdue)" once it has passed.
Sent	When it was sent.

Filter the list by status (open ones, meaning awaiting answers or review, are shown by default), by Data Liaison, by questionnaire, and with **Overdue only**. With dataPrivacy:read at "all" you see every questionnaire sent; at "own", those about records you own and those you sent. The **Data mapping questionnaires outstanding** dashboard widget summarizes them.

Status	Meaning
Awaiting answers	Sent; the Data Liaison has not submitted yet.
Awaiting review	Submitted; waiting for the privacy team to decide.
Applied	Decided; accepted values were written to the record (or a new record was created).
Rejected	The answers were rejected with a reason; nothing changed.
Canceled	Withdrawn before it was decided.

Questionnaires lists the questionnaires that can be sent, with the number of questions, how many fill record fields, whether each is **Available** or **Withdrawn**, and when it was last updated.

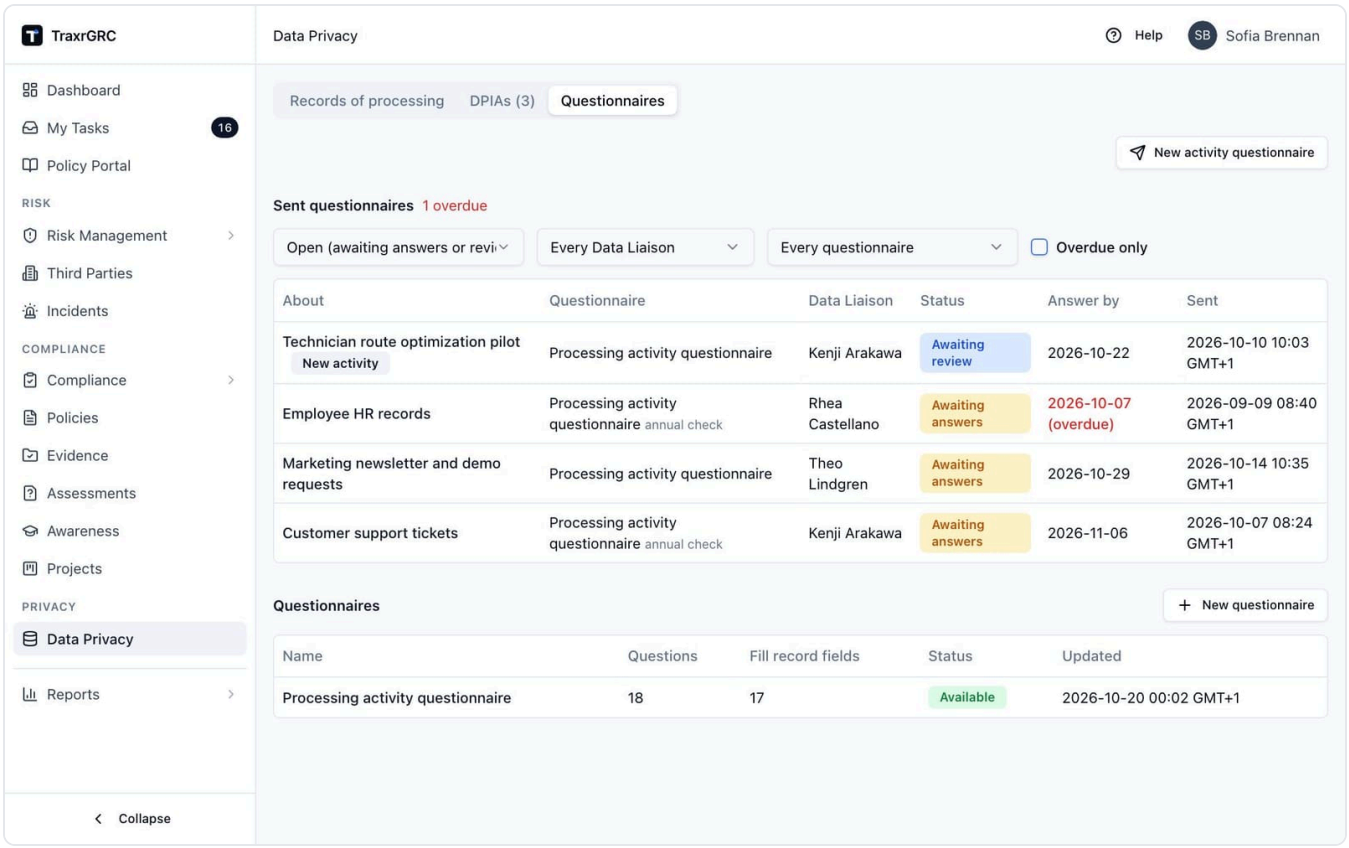


Figure 11.9. The Questionnaires tab, with the sent questionnaires and their statuses above the list of questionnaires available to send.

11.15.3 Building a questionnaire

The default **Processing activity questionnaire** asks for the Article 30 details in plain words. **New questionnaire**, or a click on a questionnaire's name, opens the editor. Changing questionnaires needs dataPrivacy:update at the "all" level; everyone with dataPrivacy:read can look at them.

1. Give the questionnaire a **Name** and an optional **Description**.
2. Click **Add question** for each question, and set:
    - **Question**: the text the Data Liaison reads.
    - **Help text**: optional guidance shown with it.
    - **Answer type**: see the table below. For **Pick from a list in the app**, choose the **List**; for single and multiple choice, enter **Choices, one per line** (at least two).
    - **Fills record field**: the record field the answer fills, or **Not mapped (kept with the answers)**.
    - **Required**: whether the question must be answered.
3. Click **Save**. The editor lists anything that still needs fixing and keeps **Save** disabled until it is right.

Answer type	What the Data Liaison sees
Short text, Long text	A text box. Long text that fills a list field takes one entry per line.
Yes or no	Two choices.
Single choice, Multiple choice	The choices you list, one per line (at least two).
Pick from a list in the app	Purposes, categories of personal data or of data subjects, sources and recipients (typed entries; the category lists and sources suggest values other records already use), countries (typed, with suggestions from countries already recorded on transfers, sub-processors, transfer assessments and systems), special categories, legal basis and processing role (fixed lists), or vendors and systems (picked by name).

A question can fill only a field that suits its answer type, and each field from one question at most; questions that fill no field are kept with the answers for the privacy team to read. Clear **Available to send** to withdraw a questionnaire without deleting it. **Delete** removes a questionnaire and needs dataPrivacy:delete at the "all" level. Requests already sent keep the questions they were sent with.

11.15.4 Sending a questionnaire

1. Start from one of two places:
    - On a record, click **Send questionnaire** (dataPrivacy:update on the record; not offered on archived records).
    - On the Questionnaires tab, click **New activity questionnaire** (dataPrivacy:create) for an activity that has no record yet, and enter the **Name of the new activity**.
2. Choose the **Questionnaire** (only available ones are offered).
3. Choose the **Data Liaison**. On a record, the record's own Data Liaison is preselected.
4. Optionally set an **Answer by** date and a **Message to the Data Liaison**.
5. Click **Send**.

The Data Liaison gets an email with a link. A record's Data Liaison is set in **Edit**; it is the default respondent and receives the annual checks.

11.15.5 Answering (the Data Liaison's view)

The Data Liaison opens **My Questionnaires** in the navigation. The page lists their questionnaires with what each is about, the questionnaire, its status and the answer-by date (red once overdue), open ones

first.

1. Click the activity name to open the questionnaire. The header shows the questionnaire, who sent it, the answer-by date and any message.
2. Answer the questions. Answers start from what the record holds today.
3. Click **Save draft** to keep work in progress, or **Submit answers** to send them for review once every required question is answered.

Whoever sent the questionnaire gets an email. After submitting, the questionnaire is read-only and says the answers have been sent to the privacy team. A Data Liaison sees only their own questionnaires; any other link shows "Page not found".

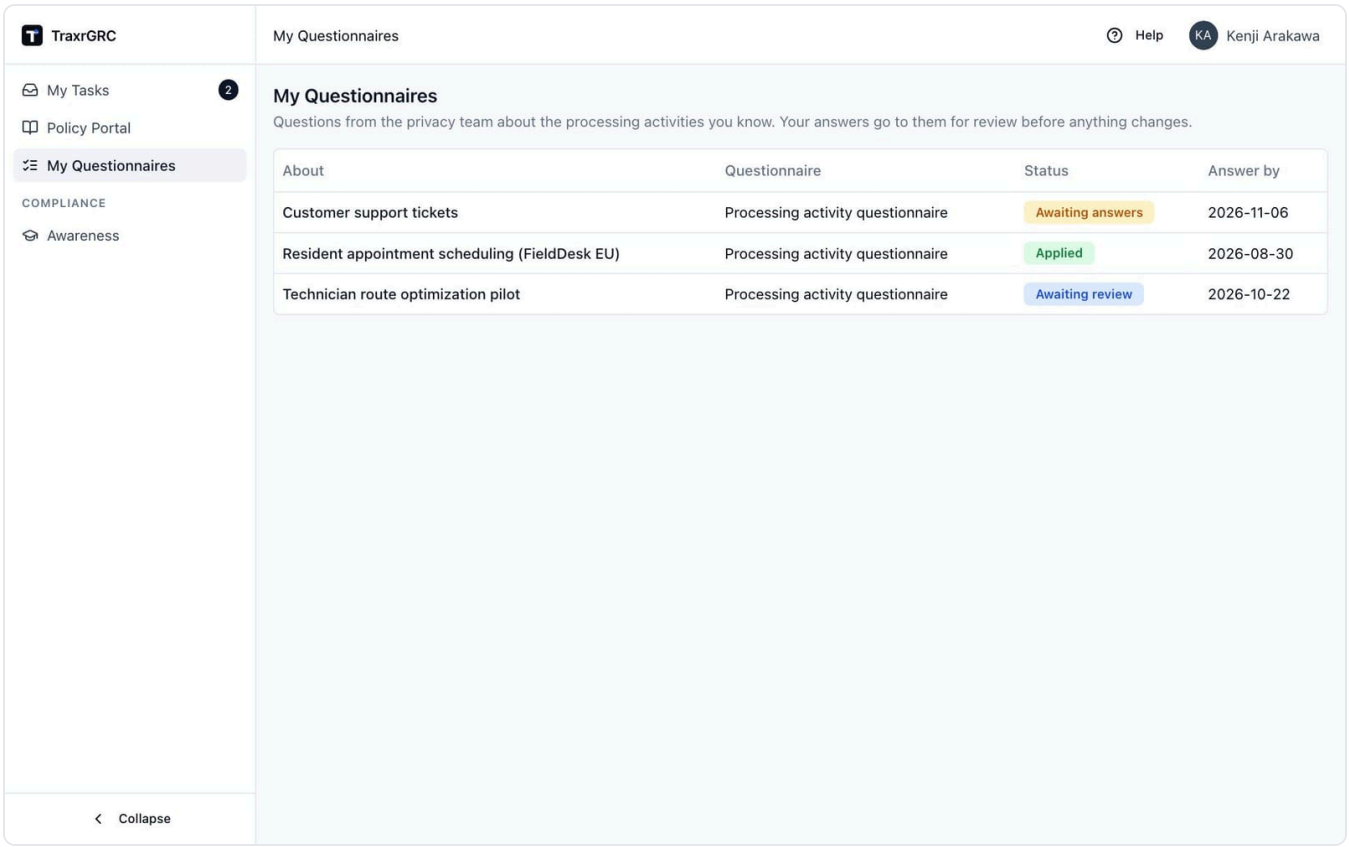


Figure 11.10. My Questionnaires as a Data Liaison sees it, listing an open annual check and earlier questionnaires with their statuses.

11.15.6 Reviewing answers

Open a sent questionnaire from the Questionnaires tab, the record's **Questionnaires** tab or the email. The page shows who it was sent to and when it was answered, with **Open the record** for questionnaires about an existing record.

Proposed changes (*N*) lists each field the answers would change, **Now** against **Proposed**. For each one choose:

Decision	Effect
Accept	Write the proposed value.
Edit	Change the value before it is written.
Reject	Keep the record's current value.

Unchanged fields lists the fields the answers leave as they are, and **Other answers** shows answers to questions that fill no field.

1. Decide each proposed change.
2. Add **Notes** if you wish (notes are required to reject the answers).
3. Click **Apply decisions** to write the accepted and edited values to the record; its History names the questionnaire. For a new activity the button reads **Create record**: it creates a draft record from the answers, with the Data Liaison set.

Alternatively, **Reject answers** (with notes giving the reason) changes nothing, and **Cancel questionnaire** withdraws one that has not been decided. Once decided, the page shows who decided it and when, with each field's decision.

Deciding needs dataPrivacy:review and the right to update the record, and is never open to the Data Liaison who gave the answers, administrators included, unless an administrator has relaxed separation of duties (see Sessions under Settings). Anyone else who can read the record sees the answers only.

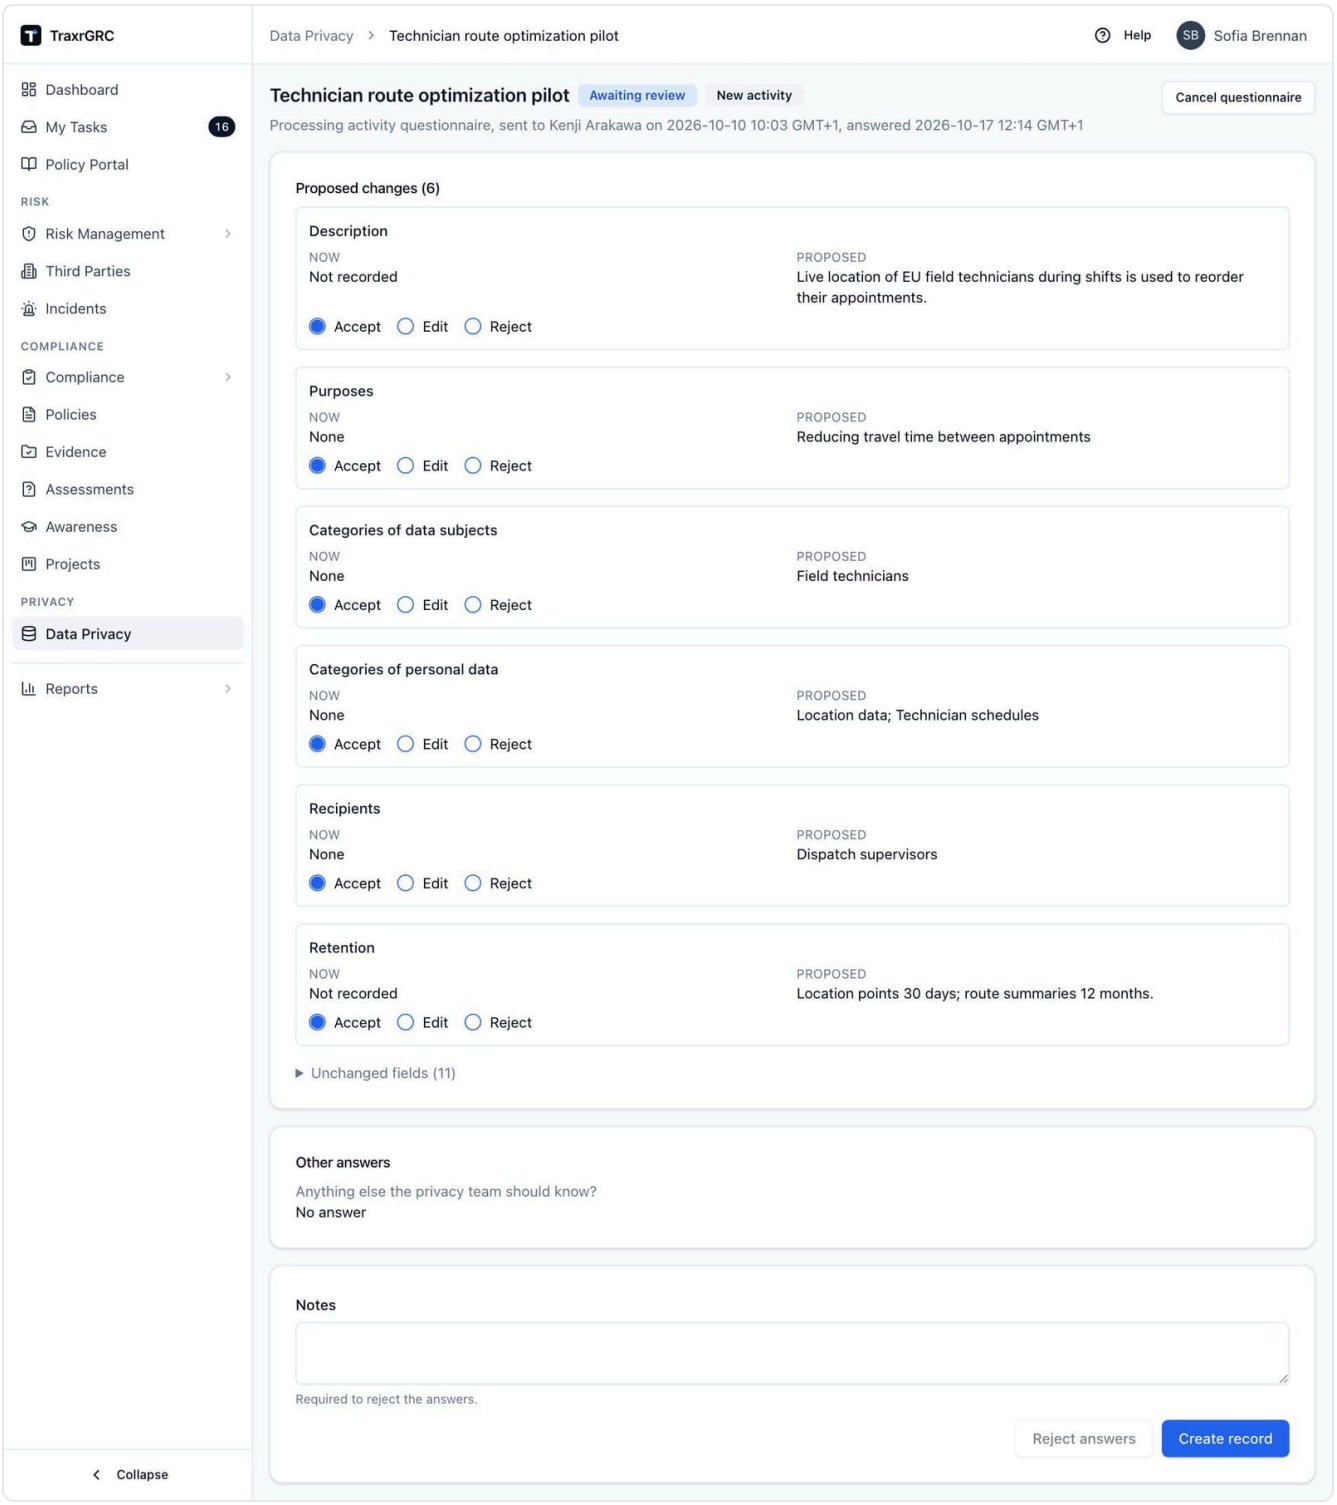


Figure 11.11. A submitted questionnaire about a new activity awaiting review, with each proposed value and the Accept, Edit and Reject choices.

11.15.7 Annual checks

Thirty days before an open review on a record that has an active Data Liaison, the daily sweep sends them an annual check: the record's last applied questionnaire (or the default one), prefilled and due on

the review date. It is marked **Annual check** throughout.

- **Nothing has changed** asks for confirmation (any changes saved as a draft are not sent), then confirms the record and completes the review with the Data Liaison as reviewer.
- Changed answers go to review like any other; applying them completes the review with the reviewer who applied them.

Each review gets one check, so a rejected or canceled check is not sent again; complete that review by hand. Completing a review by hand cancels its open annual check, so the Data Liaison is not asked to confirm a review that is already done.

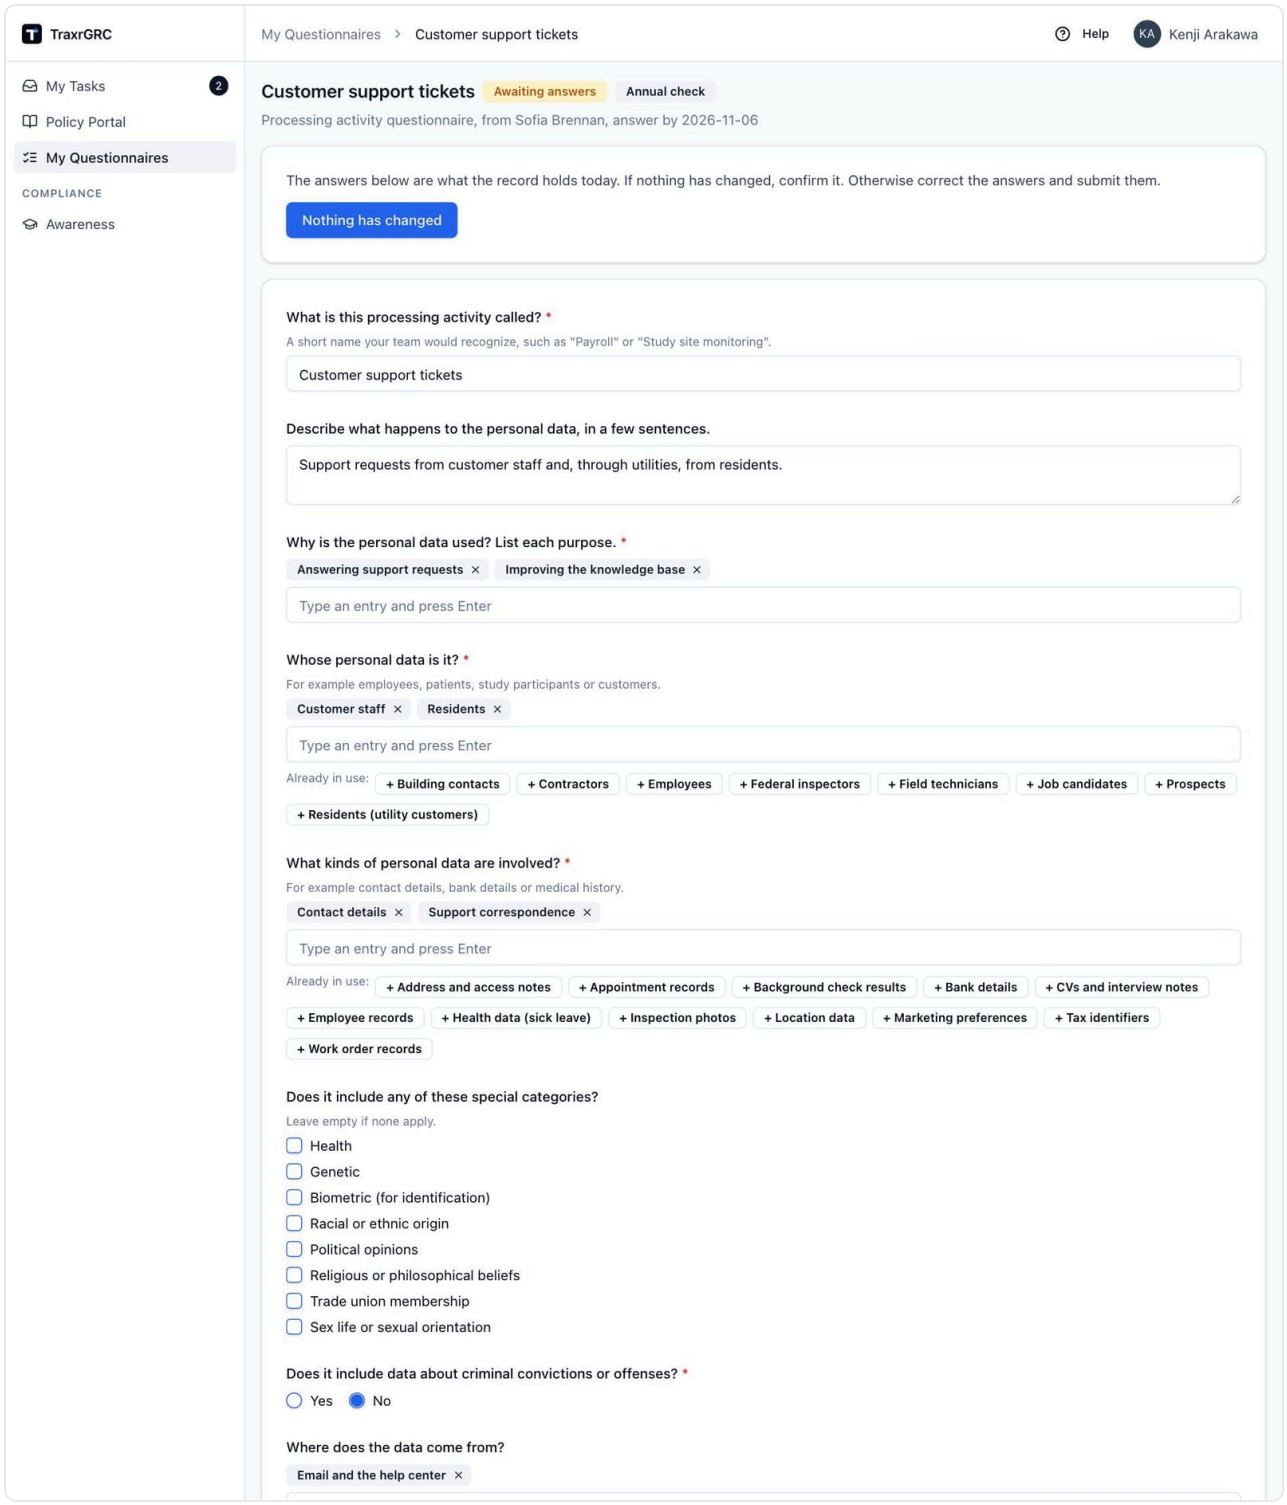


Figure 11.12. An annual check opened by its Data Liaison, prefilled from the record, with the Nothing has changed button.

11.15.8 Reminders

The daily sweep emails the Data Liaison once when a questionnaire's answer-by date is within 7 days and once when it has passed. A new answer-by date starts again.

11.16 The data map

Data map on the Data Privacy page draws where personal data lives and flows, in four columns:

Column	What it shows
Systems	The assets linked to records. A system's hosting country is written on the system.
Records of processing	The records themselves.
Vendors	The third parties that process the data for them.
Countries	The countries the data goes to, from transfers and the transfer summary; entries such as "None" are ignored.

Each item sits beside the records it touches. Point at an item, or move to it with Tab, to light up its links, and click one to open it.

Filter the map by record status (**Draft and active** by default, or Active, Draft, Archived or Every status), by data category, by vendor, or with **Special category data only**. A line under the filters counts the records shown and names what is left off because you cannot read it. On a narrow screen the map scrolls sideways inside its own box. Opening the map needs dataPrivacy:read; systems need assets:read and vendors thirdParties:read.

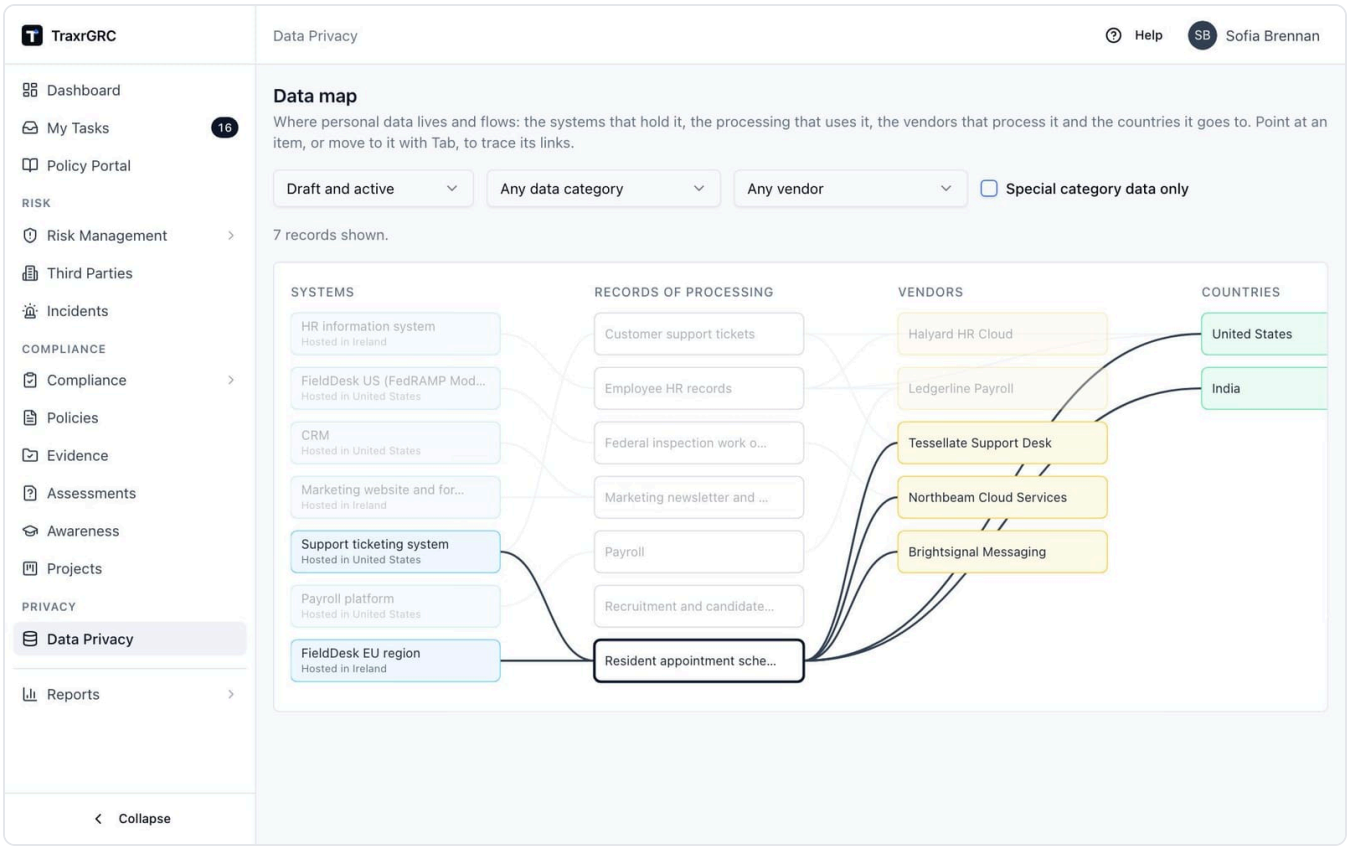


Figure 11.13. The data map, tracing personal data from systems through records of processing to vendors and destination countries.

11.17 Data privacy permissions

Action	Permission needed
See the list and open a record	dataPrivacy:read
Create a record	dataPrivacy:create
Edit documentation, change status, complete or schedule reviews	dataPrivacy:update on the record; at "own" level you must own it (or its owner group)
Post comments, upload attachments	dataPrivacy:update on the record
Add or remove transfers, link assets, risks and controls	dataPrivacy:update on the record, plus read on the linked module
Start, screen, edit, submit or reopen a DPIA	dataPrivacy:update on its record
Give DPO advice on a DPIA	Being the Data Protection Officer named in Settings → Privacy, with dataPrivacy:read on its record
Approve or reject a DPIA	dataPrivacy:approve on its record; DPO advice first when it is required; not its creator or submitter while "Require a second person to approve DPIAs" is on, unless separation of duties is relaxed
Delete a record or a DPIA	dataPrivacy:delete on the record
Download the Article 30 register	dataPrivacy:read (the register lists only records you can read)
Open the data map	dataPrivacy:read (systems need assets:read, vendors thirdParties:read)
Create or change data mapping questionnaires	dataPrivacy:update at the "all" level
Delete a data mapping questionnaire	dataPrivacy:delete at the "all" level
Send a questionnaire about a record, or cancel one	dataPrivacy:update on the record

Action	Permission needed
Send a questionnaire about a new activity	dataPrivacy:create
Answer a questionnaire	Being its Data Liaison, with the Data Liaison role (dataLiaison:update)
Accept, edit or reject a Data Liaison's answers (never your own)	dataPrivacy:review, plus dataPrivacy:update on the record (dataPrivacy:create for a new activity)

Roles are assigned to groups and inherited through group membership; administrators bypass all of the above (see Managing Users, Groups & Roles).

CHAPTER 12

Account Reviews

Account reviews are periodic attestations that the accounts on your systems should still exist, for example a quarterly privileged access review in which a manager confirms every administrator account on every system. An account review program groups the accounts that must be reviewed together. Reviewers work through the list, keeping the access that is still needed and disabling the rest, and TraxrGRC records the reviewer's name and the time on every decision.

The module is used by the people who run access reviews (typically the GRC team or IT security), by the system and team owners who make the decisions, and by auditors who read the result as evidence. This chapter covers creating programs, loading accounts one by one or from CSV, recording decisions, and exporting the evidence.

12.1 The programs list

Open **Account Reviews** from the navigation. The page lists every program you can see.

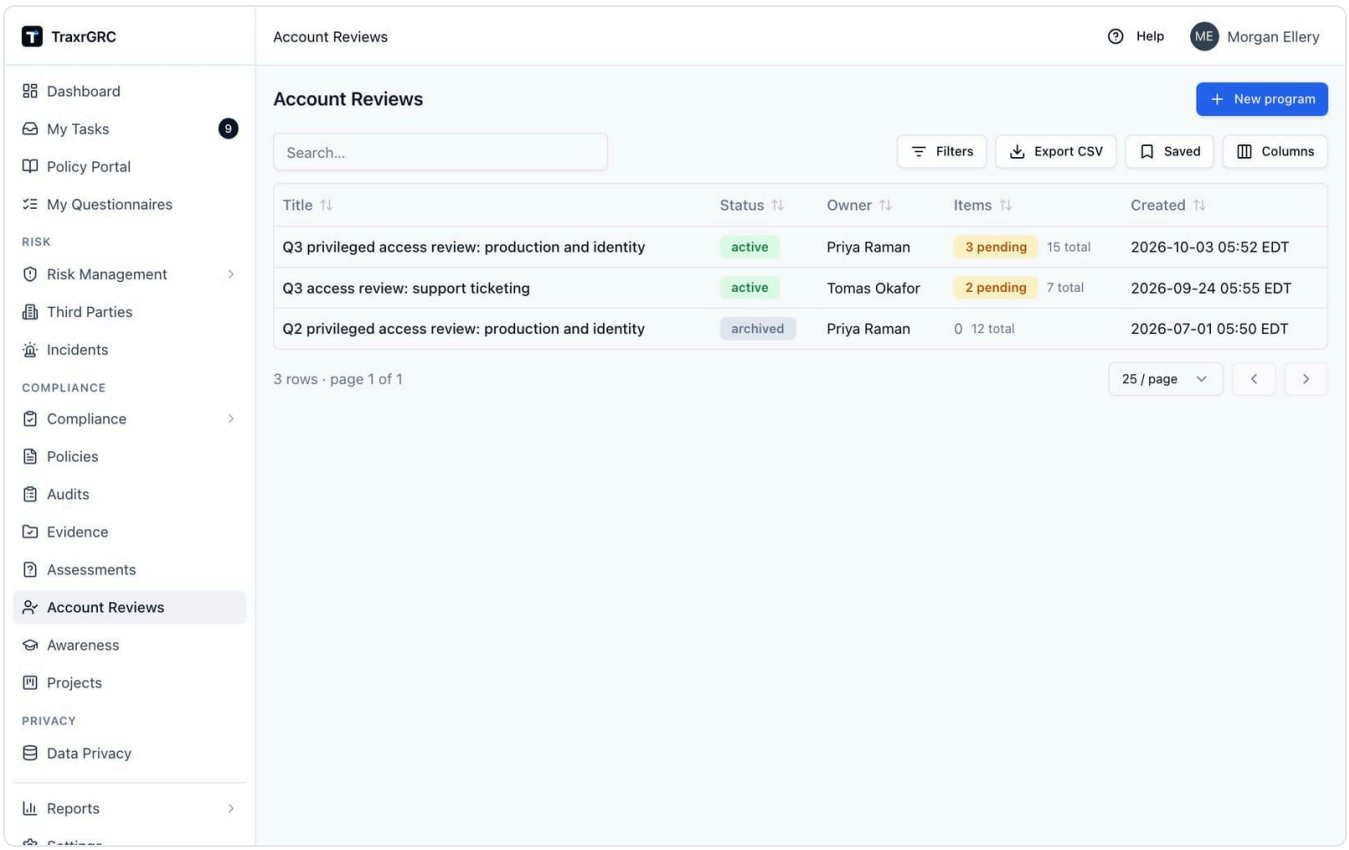


Figure 12.1. The Account Reviews list with the two Q3 programs and the archived Q2 program, showing pending and total item counts.

Column	What it shows
Title	The program's name. Click it to open the program.
Status	`Active` (green badge) or `Archived` (gray badge). Filterable.
Owner	The owning person, or the owning group's name when no person is set.
Items	An amber **N pending** badge plus the total item count (for example "3 pending 15 total"). A zero pending count shows as a plain `0`.
Created	The date and time the program was created.

The search box filters by title or description. **Export CSV** downloads the programs that match your current filters, which is useful as review evidence. **Saved** stores and reapplies filter sets, as on every list in TraxrGRC.

The dashboard widget **Account review items** charts pending, kept and disabled items across all the programs you can see (see Getting Started, The Dashboard).

12.2 Creating a program

1. Click **New program**. The **New account review program** dialog opens.
2. Complete the fields:

Field	What it means
Title	Required. The program's name, for example "Quarterly privileged access review".
Description	Free text explaining scope, frequency or method.
Owner	The person accountable for the review. Defaults to "None".
Owner group	A group accountable for the review. You can set either or both.

3. Click **Create program**.

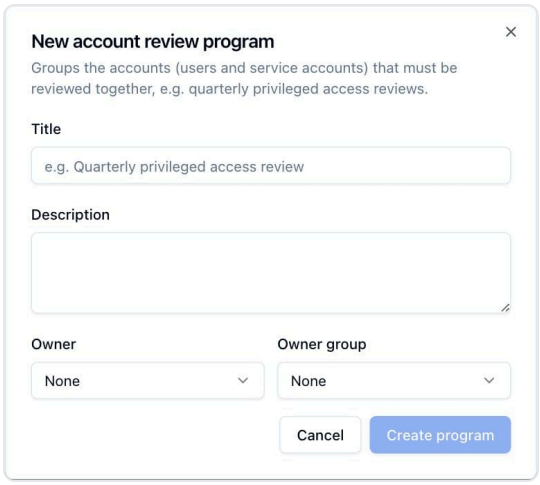


Figure 12.2. The New account review program dialog with its Title, Description, Owner and Owner group fields.

Creating, editing and deleting programs is recorded in the audit log (see Administration & Settings).

12.3 The program page

Open a program from the list.

Under the title and status badge, four summary cards show:

Card	What it shows
Owner	The owning person, with the owning group underneath.
Pending items	The number of items still to decide (amber when above zero), with the kept and disabled counts underneath.
Status	`Active` or `Archived`.
Created	When the program was created.

Below the cards are the **Items**, **Comments**, **Attachments** and **History** tabs, followed by **Custom Fields**, which shows any custom fields your administrator has defined for review programs. **Items** opens by default.

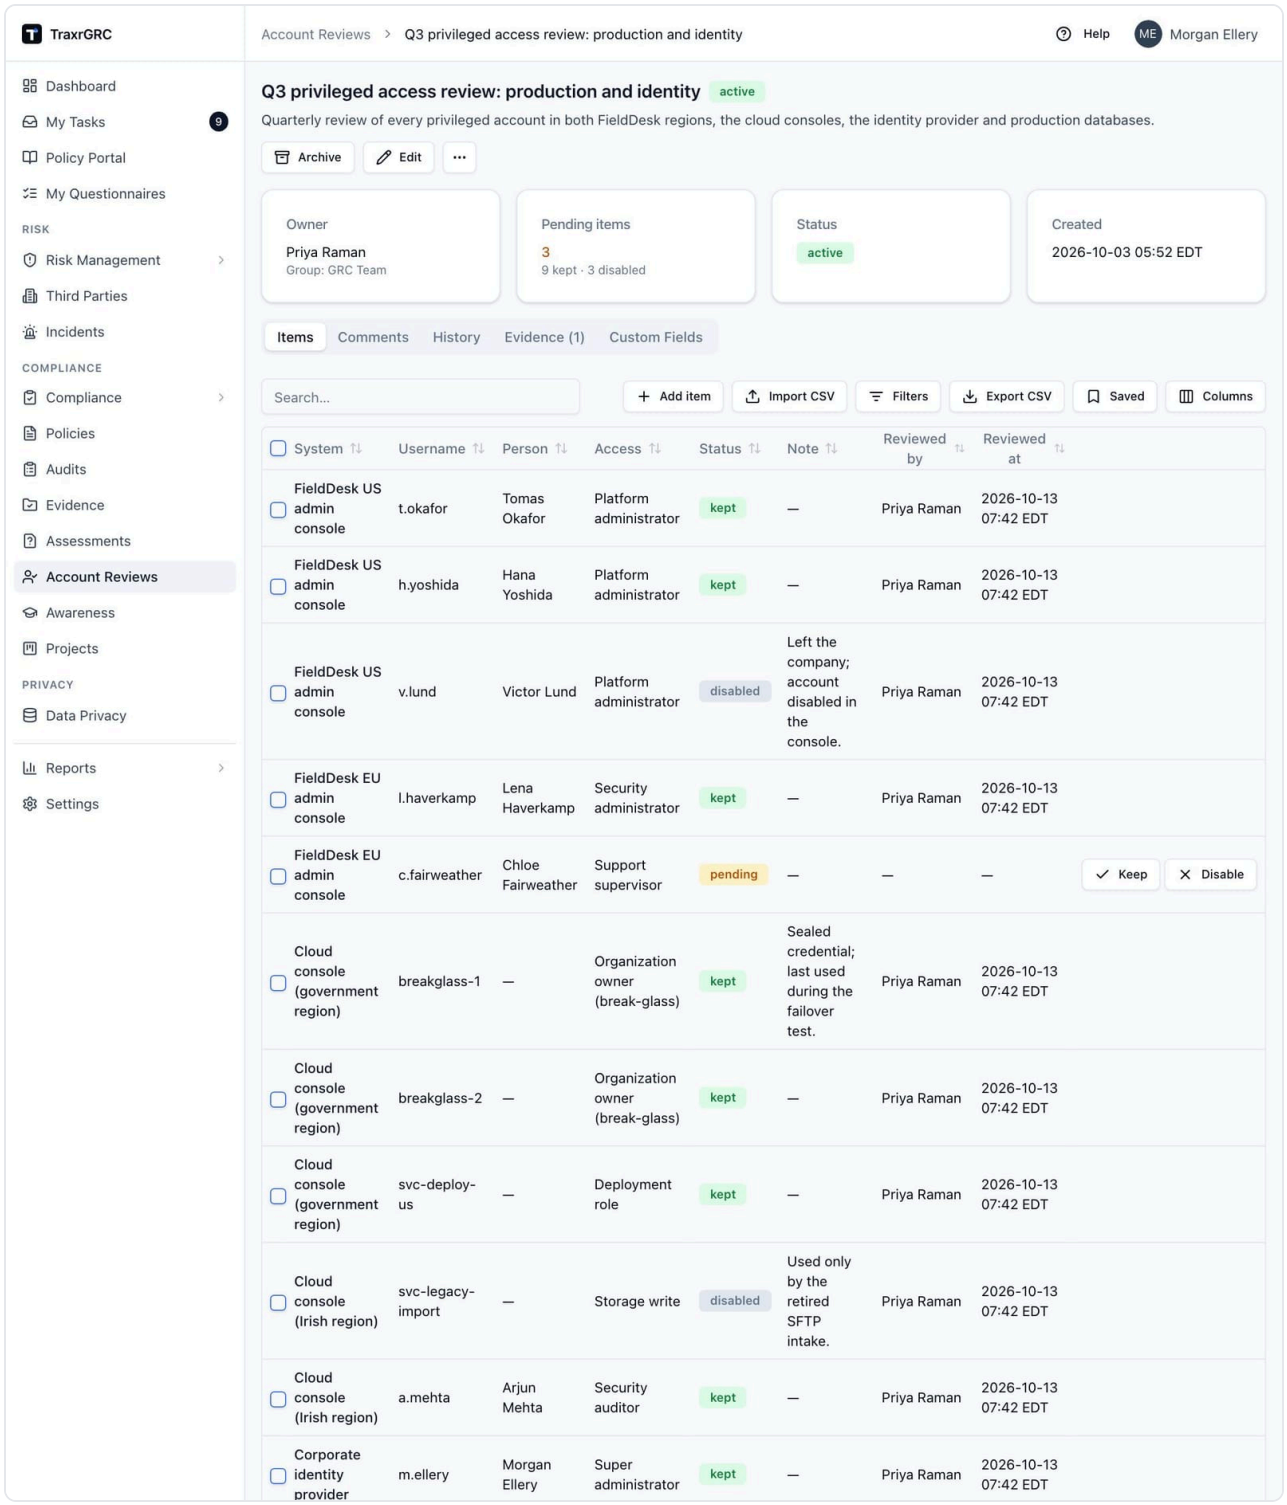


Figure 12.3. The Q3 privileged access review program on its Items tab, with the summary cards and a mix of pending, kept and disabled accounts.

The item table has these columns: **System**, **Username**, **Person**, **Access**, **Status**, **Note**, **Reviewed by** and **Reviewed at**. With update rights you also see a selection checkbox on each row and, on pending rows, the **Keep** and **Disable** buttons. You can filter by status (**Pending**, **Kept**, **Disabled**) and search by system, username or person.

12.4 Adding items to a program

You need the `accountReviews:update` permission to add or import items. Both buttons are in the toolbar above the item table on the **Items** tab.

12.4.1 Adding a single account

1. Click **Add item**. The **Add account** dialog opens.
2. Complete the fields:

Field	What it means
System	Required. The system the account exists on, for example "GitHub".
Username	Required. The account's login, for example "jsnow".
Person	Optional. The person behind the account, for example "Jane Snow".
Access	Optional. A short summary of what the account can do, for example "Organization owner".

3. Click **Add item**. The new item starts as **Pending**.

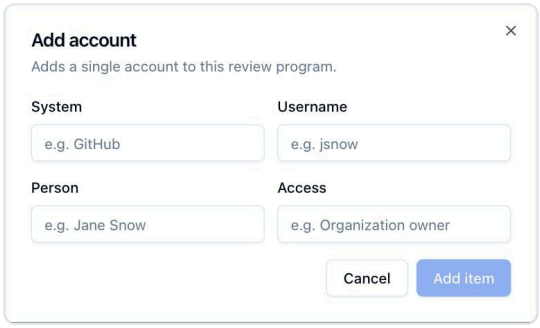


Figure 12.4. The Add account dialog with the System, Username, Person and Access fields.

12.4.2 Importing accounts from CSV

1. Click **Import CSV**. The **Import accounts from CSV** dialog opens.
2. Either choose a `.csv` file or paste CSV text into the **Or paste CSV here** box. A chosen file's contents appear in the box, so you can check them before importing.
3. Click **Import items**. A successful import reports how many items were added; every imported item starts as **Pending**.

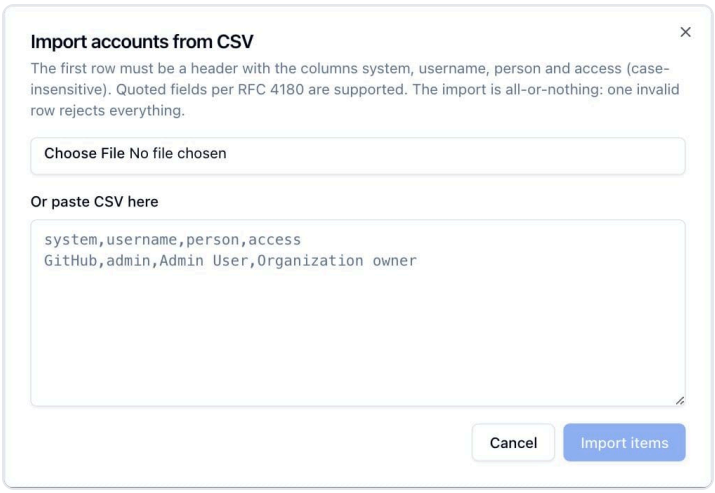


Figure 12.5. The Import accounts from CSV dialog, with a file picker and the paste box showing the expected header row.

The format rules:

- The first row must be a header naming the columns `system`, `username`, `person` and `access`. Matching ignores case, spaces and underscores, and `personname` and `accesssummary` are also accepted. Other columns are ignored.
- `system` and `username` are required on every data row; `person` and `access` may be empty.
- Quoted fields are supported (standard RFC 4180 CSV quoting).
- Completely blank lines are ignored.
- The import is **all-or-nothing**: every row must be valid or nothing is imported. Problems are listed per line with 1-based line numbers, where the header row is line 1 (for example `Line 4: username: ...`). Fix every problem and try again.

12.5 Reviewing items: keep or disable

Each item is decided as **kept** (the access is still required) or **disabled** (the access should be removed). You need `accountReviews:update` to record either decision, per row or in bulk.

12.5.1 Deciding one item

1. On a pending row, click **Keep** or **Disable**. The **Keep access** or **Disable access** dialog opens.
2. Optionally enter a **Note**. A note is recommended; the placeholders show the level of detail expected ("e.g. Access still required for the role." or "e.g. Left the team; access revoked in the IdP.").
3. Click **Mark kept** or **Mark disabled**.

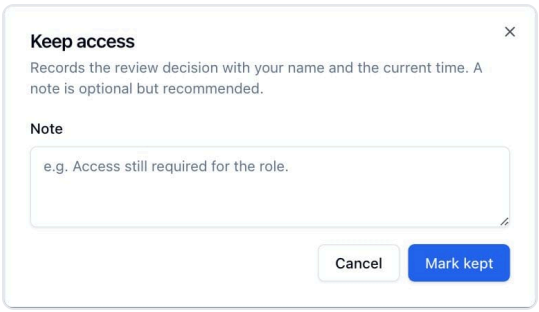


Figure 12.6. The Keep access dialog for a pending account, with the optional note field.

12.5.2 Deciding items in bulk

1. Tick the checkboxes on the rows you want. The header checkbox selects every row on the current page.
2. Click **Mark kept (N)** or **Mark disabled (N)** in the toolbar. A badge next to the buttons shows how many rows are selected.

Bulk decisions do not ask for a note. A bulk action applies to every selected row, including rows already decided: the new decision replaces the earlier one, with your name and the current time, and clears any earlier note.

12.5.3 What a decision records

Every decision records your name and the current time on the item, plus the note if you entered one. The item's status badge changes to `Kept` (green) or `Disabled` (gray), the **Note**, **Reviewed by** and **Reviewed at** columns fill in, and the per-row buttons disappear (they appear only while an item is pending). The decision is written to the audit log against the program, and the pending counts on the program's summary card and on the programs list update immediately.

Recording a decision in TraxrGRC does not change the account on the source system. Disabling the account in the system itself remains your team's task; the note is the place to record that it was done.

Export CSV above the item table downloads all items matching the current filters, including who reviewed each one and when. This file is the usual evidence artifact for an auditor.

The web interface has no button for deleting a single item. An item added by mistake can be removed through the REST API (see REST API), or decided with a note explaining the error.

12.6 Comments, attachments and history

Every program has three more tabs:

- **Comments**: a discussion thread on the program, for example to chase a reviewer for outstanding decisions. Adding comments requires update access to the program.

- **Attachments**: files such as exports from your identity provider or sign-off sheets. Uploading requires update access.
- **History**: the audit entries for this program, newest first. Each row shows the action (`create`, `update`, `keep`, `disable` and so on), who did it, when, and which fields changed.

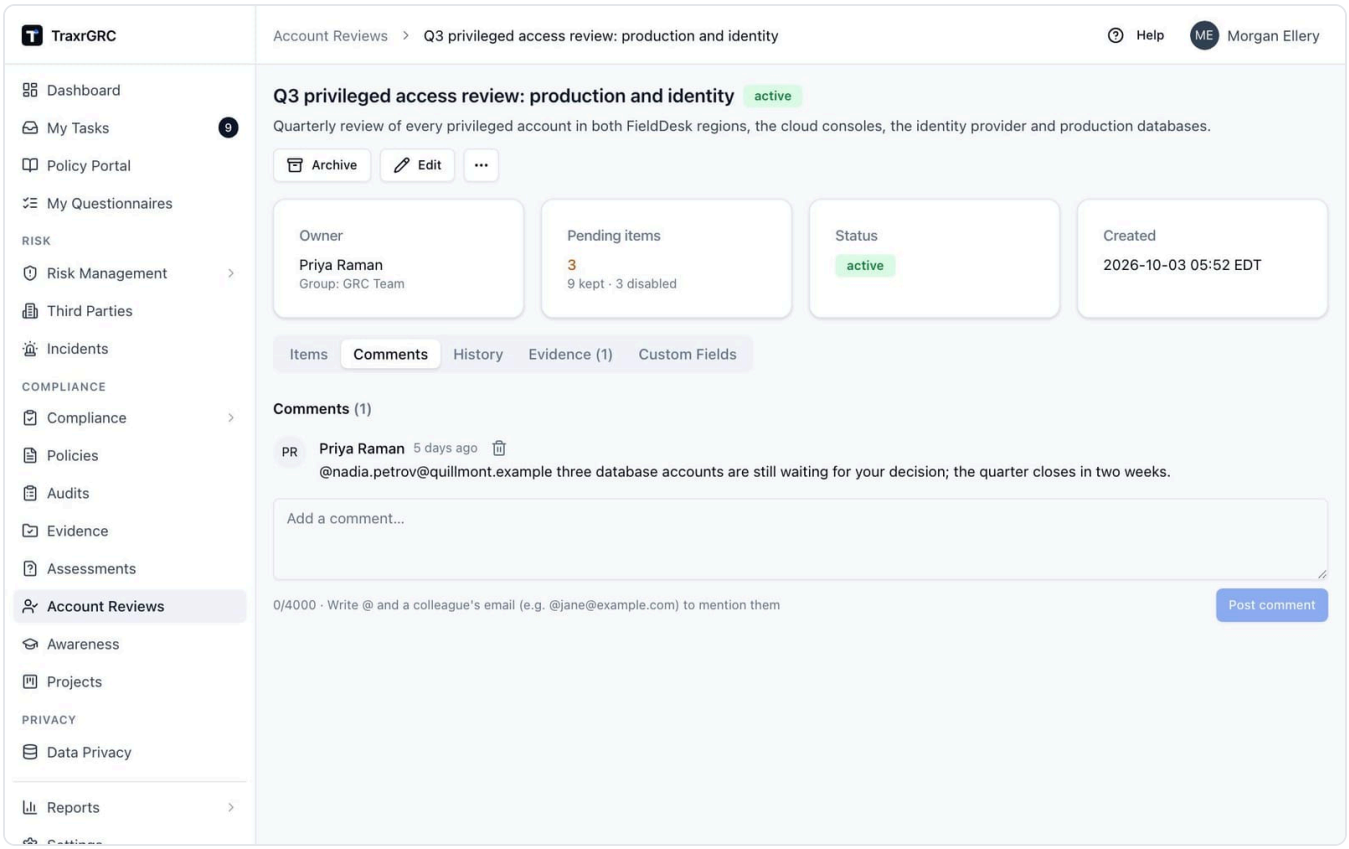


Figure 12.7. The Comments tab of the Q3 privileged access review, with a reminder that three database accounts still await a decision.

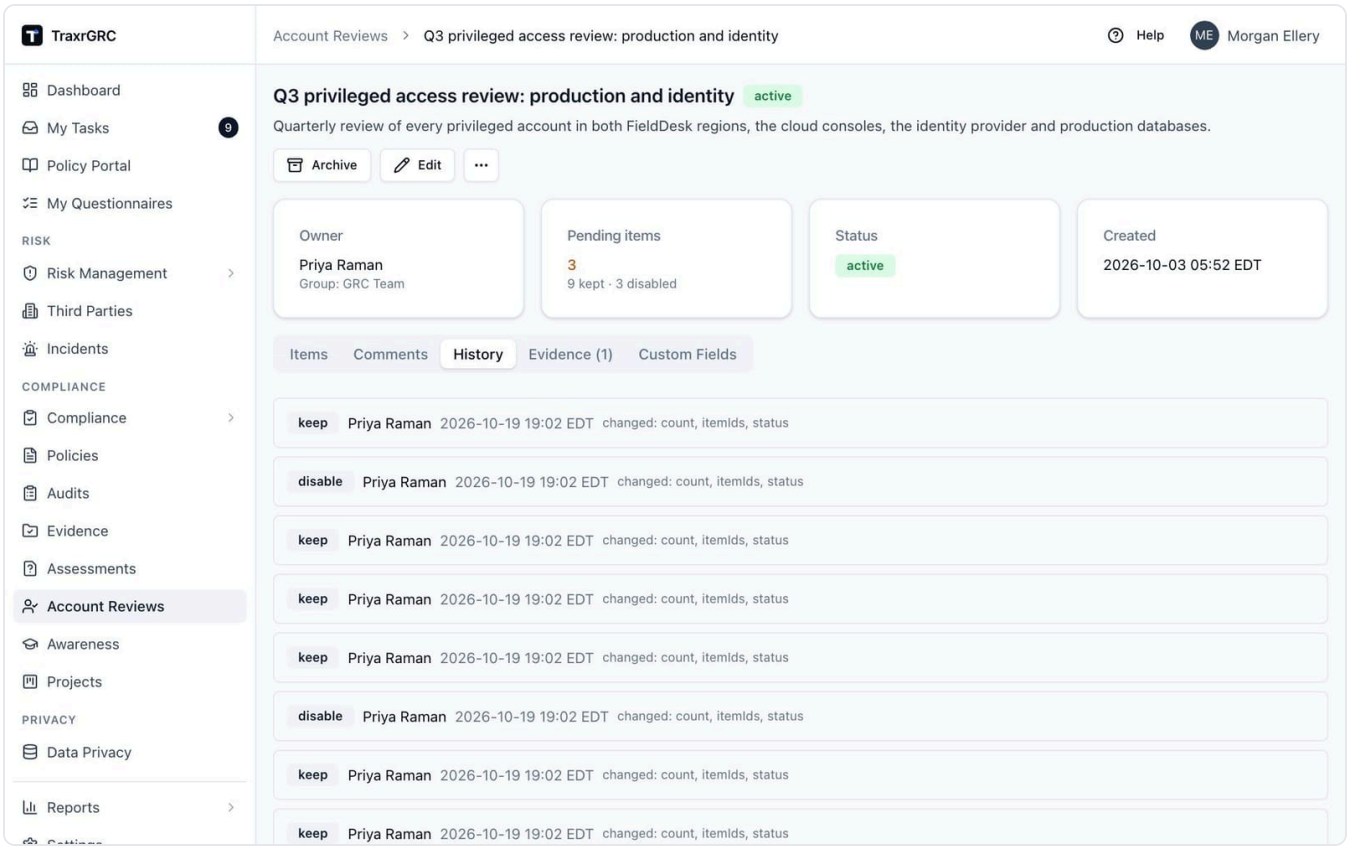


Figure 12.8. The History tab of the Q3 privileged access review, listing each keep and disable decision with its actor and time.

12.7 Archiving, reactivating and deleting

The buttons at the top right of a program's page depend on your rights.

With `accountReviews:update`:

- **Archive** (on an active program) marks a finished program as archived; its status badge turns gray. Nothing is deleted, and the program stays visible in the list under the `Archived` status.
- **Reactivate** (on an archived program) makes it active again.
- **Edit** opens the **Edit review program** dialog to change the title, description, owner or owner group. Click **Save changes** to apply.

Archiving does not lock a program: anyone with update rights can still add and decide items on it.

With `accountReviews:delete`, **Delete** in the **More actions** menu (**...**) asks for confirmation ("Delete account review program?") and then permanently removes the program **together with all its review items**. This cannot be undone.

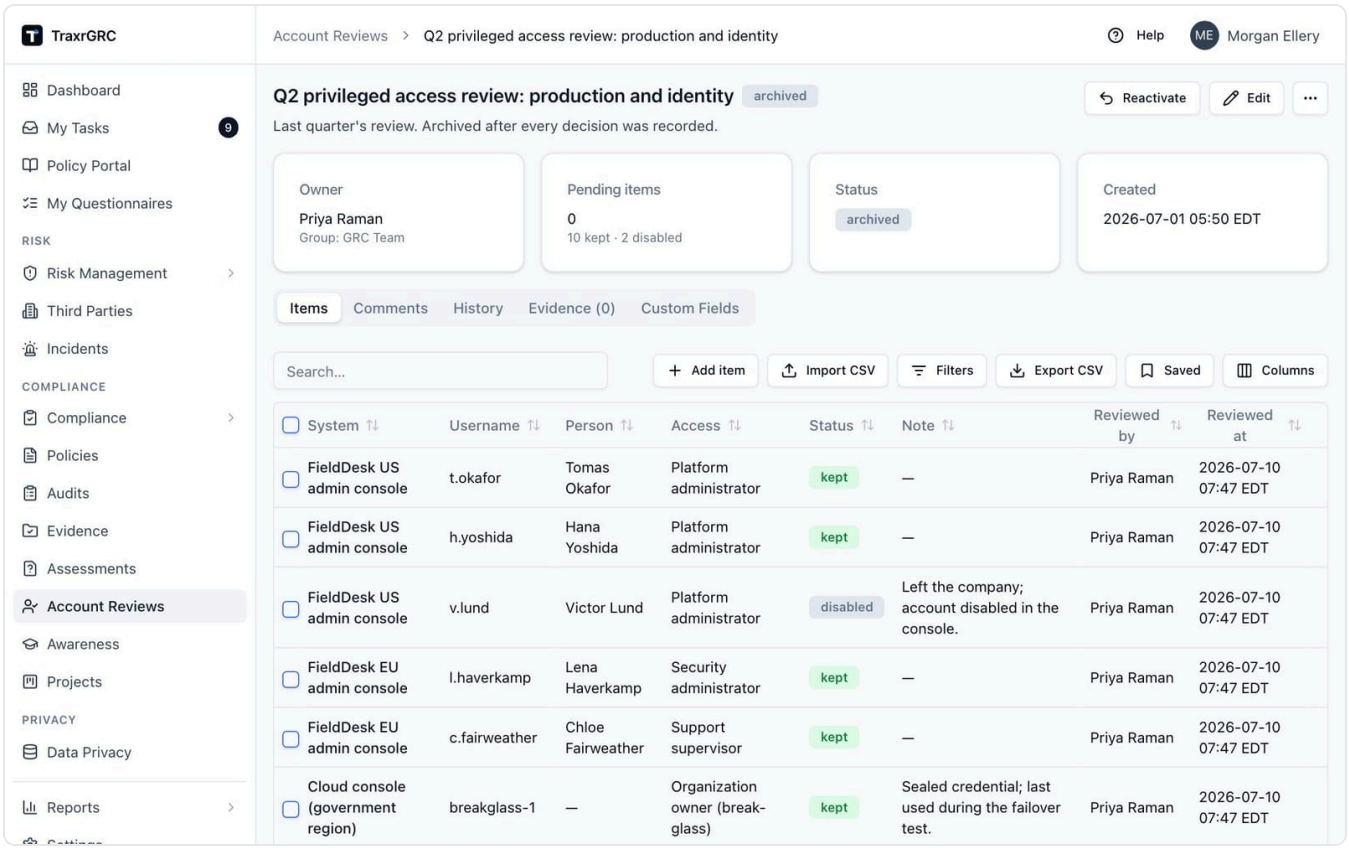


Figure 12.9. The archived Q2 privileged access review, showing the gray Archived badge and the Reactivate button.

12.8 Who can do what

Action	Permission
See the programs list and a program's items	`accountReviews:read` (at the `own` level you only see programs owned by you or your groups)
Create programs	`accountReviews:create`
Add or import items, keep or disable items, comment, attach, edit, archive, reactivate	`accountReviews:update` (at the `own` level, on programs you own)
Delete programs	`accountReviews:delete`

Administrators can do everything.

CHAPTER 13

Evidence

Evidence is what shows a control operating or a third party meeting its obligations: the screenshot of a quarterly access review, the notes from a monthly vulnerability scan review, a vendor's SOC 2 report or ISO certificate. Each piece of evidence is one record with its files, the period it covers, when it was collected and how long it stays valid, and links to every record it supports. One access review export can support a control, the test result it was collected for and several compliance requirements at once, and an auditor can open it from any of them.

There are two kinds:

Kind	Used for	Extra fields
Internal	Evidence your organization produces: access reviews, vulnerability reviews, configuration screenshots, log reviews, training records, meeting minutes, test results, and reports about your own organization (an **Audit or assessment report** such as your FISMA assessment, an internal audit or your own SOC 2 report, and a **Certification** such as your own ISO/IEC 27001 certificate)	None
Third-party report	Assurance a third party provides: SOC 1 and SOC 2 reports, ISO/IEC 27001, 27701 and 22301 certificates, PCI DSS attestations, HITRUST, FedRAMP packages, penetration test summaries, bridge letters, security questionnaires (SIG, CAIQ)	The third party, who issued it, the auditor's opinion, and exceptions noted with your response

13.1 The Evidence page

Open **Evidence** in the sidebar under Compliance. The table lists every piece of evidence you can see, newest first:

Column	What it shows
Title	Click to open the evidence.
Kind	Internal or Third-party report. Filter by it.
Type	Access review, SOC 2 Type II and so on; an Other type reads "Other: " and its description, for example "Other: DR test report". Filter by it.
Other type	Hidden by default: the description of an Other type. Filter by typing part of it to narrow "Other internal evidence" to one kind.
Third party	For a report, the third party it is about. Filter by typing part of the name.
Period covered	The dates the evidence covers, such as the audit period of a SOC 2 report or the quarter of an access review.
Collected	When it was produced or obtained (for a report, the report date).
Valid until	When it stops counting: a certificate's expiry, or when the next review or report is due.
Status	See the table below. Filter by Current, Expiring or Expired.
Supports	How many records it is linked to and how many files it holds.
Owner	Hidden by default; turn it on from the column menu.

The status comes from the dates:

Status	Meaning
Current	Valid until is more than 30 days away, or not set.
Expiring	Valid until is within 30 days.
Expired	Valid until has passed.
Superseded	Later evidence replaces it (its **Replaces** field names this one). See Replacing earlier evidence.

The status filter matches the dates only, so a superseded report still appears under the status its own dates give it. Searching the table also matches the description of an Other type.

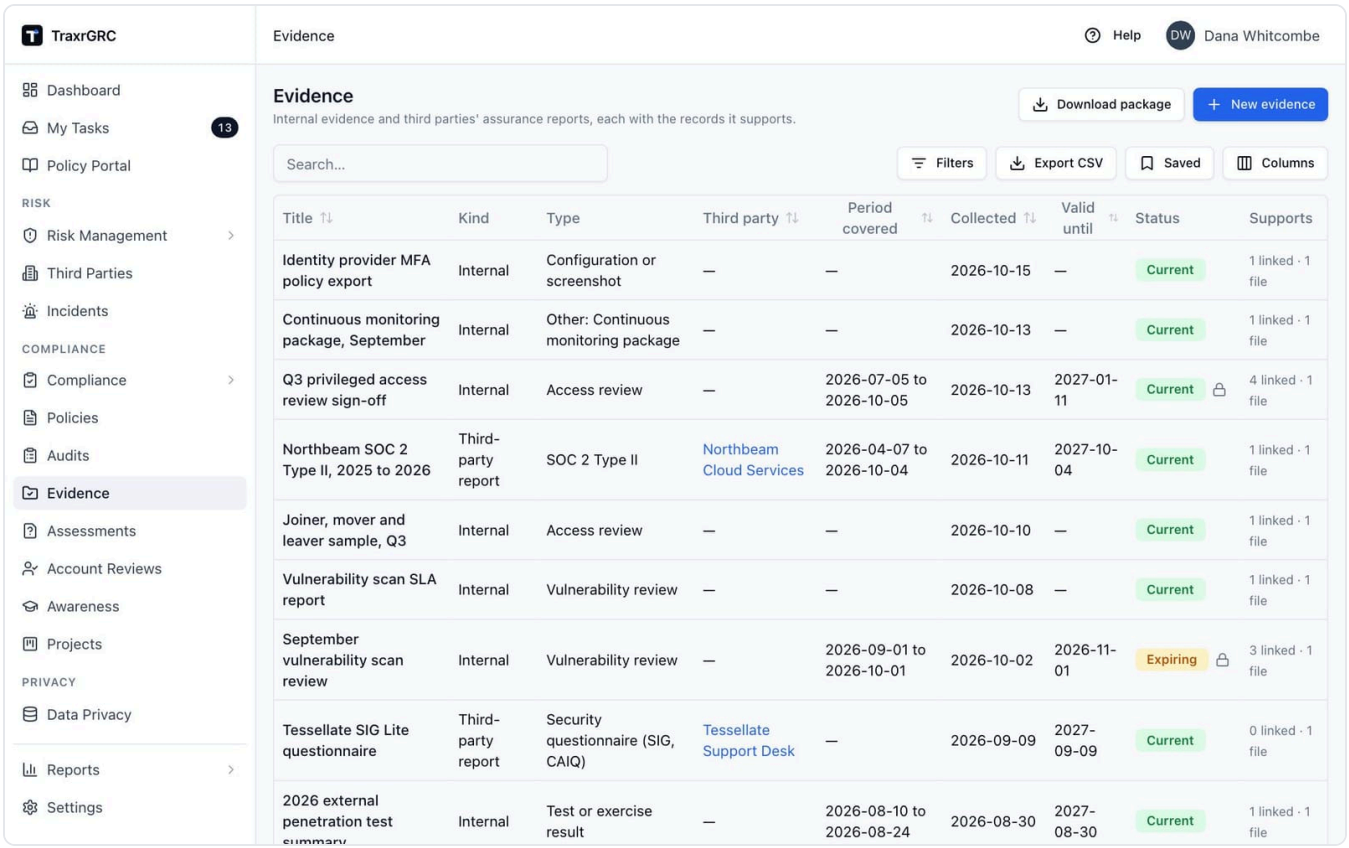


Figure 13.1. The Evidence page with internal evidence and third-party reports, their periods, valid-until dates and status, and the Download package button.

13.2 Adding evidence

1. Click **New evidence** on the Evidence page, or **Add evidence** on a control, requirement, account review or third party (see Evidence on other records). From a record, the new evidence is linked to that record.
2. Choose the **Kind** and the **Type**. For **Other internal evidence** or **Other third-party report**, **Describe the type** in a few words (for example "DR test report"); it is required, and the type then reads "Other: DR test report" wherever it appears.
3. Enter a **Title**, for example "Q3 privileged access review" or "SOC 2 Type II report, 2026".
4. For a third-party report, choose the **Third party** (required) and enter who **Issued by**: the audit firm or certification body.
5. Enter the **Period covered from** and **to** if the evidence covers a period, the date it was **Collected on** (for a report, the **Report date**), and **Valid until** if it expires.
6. For a third-party report, choose the **Auditor's opinion** (Unqualified, Qualified, Adverse or Disclaimer of opinion, or Not applicable for a certificate) and record **Exceptions noted and our response**: exceptions, nonconformities or complementary user entity controls that need follow-up.
7. Check **Replaces**: the earlier evidence this one replaces, if any (see Replacing earlier evidence). The form may have suggested one.

8. Add a **Description**, and choose the **Owner** (you, unless you choose someone else) and **Owner group**.
9. Click **Add evidence**. The evidence opens; add its files on the **Files** tab.

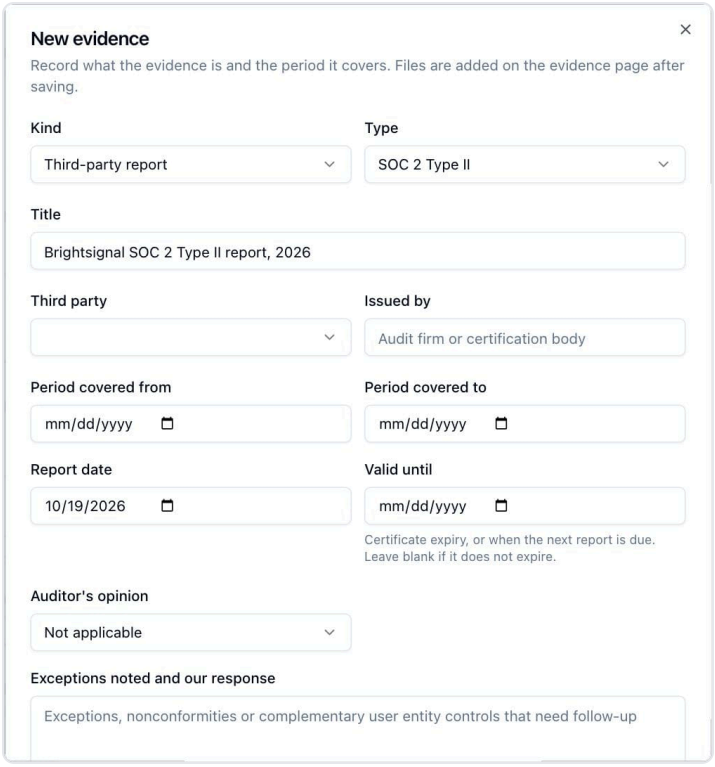


Figure 13.2. The New evidence dialog for a third-party report.

13.3 An evidence record

The evidence page shows its kind, type and third party, its status, and its dates: **Period covered**, **Collected on** (or **Report date**), **Valid until**, **Collected by** (who recorded it, and when), the **Owner** and, for a report, **Issued by**, **Auditor's opinion** and the exceptions noted. Below the type, "Superseded by" links to the evidence that replaces it and "Replaces" to the evidence it replaces (evidence you cannot read is named as such, without its title).

The **Supports** card lists every record the evidence is linked to, each labeled with its type (Control, Control test, Requirement, Risk, Audit finding, Policy or Account review) and linking to it. Records you cannot read are counted but not named. If you can edit the evidence, the buttons above the list link more records of each type (a searchable checklist, as for compliance mappings), and the **x** beside a record unlinks it. Linking needs only read access to the record; it changes the evidence, not the record.

The tabs below hold the **Files**, **Comments** and **History**. Every change, link, unlink, upload, lock, unlock, file check and deletion is recorded in the audit log.

Edit changes the fields (the kind is fixed once created). **Delete**, in the **More actions** menu (**...**), removes the evidence with its files, comments and links after you confirm; the records it supported are

not changed. Locked evidence cannot be deleted: the menu shows **Delete** disabled with "Locked evidence cannot be deleted. Unlock it first." (see Locked evidence and file fingerprints).

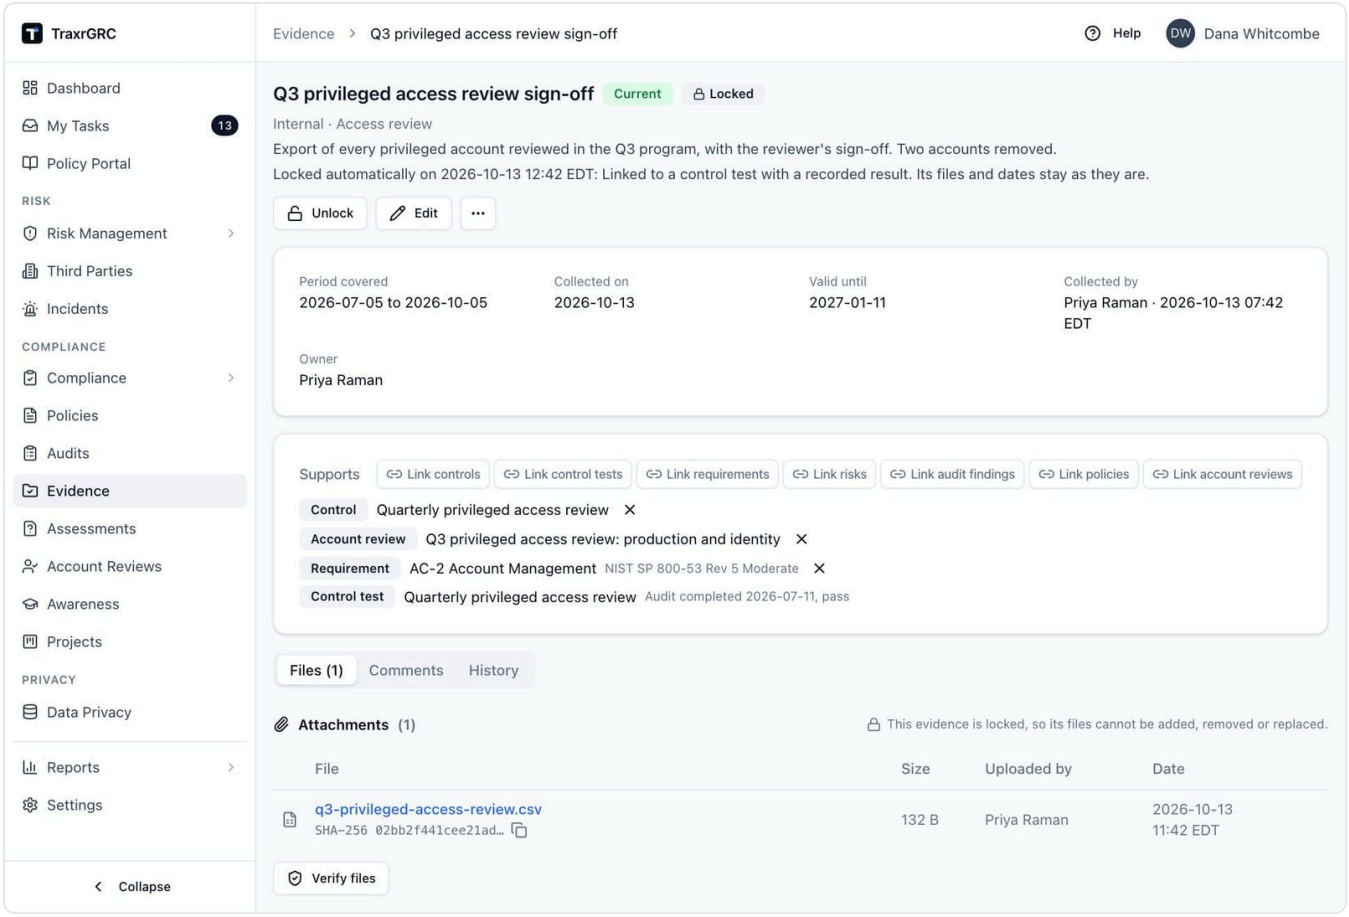


Figure 13.3. An access review's evidence record with the control, control test, account review and requirement it supports.

13.4 Evidence on other records

Evidence also appears where you work:

Record	Where	What it lists
Control	**Evidence** tab	Evidence linked to the control or to any of its tests (marked "via" the test). Each recorded result has an **Add evidence** button, and recording a result offers **Add evidence** in the confirmation, so the evidence is linked to that test and the control as you record it.
Compliance requirement	**Evidence** tab, **Evidence supporting this requirement**	Evidence linked to the requirement, and evidence of its mapped controls and their tests (marked "via" the control or test).
Account review	**Evidence** tab	Evidence linked to the review program. **Add evidence** starts an access review.
Third party	**Assurance reports** tab	Every report about the third party, with its status, so an expired or superseded report is visible at a glance. **Add evidence** starts a SOC 2 Type II report for this third party.

On a control, requirement or account review, **Link existing** links evidence you can edit that was recorded elsewhere.

13.4.1 Files on controls, requirements and account reviews

Controls, compliance requirements and account reviews take evidence rather than loose files: add a file through **Add evidence**, then the evidence's **Files** tab. A file sent to one of these records another way (through the AI assistant's upload tool or an upload link) becomes new evidence linked to the record, titled with the file name and typed "Other: Uploaded file", owned by whoever uploaded it; edit its type and dates afterwards. This needs `evidence:create` and read access to the record, the same as linking evidence.

Files attached to these records before they took evidence move onto evidence by themselves shortly after the update: each becomes evidence linked to the same record, titled with the file name, typed "Other: Migrated from control" (or requirement, or account review), collected on the day it was uploaded, and owned by the record's owner. The file keeps its fingerprint, and each move is recorded in the audit log. Until then the record shows an **Attachments** tab (on a requirement, a list under its evidence) with the old files and no upload button. Evidence moved this way that supports a recorded test result locks the night after (see Locked evidence and file fingerprints).

13.4.2 A third party's files

A third party keeps its **Attachments** tab for contracts, DPAs and other documents. A report that was attached there as a plain file can become evidence: **Convert to evidence** beside the file opens the New

evidence form as a report about that third party, titled with the file name and dated the day it was uploaded. Set the type, period, opinion and the rest, and save: the file moves onto the new report (with its fingerprint) and leaves the Attachments tab. Converting needs `evidence:create` and `thirdParties:update` on the third party.

13.5 Replacing earlier evidence

Evidence can replace earlier evidence: this year's SOC 2 report replaces last year's, this year's FISMA assessment replaces last year's. The earlier one then shows as **Superseded**, sends no more reminders, and drops out of the **Evidence expiring** widget and the requirement counts on the Compliance workspace. Each piece of evidence replaces at most one other, and is replaced by at most one.

Set it with **Replaces** on the evidence form. It offers evidence of the same kind that nothing else replaces yet (for a third-party report, that third party's reports), the same type first. On new evidence the form suggests one:

Evidence	Suggested
A third-party report	The newest report of the same type from the same third party.
An Audit or assessment report, or a Certification	The newest of the same type linked to the same control (when you add it from a control or its test).
Anything else	Nothing. Access reviews, log reviews and the like are a history where each entry still counts, so a new one replaces nothing unless you choose.

Clear the field to replace nothing. Changing **Replaces** later is recorded in the audit log; it stays editable on locked evidence, since it says nothing about the evidence's own contents. Reports recorded before Replaces existed were linked by the old rule: each report replaces the previous one of the same type from the same third party.

13.6 Locked evidence and file fingerprints

Every uploaded file gets a SHA-256 fingerprint, taken from its bytes when it is uploaded. The **Files** tab shows the start of it under each file name; the copy button beside it copies the whole value, for an auditor to compare with their own copy of the file. Files uploaded before fingerprinting was added are fingerprinted in the background and show "SHA-256 not recorded yet" until then.

Verify files, below the files, reads each stored file again and compares it with its fingerprint. Each file is reported as **Matches**, **Changed since upload** or **File missing** (a file with no fingerprint yet gets one, shown as **Fingerprint recorded**). Anyone who can read the evidence can run it, and the result is

recorded in the audit log. A changed or missing file means the stored copy is no longer the one that was uploaded: tell your administrator.

Locked evidence is evidence someone has relied on. While it is locked:

- its files cannot be added, removed or replaced;
- it cannot be deleted;
- its type, third party, issuer, period, collected (or report) date, valid until date, opinion and exceptions cannot change;
- a link to a control test with a recorded result cannot be removed.

The title, description and owner can still change, and more records can still be linked. A **Locked** badge shows beside the status, with who locked it, when and why; the evidence list and the Evidence cards on other records show a lock icon.

Evidence locks in three ways:

How	When
By hand	**Lock** on the evidence, with an optional note (for example "Given to the external auditor"). Anyone who can edit the evidence can lock it.
A result is recorded	Recording a control test's result locks the evidence linked to that test, if it has files. Linking evidence that has files to a test whose result is already recorded locks it too.
A day after its files	Evidence added from a recorded result has no files yet, so it stays open while you upload them. Once it supports a recorded result and its newest file is a day old, the nightly run locks it.

Unlock needs `evidence:delete` at the `all` level (GRC Manager, or an administrator) and a reason, which is recorded in the audit log.

13.7 Expiry and reminders

Evidence with a **Valid until** date reminds its owner by email when it comes within 30 days of that date (**Evidence expiring**) and again when the date passes (**Evidence expired**), once each per date: entering a new date starts again. A third-party report also reminds the third party's owner in your organization (the **Owner** on the third party, who owns the relationship and usually has to chase the new report); when that is the same person, they get one email. Each recipient must be able to read the evidence. Nobody at the third party is emailed. Superseded evidence, and a report about a terminated third party, sends no reminder. The **Evidence expiring** dashboard widget lists what is expired and expiring.

13.8 Evidence packages for auditors

An evidence package is one zip an auditor can work from offline: the selected evidence's files, one folder per evidence record, with a manifest.

- **From a record: Download package** on the Evidence tab of a control or requirement packs everything that tab shows (for a requirement, the evidence of its mapped controls and their tests too). On a third party's **Assurance reports** tab and an account review's Evidence tab it packs the evidence listed there.
- **By period: Download package** on the Evidence page asks for **From** and **To** dates and packs the evidence you can read whose period covered overlaps them, or, for evidence with no period, whose collection date falls between them. **Include superseded reports** keeps reports a later one replaces; they are left out otherwise.

The zip holds:

File	What it is
`manifest.json`	Every evidence record: type, third party, issuer, period, dates, opinion, status, lock state (who locked it, when and why), the evidence it replaces or that replaces it, the records it supports, and each file.
`manifest.csv`	The same, one row per file, for a spreadsheet.
`README.txt`	What the columns mean and how to check a hash yourself.
`evidence/`	The files, in a folder per evidence record.

Each file is listed with two SHA-256 values: the fingerprint taken when it was uploaded, and the hash of the bytes in the package. **file_check** says `match` when they agree, `mismatch` when the stored file changed, `missing` when it is gone from storage, and `recorded` for an older file that had no fingerprint yet. Mismatched and missing files are also listed together under **problems**.

A package contains only what you can read: evidence you cannot read is left out, and records the evidence supports that you cannot read are counted, not named. It holds at most 500 evidence records. Downloading a package needs `evidence:read` and is recorded in the audit log with what it held.

13.9 Who can do what with evidence

Action	Permission
See the Evidence page and evidence on records	`evidence:read` (at the `own` level, evidence owned by you or your groups)
Add evidence	`evidence:create`
Edit, link, unlink, upload files, comment	`evidence:update` (at the `own` level, on evidence you own)
Delete evidence	`evidence:delete` (unlocked evidence only)
Lock evidence	`evidence:update` on that evidence
Unlock evidence	`evidence:delete` at the `all` level
Verify files	`evidence:read`
Download an evidence package	`evidence:read`

The built-in GRC Manager role has every evidence right, Risk Owner can read all evidence, add it and edit their own, and Auditor can read it. Administrators can do everything.

CHAPTER 14

Online Assessments

Online assessments let you send a questionnaire to someone outside your organization, typically a vendor or other third party, without giving them an account. You build reusable templates, create a request from a template, and the contact answers through a personal link. Choice answers are scored automatically, files can be attached as evidence, and the answers come back to you for review answer by answer. A request can repeat every year (or every N months) against the same third party, offering last time's answers for the contact to confirm.

The module is used mainly by vendor and third-party risk managers, with security and privacy staff reviewing answers in their area. This chapter covers both sides: managing templates and requests inside TraxrGRC, and what your contact sees on the public assessment page.

14.1 The Assessments page

Open **Assessments** from the navigation. The page has two tabs:

- **Templates** (the default): the reusable questionnaires, with the columns **Title**, **Questions** (the number of questions), **Requests** (how many requests use the template), **Owner** and **Created**.
- **Requests**: every questionnaire sent to a contact (see Requests below).

Both lists have a search box, **Export CSV** and **Saved** filters. With `assessments:create` you see **New template** at the top of the page and **New request** on the **Requests** tab.

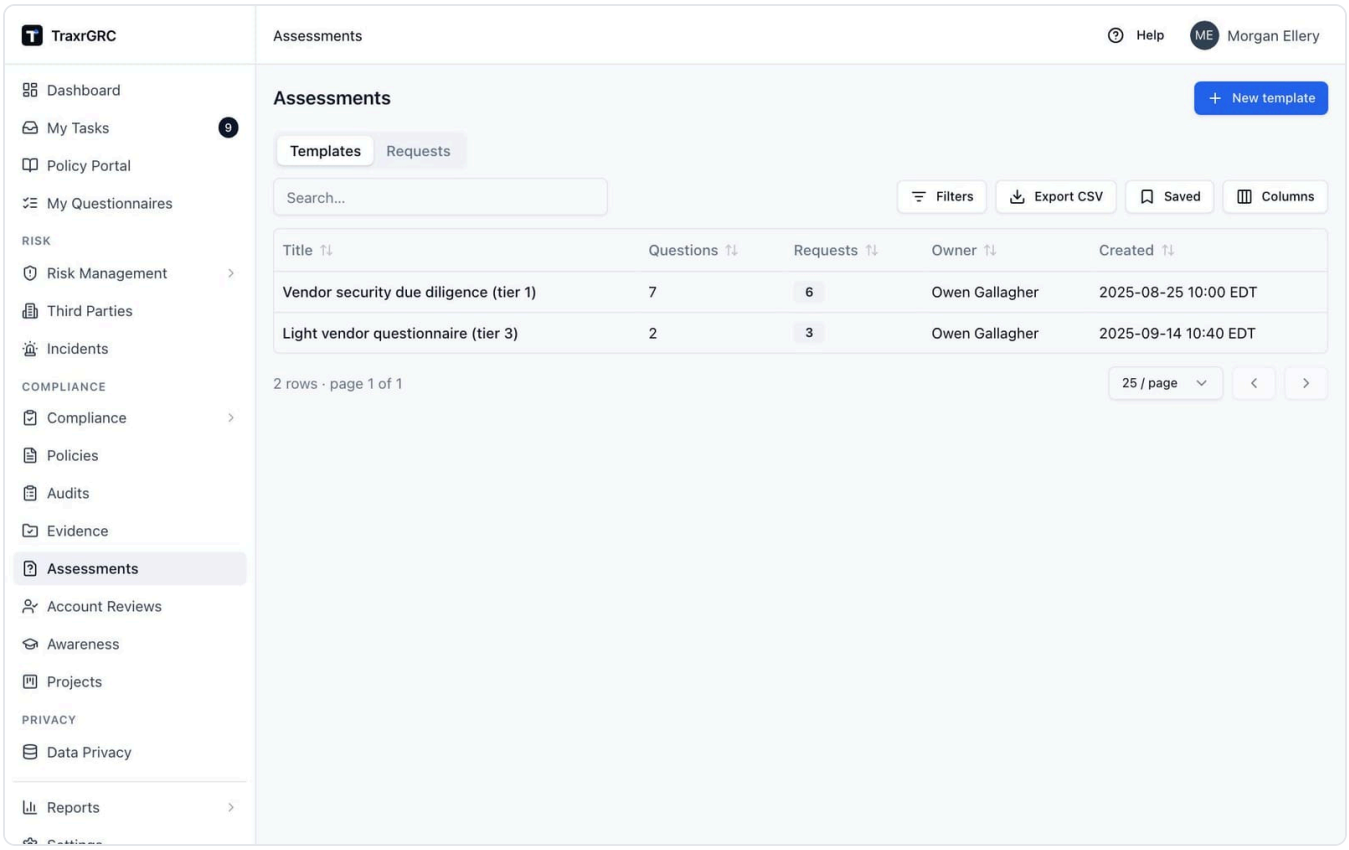


Figure 14.1. The Assessments page on the Templates tab, listing the tier 1 due diligence and tier 3 light questionnaires with their question and request counts.

The dashboard widget **Assessment requests** counts requests by status (see Getting Started, The Dashboard).

14.2 Templates and the question editor

14.2.1 Creating a template

1. Click **New template**. The **New assessment template** dialog opens.
2. Fill in **Title** (required, for example "Vendor due diligence"), an optional **Description**, and an optional **Owner**. If you leave the owner as "None", you become the owner.
3. Click **Create template**.
4. Open the template from the list to add its questions.

14.2.2 The template page

The template page shows the title, owner, creation time and the number of requests that use it. **Edit details** opens the **Edit template** dialog for the title, description and owner (questions are edited on the page itself); click **Save changes** to apply. Below the header is the question editor.

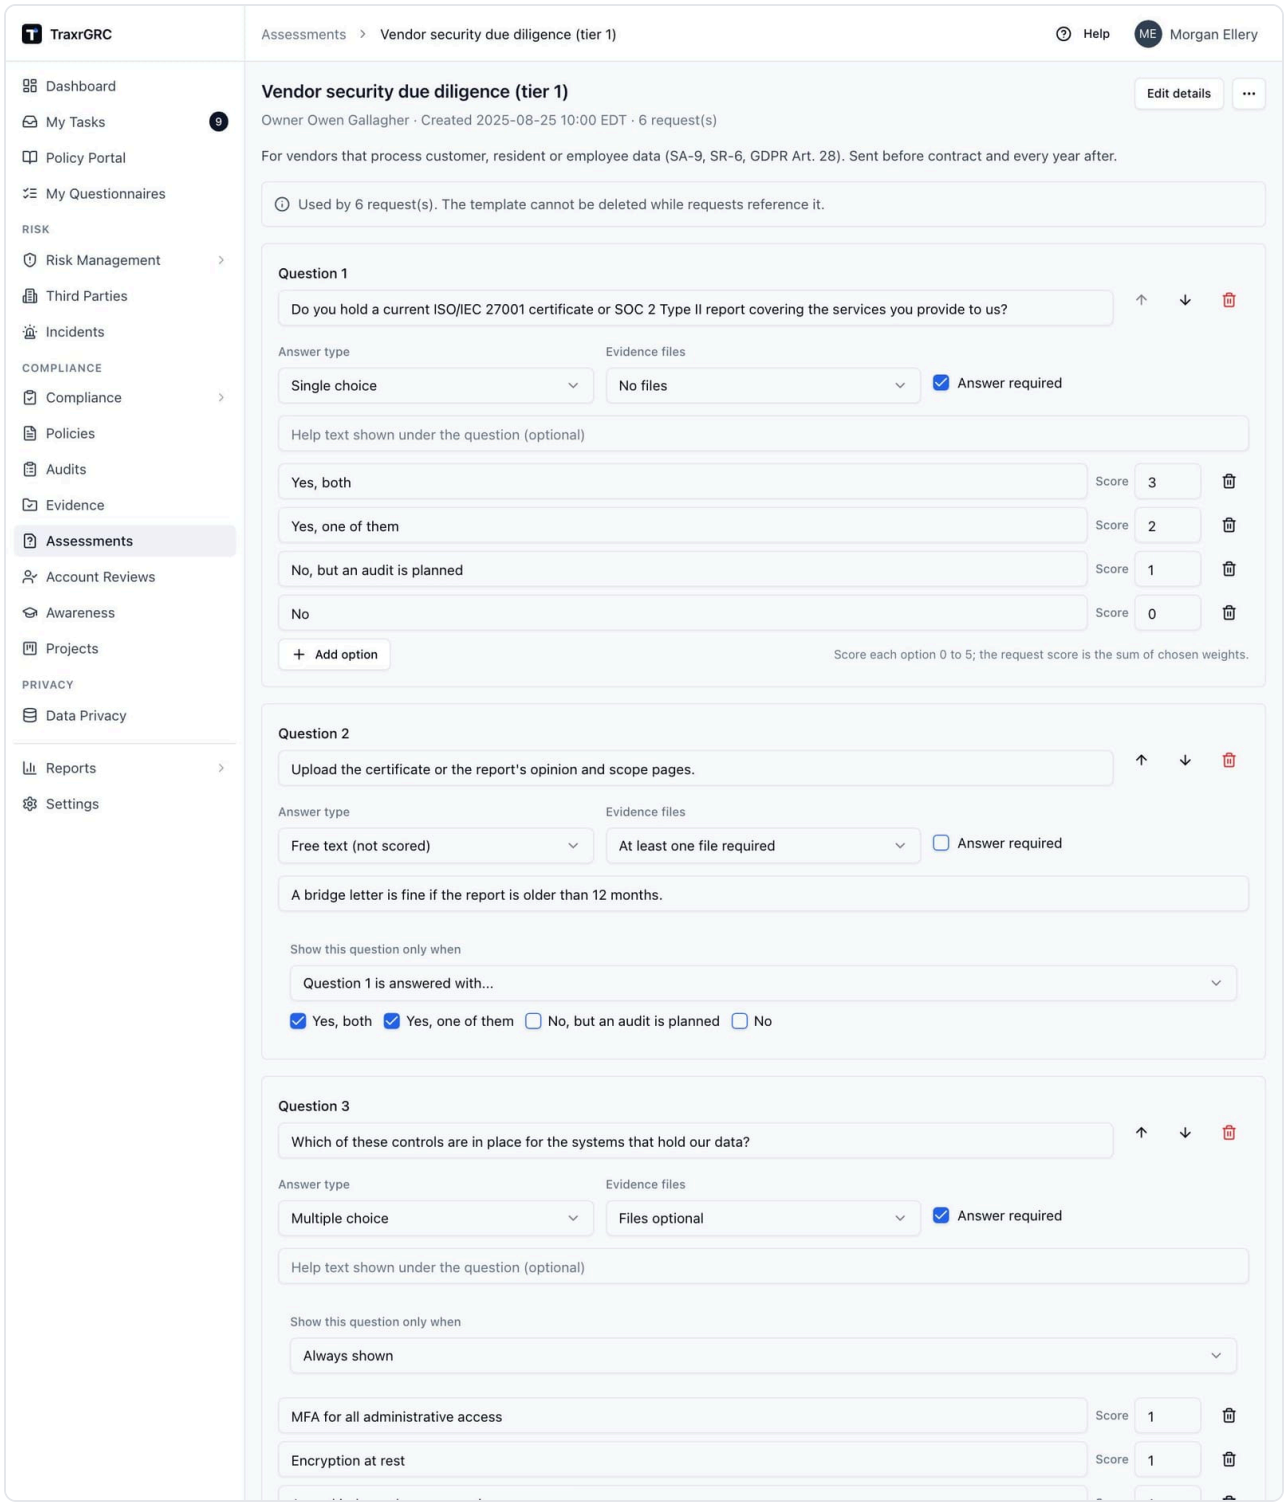


Figure 14.2. The Vendor security due diligence (tier 1) template, showing scored single and multiple choice questions, an evidence question and a follow-up question with its condition.

Each question has these settings:

Setting	What it does
Question text	Required.
Answer type	**Single choice** (one option), **Multiple choice** (any number of options) or **Free text** (not scored).
Options	Choice questions only: 2 to 20, each with a label and a **Score** from 0 to 5. A single choice earns its option's score; a multiple choice earns the sum of the options picked.
Evidence files	**No files**, **Files optional** or **At least one file required**: whether the contact can or must attach documents (certificates, reports, policies) to this question.
Answer required	On by default. An optional question can be left blank.
Help text	Optional guidance shown under the question.
Show this question only when	Optional condition. **Always shown** by default; otherwise pick an earlier single or multiple choice question and tick the options that show this one (one option means "equals", several mean "one of").

To build the questionnaire:

1. Click **Add question** for each question and complete its settings.
2. On choice questions, use **Add option** to add options and the trash button to remove one. The last two options of a choice question cannot be removed.
3. Use the arrow buttons to move a question up or down, and its trash button to delete it.
4. Click **Save questions**. Nothing is stored until you save.

A template holds at most 100 questions. Contacts never see the scores; grading happens inside TraxrGRC. A request's score is the sum earned over the questions the contact was shown, as a percentage of what those questions could earn: a question hidden by its condition counts neither way, and free-text questions are not scored.

Conditions. A question can depend only on an earlier choice question. If you move questions so that a condition would point at a later question, or remove the option it names, the condition is dropped. A question that depends on a hidden question is hidden too. Hidden questions are never required, and anything typed into them before they were hidden is dropped when the contact submits.

Saving questions replaces the whole set. Each request keeps the questions it was sent with, so editing a template never changes, regrades or redisplays assessments already sent or completed. A repeat (see Repeating assessments) asks the template's questions as they are when it is sent. Templates built before answer types existed keep working unchanged: their questions read as required single choice without evidence.

14.2.3 Deleting a template

A template that has requests cannot be deleted; the template page says how many requests use it. Once the last referencing request is gone, **Delete** in the **More actions** menu (**...**, requires `assessments:delete`) asks for confirmation ("Delete template?") and removes it permanently.

14.3 Requests

14.3.1 Creating a request

1. On the **Requests** tab, click **New request**. The **New assessment request** dialog opens.
2. Complete the fields:

Field	What it means
Template	Required. Choosing one prefills the title.
Title	Defaults to the template title. Override it for a per-request name.
Third party (optional)	Links the request to a third-party record (see Third-Party Management). Required for a repeating request and for raising risks from answers.
Contact name	Optional. Shown in the public page greeting and in emails.
Contact email	Required. The person you are sending the questionnaire to.
Repeat	**Does not repeat**, **Annually**, **Every 2 years**, **Every 6 months**, **Every 3 months** or **Every month**. A repeating request needs a third party.
Email the link to the contact, with a reminder a week before the due date	Off by default. When ticked, TraxrGRC emails the link now, a reminder a week before the due date, and a notice if you send answers back for rework. Always on, and locked, for a repeating request.
Due date (optional)	After this date the link stops accepting answers.

3. Click **Create request**. The **Request created** dialog shows the public link with a **Copy** button.
4. If TraxrGRC is not emailing the contact, send them the link yourself.
5. Click **Done**.

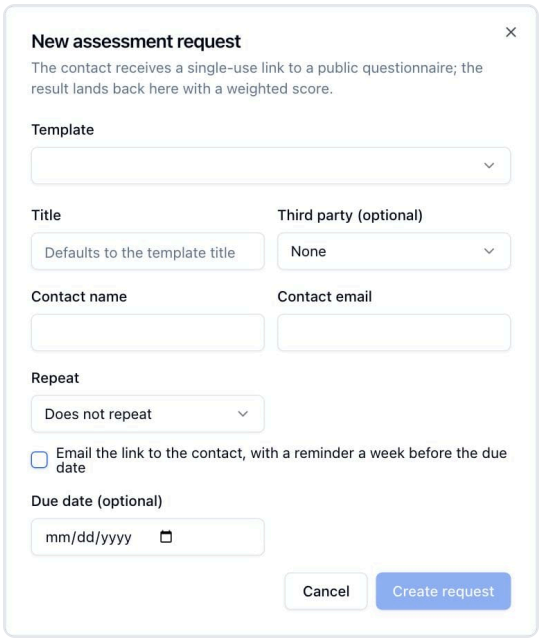


Figure 14.3. The New assessment request dialog with the template, third party, contact, repeat, email and due date fields.

The link contains a random, unguessable token. Anyone holding it can answer, so share it only with the intended contact.

14.3.2 The requests list

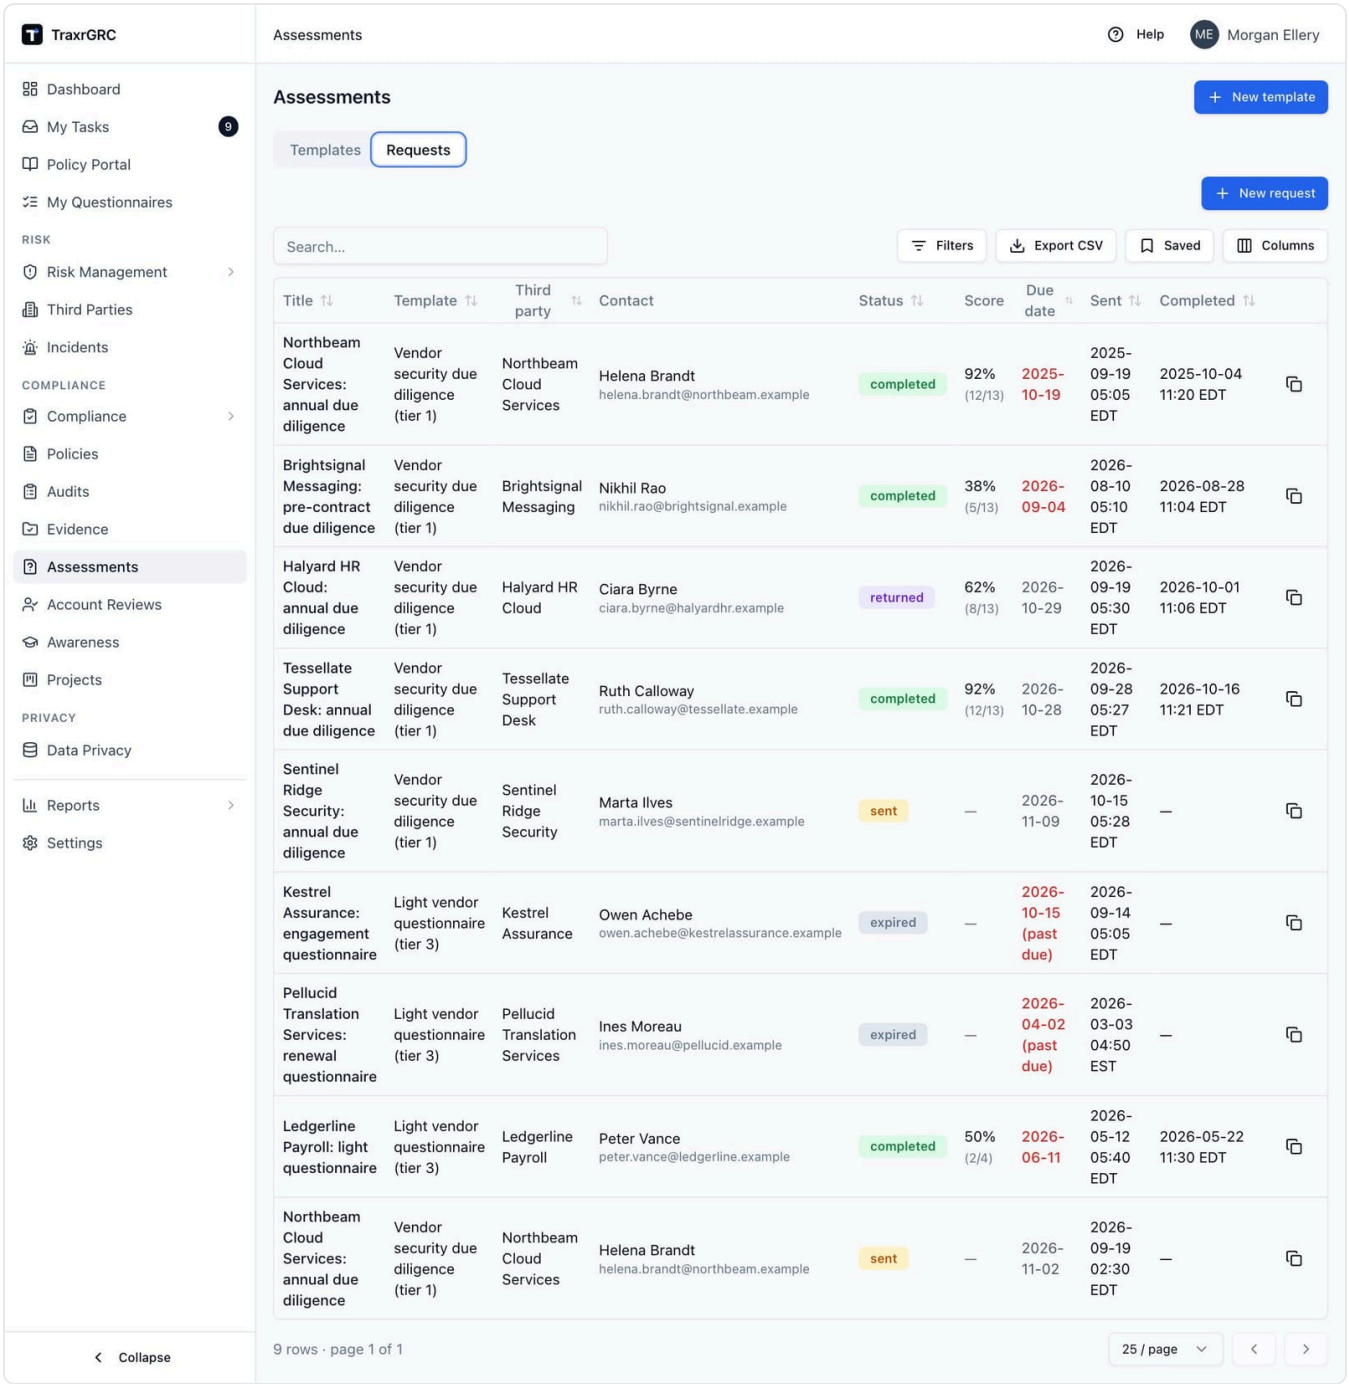


Figure 14.4. The Requests tab, with completed, returned, sent and expired due diligence requests and their scores and due dates.

Column	What it shows
Title	Click to open the request.
Template	The template the request was created from.
Third party	The linked third party, if any.
Contact	The contact's name, with the email underneath.
Status	The status badge (see below). Filterable.
Score	For a submitted request, `pct% (score/max)`.
Due date	The due date. A past date shows in red with "(past due)" unless the request is completed.
Sent	When the request was sent.
Completed	When the contact submitted.

The search box matches the title, contact name and contact email.

Status badges:

Status	Meaning	Badge
Sent	Created; the contact can still answer.	Amber
Completed	The contact submitted.	Green
Returned for rework	You sent flagged answers back; the contact can change those until the new due date.	Violet
Expired	Either explicitly marked expired, or a *Sent* or *Returned* request past its due date. The badge changes as soon as the due date passes, without anyone touching the record.	Gray

The **Status** filter matches the status stored on the request. A request that shows **Expired** only because its due date has passed is still stored as *Sent* or *Returned*, so the filter lists it under that status.

Row actions on each request:

- **Copy link** (copy icon, only with `assessments:update` on the request) copies the public link again; the icon turns into a checkmark when copied.
- **Mark expired** (ban icon, on requests stored as *Sent*) asks for confirmation and then makes the link stop accepting answers immediately.

- **Delete** (trash icon) asks for confirmation and permanently removes a request that has never been submitted, together with its link and any files the contact uploaded. A submitted request cannot be deleted: its answers are the record.

14.3.3 The request page

Open a request from the list. The page has no tabs; a **Custom Fields** card below the answers shows any custom fields defined for assessment requests.

The header shows the title, status badge, template, third party, sent time and owner. Below it are the contact, the due date (red once expired), the completion time and, for people with `assessments:update` on the request, the **Public link** with a **Copy** button. The link is hidden from everyone else because it is the contact's only credential.

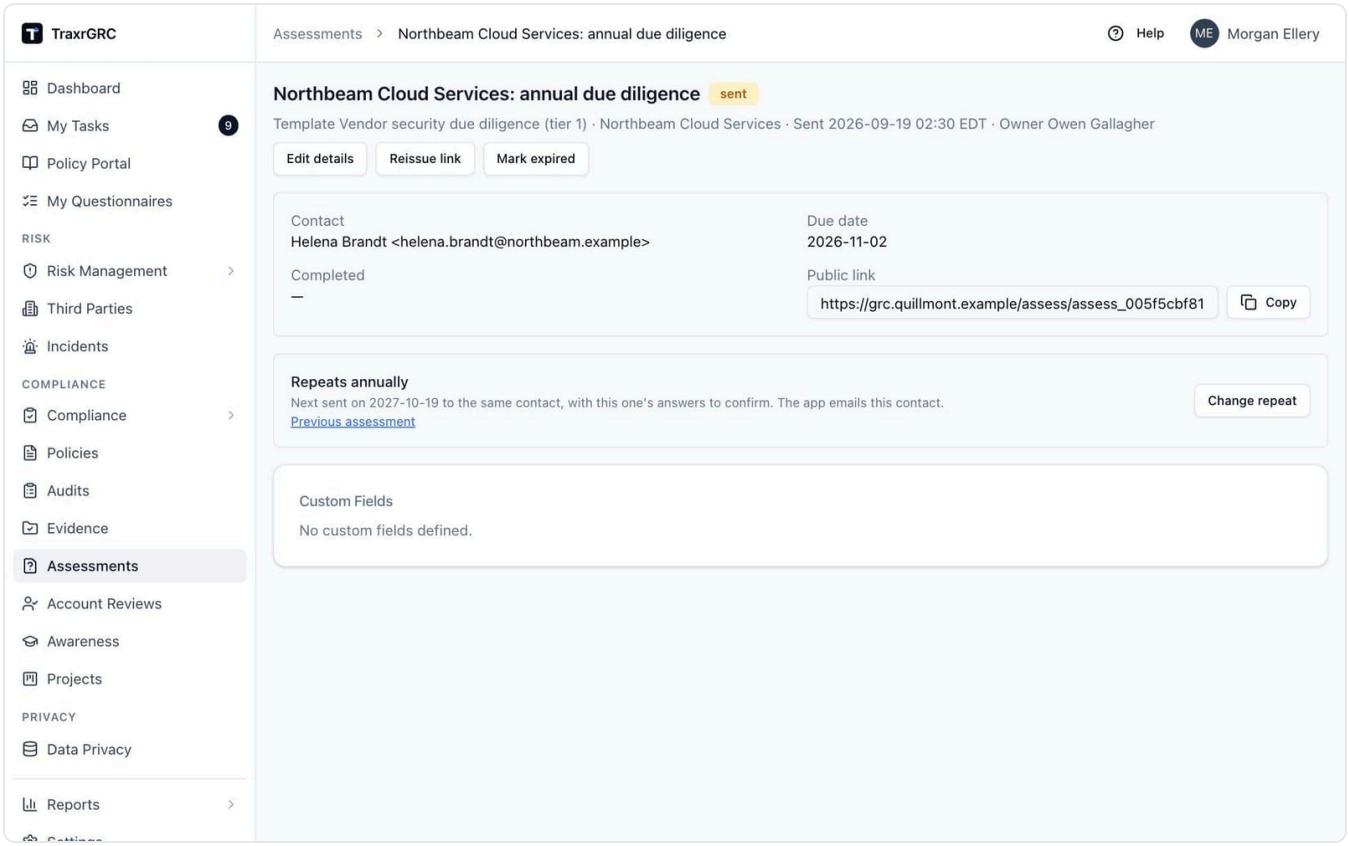


Figure 14.5. The current year's Northbeam Cloud Services due diligence request, sent and awaiting answers, with its public link and repeat panel.

The buttons at the top right depend on the request's state and your rights (`assessments:update`):

Button	When it shows	What it does
PDF	Once the contact has submitted. Needs only read access.	Downloads the submitted assessment (see Reviewing answers).
Edit details	Until the contact first submits.	Opens **Edit request** for the title, contact name, contact email, due date and third party. The public link and any answers are not affected.
Reissue link	On *Sent* and *Returned* requests.	Replaces the public link (see below).
Mark expired	On *Sent* and *Returned* requests.	After confirmation, the link stops accepting answers immediately.

Once the contact has submitted, a score summary appears with the percentage and the raw score, followed by the **Answers review**.

14.3.4 Reissuing the public link

Use **Reissue link** when the link reached someone other than the contact.

1. On the request page, click **Reissue link**. The **Reissue the public link?** dialog opens.
2. Leave **Email the new link to** *contact* ticked to have TraxrGRC send the new link, or untick it to send the link yourself. The box starts ticked when the request was set up to email the contact.
3. Click **Reissue link**.

The old link stops working at once. Answers and evidence already given stay on the request, and the new link appears in the **Public link** field. The new link is never written to the audit log; the log records only that the link was reissued and to whom it was emailed. If someone else already submitted through a leaked link, return the request for rework first (see Reviewing answers), then reissue the link.

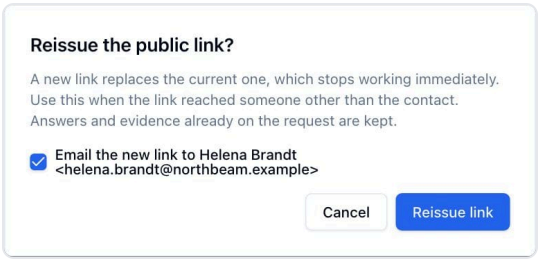


Figure 14.6. The Reissue the public link dialog, with the option to email the new link to the contact.

14.4 Reviewing answers

The **Answers review** lists each answer with:

- its points (for scored questions);
- the files attached to it, by name and the start of its SHA-256 fingerprint (click a name to download; files added during rework are marked "(rework)");
- "Unchanged from last time" where the contact confirmed last time's answer;
- a decision badge: **Not reviewed**, **Accepted**, **Flagged** or **Reworked, review again**;
- the review comment, and a **History (N)** list once an answer has more than one decision.

The status line above the answers reads **Review in progress**, **Flagged answers are back with the respondent** or **Review finished**, followed by how many answers are accepted, flagged and still to decide, and who finished the review and when. Below the answers, each submission (the first and any rework) is listed with its time and score.

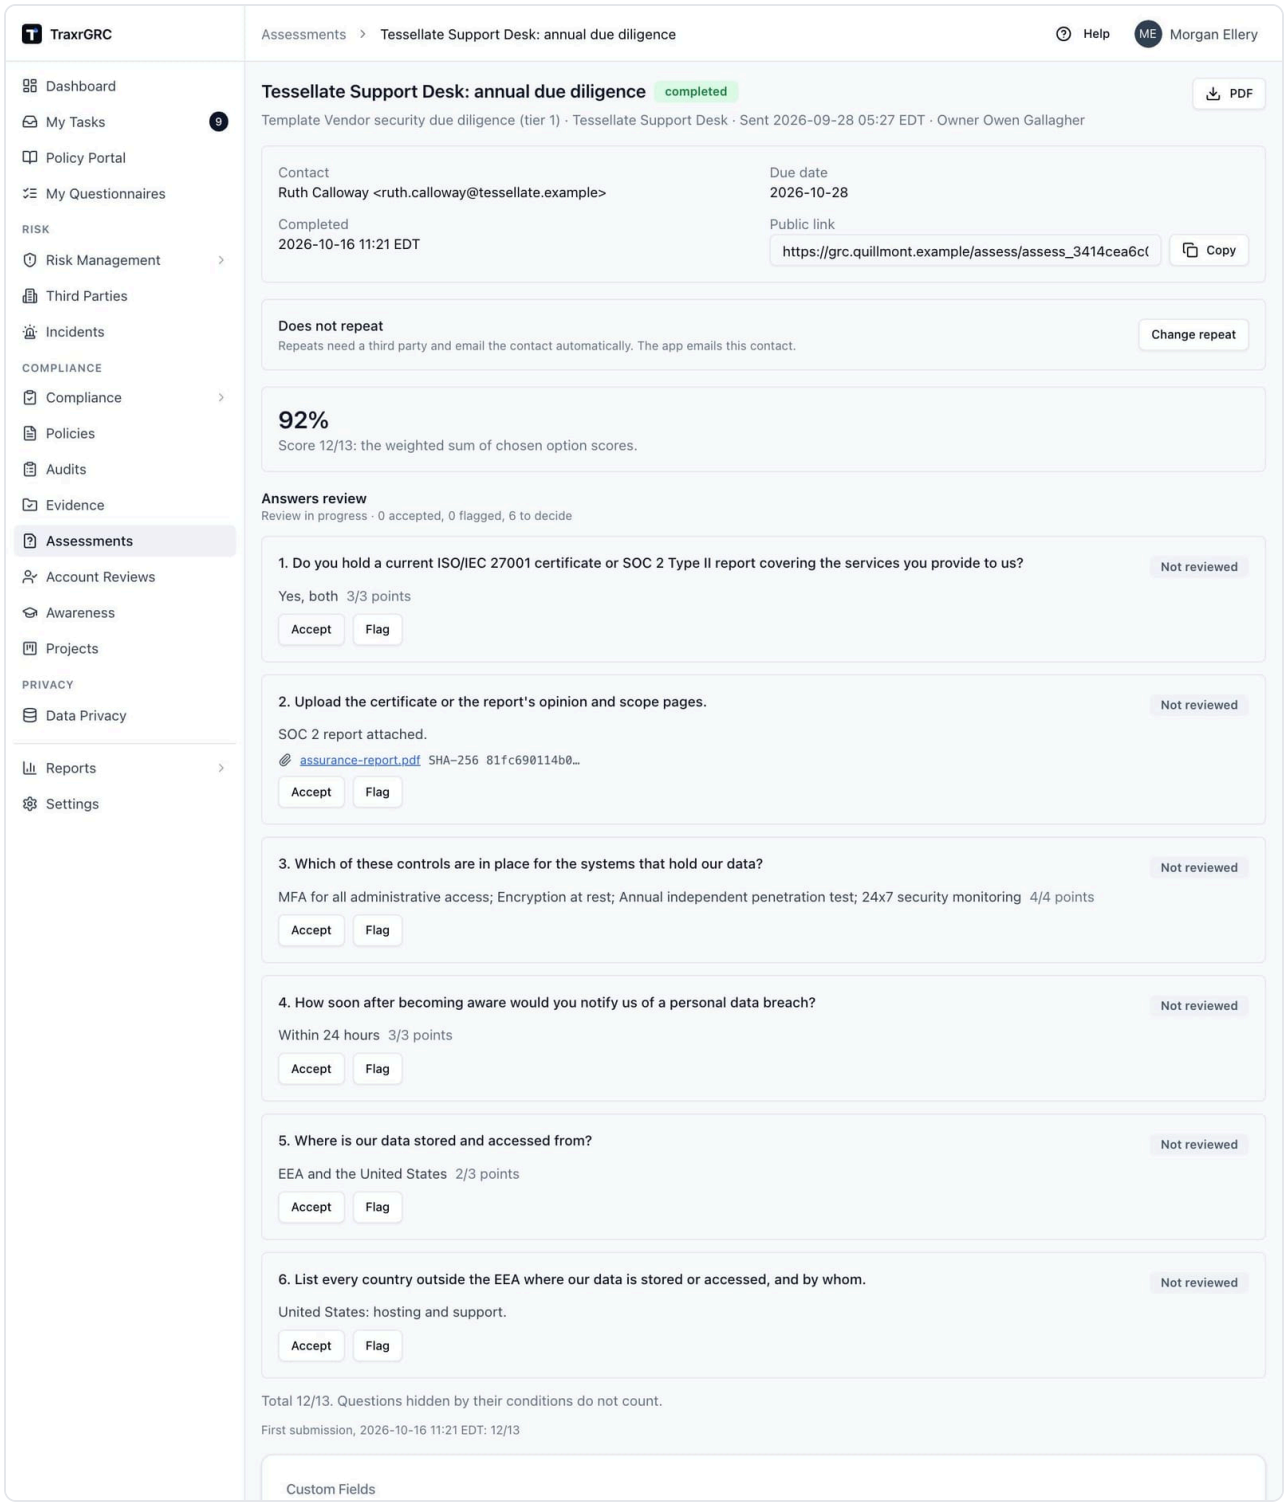


Figure 14.7. The Tessellate Support Desk due diligence answers awaiting review, each with Accept and Flag buttons and the attached assurance report file.

With `assessments:update` on the request, you review as follows:

1. Click **Accept** on each answer you are satisfied with, or click **Flag**, enter what needs to change (a comment is required; the contact sees it if the answers go back) and click **Flag answer**. You can change a decision until you finish; every decision is kept in the answer's **History**.

2. Optionally, on a flagged answer of a request linked to a third party, click **Raise third-party risk**. In the **Raise a third-party risk** dialog, adjust the **Risk title** and click **Raise risk**. TraxrGRC creates a risk in the **Third-party** classification, attached to that third party and owned by you, at the lowest likelihood and impact for you to rate. Its description carries the question, the answer and your comment, and the answer links to it ("Risk raised: ..."). One risk per answer. This needs `risks:create` and read access to the third party.
3. Optionally, click **Send flagged answers back** (once per request). In the dialog, set a **New due date** (14 days from today if left empty) and click **Send back**. The request becomes **Returned for rework**. The contact sees your comments and can change only the flagged answers, add files to them, and answer any question their changes reveal; everything else stays as it was. If TraxrGRC does not email this contact, the dialog reminds you to tell them the link is open again. While the answers are back with the contact, the flagged answers keep their **Flagged** badge and your comment, and the status line reads **Flagged answers are back with the respondent**.
4. When the contact resubmits, each reworked answer shows **Reworked, review again** until you decide on it again. Flags go stale only at this point; if the link expires during rework without a resubmission, the flags stay current.
5. Click **Finish review** once every answer has a current decision. Flagged answers can stay flagged when you finish; the flag is the decision.

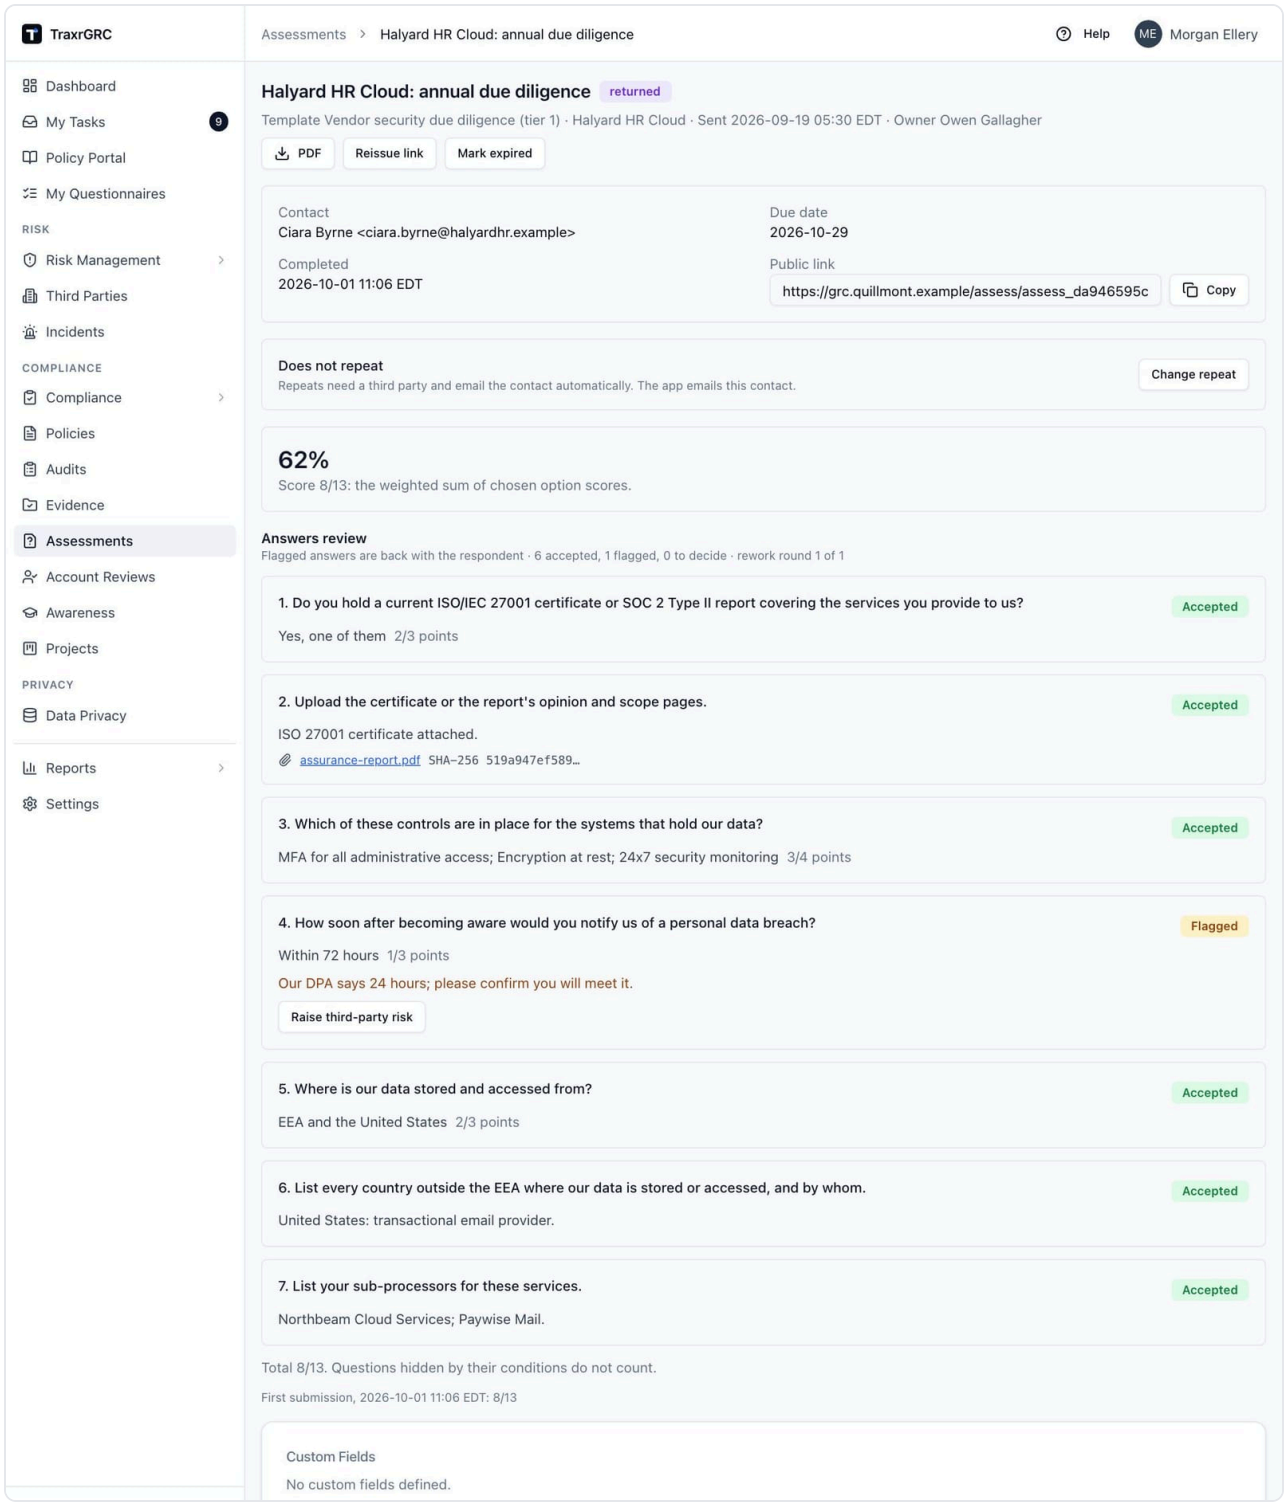


Figure 14.8. The Halyard HR Cloud due diligence request returned for rework, with the reviewer's comment on the breach notification answer and the review status showing the answers are back with the respondent.

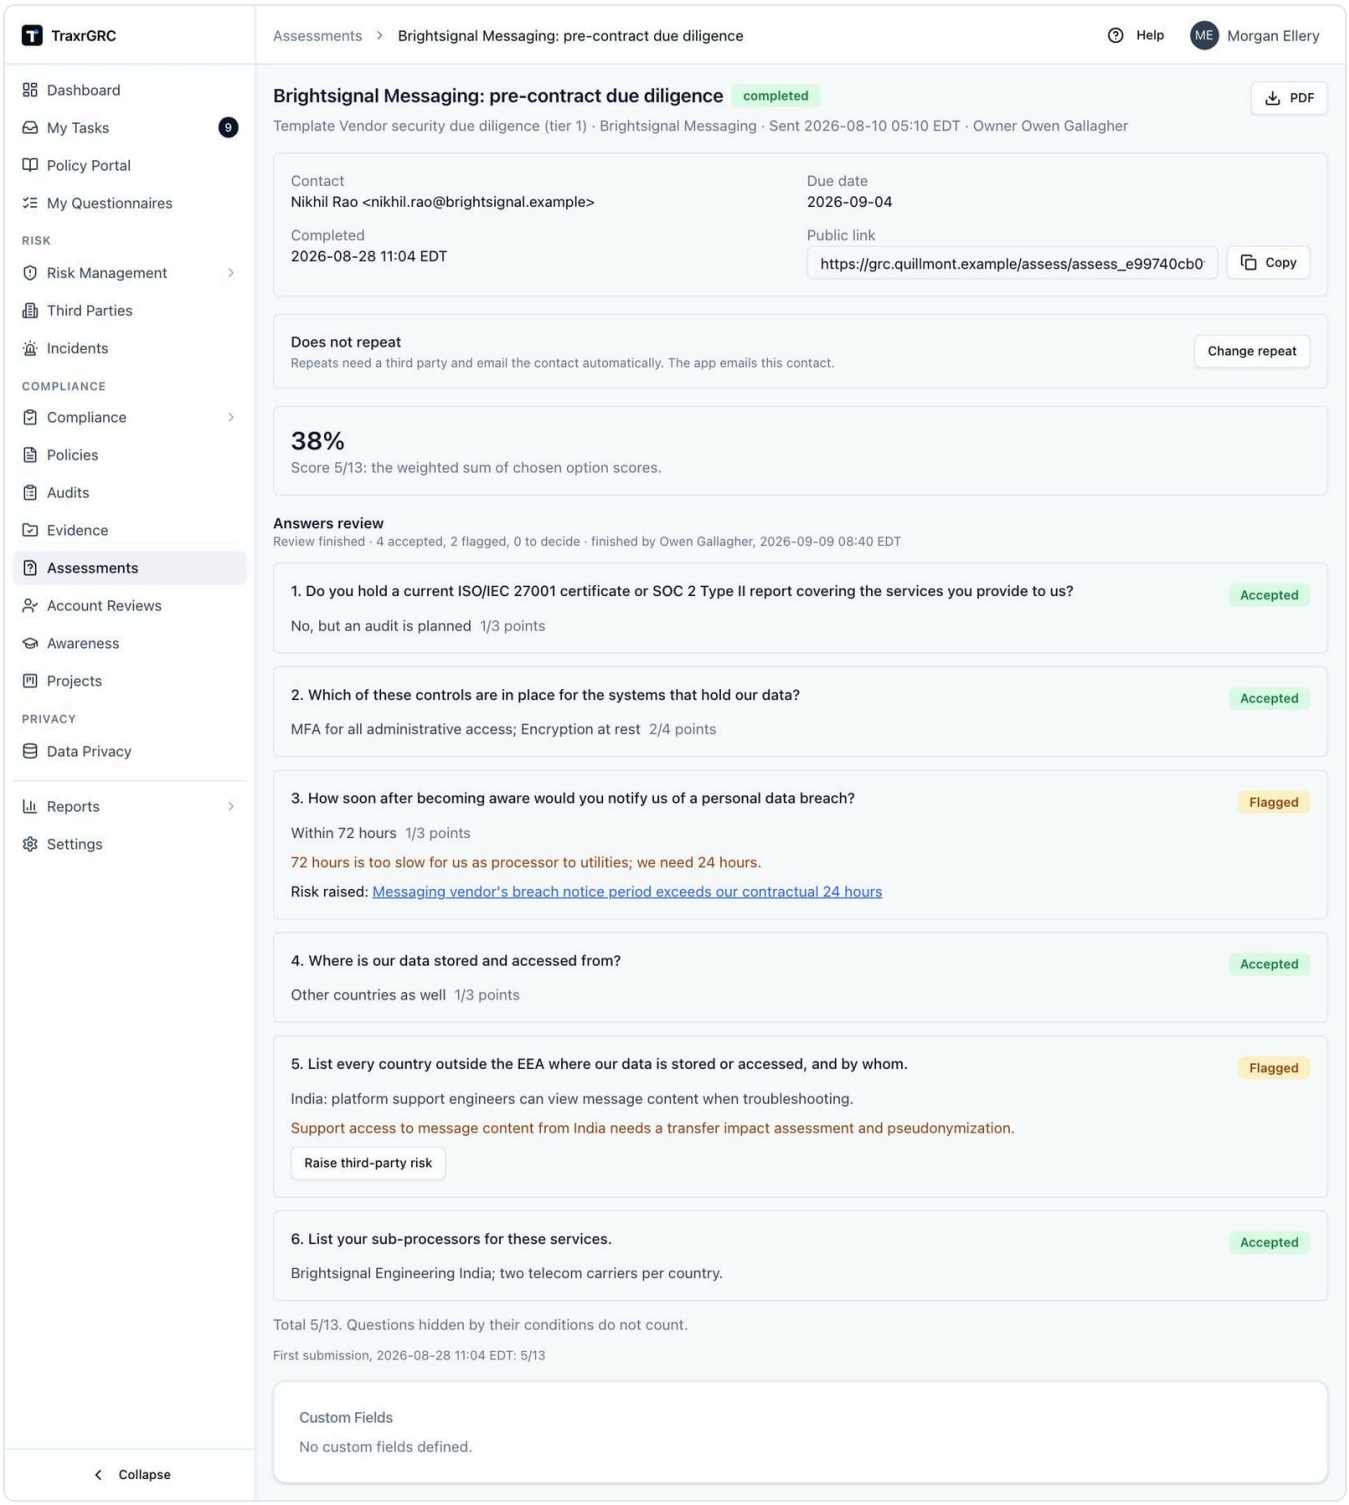


Figure 14.9. The finished review of the Brightsignal Messaging due diligence, with two flagged answers, their comments, and the third-party risk raised from one of them.

PDF. The **PDF** button downloads the submitted assessment: every question in order with the answer, points, whether it was confirmed unchanged, the review decision with the reviewer, time and comment, earlier decisions, any risk raised (by title if you may read it), and each evidence file by name, size and SHA-256 fingerprint. The files themselves are not embedded; download them from the request page.

Questions the contact was not shown are marked as such. Times are in your time zone. At most two PDFs render at once; if the server is busy, try again in a few seconds.

14.5 Repeating assessments

You set a repeat in the **New assessment request** dialog, or with **Change repeat** on the request page. Only the latest request in a chain can carry the repeat, so **Change repeat** appears only there. The repeat panel on the request page shows the schedule ("Repeats annually" and the date the next one will be sent) and, in a chain, **Previous assessment** and **Next assessment** links.

To change or stop a repeat:

1. On the latest request in the chain, click **Change repeat**. The **Repeat this assessment** dialog opens.
2. Choose an interval, or **Does not repeat** to stop.
3. Click **Save**.

The daily notification sweep sends the repeat automatically, counted from when the request was sent: annually means a year after the send date. Setting a repeat on an older request whose date has already passed sends the repeat at the next sweep. The repeat:

- goes to the same contact and third party, owned by the same person;
- asks the template's current questions;
- gets as many days to answer as the previous request had (30 if it had no due date);
- is emailed to the contact, who sees last time's answers under each question with **Unchanged from last time** to confirm them;
- carries the schedule on, so the previous request links to the next one and vice versa.

14.6 The public assessment page (for external contacts)

This page is for the people you send links to; they need **no TraxrGRC account**. The URL ends in the link's token. The page carries only a header with your logo, product name and "Online assessment": no navigation and no sign-in.

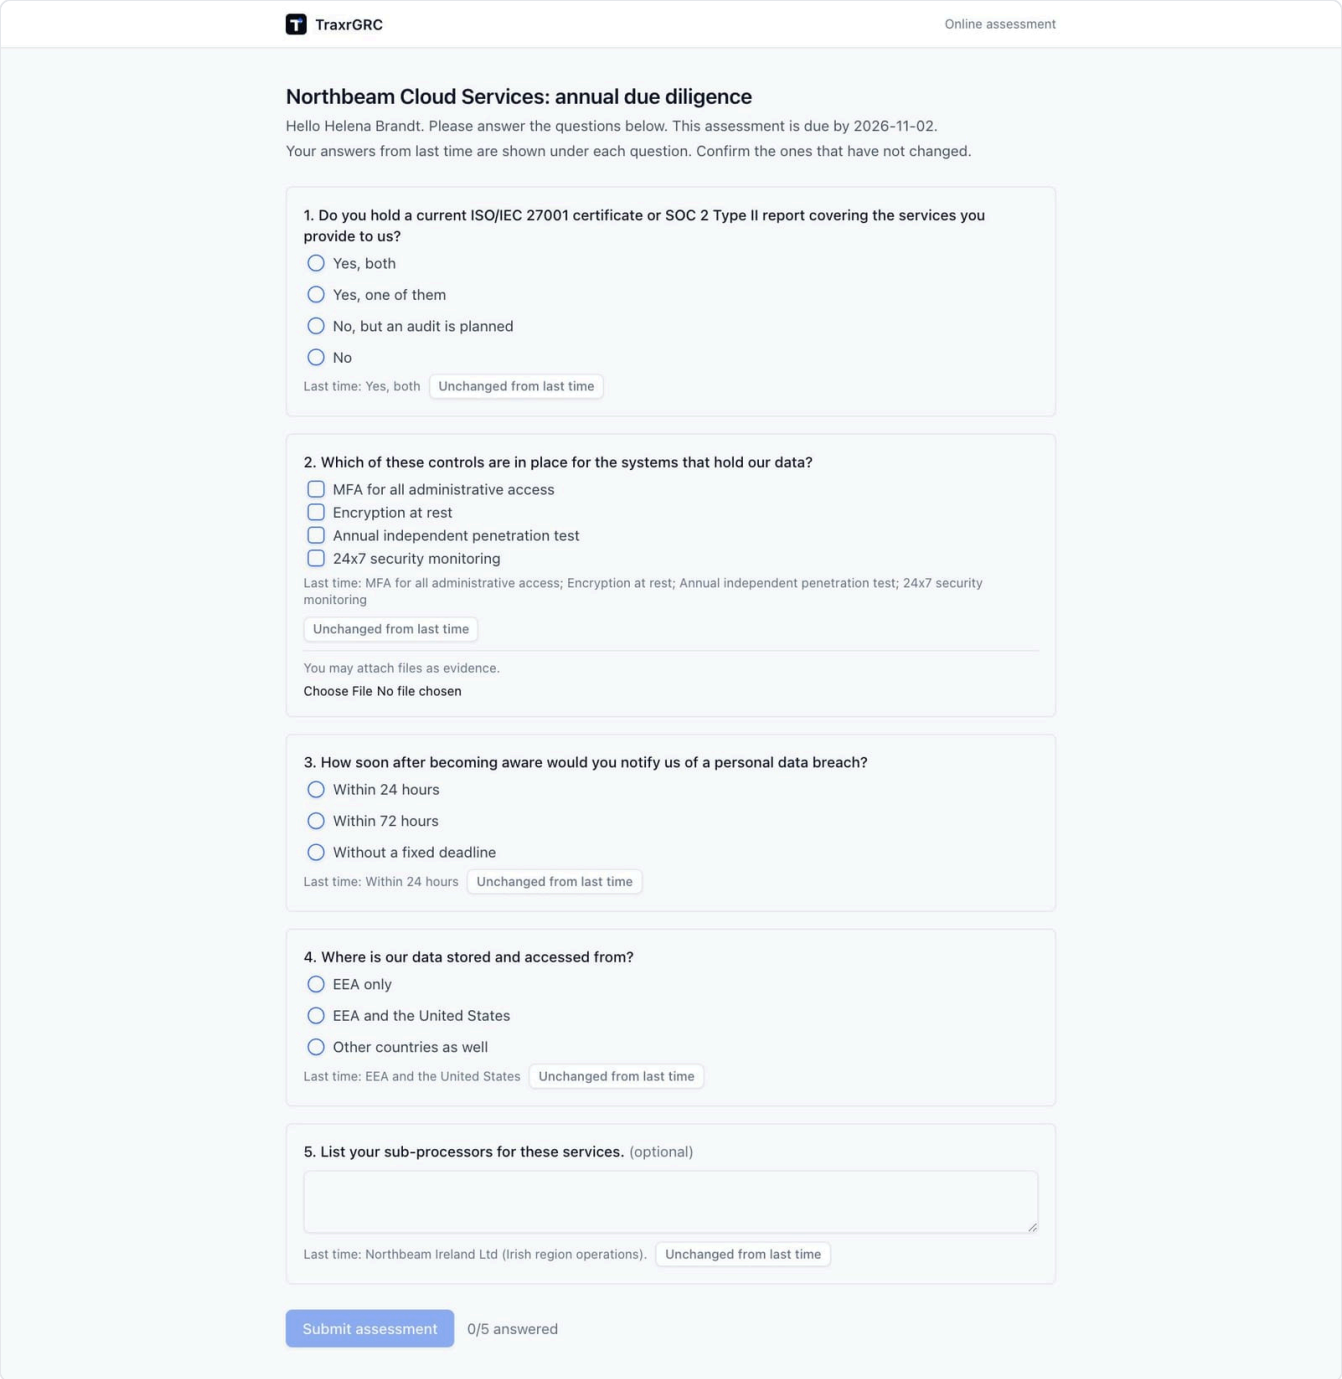


Figure 14.10. The public assessment page for the Northbeam Cloud Services repeat, greeting the contact and offering last time's answers to confirm.

- **Answering.** The page greets the contact by name (if given), says "Please answer the questions below." and shows the due date. Single choice questions offer radio buttons, multiple choice questions checkboxes, and free-text questions a text box. Optional questions are marked "(optional)", and help text appears under the question. Follow-up questions appear and disappear as the answers they depend on change. A counter next to the button shows how many shown questions are answered. **Submit assessment** stays disabled until every shown required question

has an answer and every shown question that needs a file has one. Option scores are never shown.

- **Evidence files.** Where a question takes files, the page says "Attach at least one file as evidence." or "You may attach files as evidence." The contact picks a file and it uploads straight away (PDF, images, Office documents, text, CSV or ZIP; the content must match the extension). Limits per link: 10 MB per file, 5 files per question, 40 files and 100 MB in all. Files can be removed until the answers are submitted; after that they are kept as evidence.
- **Last time's answers.** On a repeat, each question shows "Last time:" with the earlier answer and an **Unchanged from last time** button. Clicking it fills the answer in and shows "Confirmed unchanged". Changing the answer removes the mark.
- **Rework.** When answers are sent back, the page explains that the reviewer has asked the contact to look again at the answers marked below, and shows the reviewer's comment ("Reviewer: ...") on each flagged answer. Only those answers can be changed, along with any question the changes reveal. **Send revised answers** resubmits.
- **Thank-you screen.** On submit, the page shows "Thank you!" and the contact's score as a percentage. Reopening the link shows "This assessment has been submitted" with the score and completion date.
- **Expired link.** Past the due date, or after someone marks the request expired, the page shows "This assessment link has expired", states the due date, and asks the contact to request a new link from the person who sent it. Submissions and uploads are refused as well.
- **Unknown links and rate limits.** An unknown or mistyped link shows a generic "not found" page. The page and its uploads are rate-limited per address; an address that makes too many requests, or tries too many unknown links, sees "Too many requests" and is asked to wait a few minutes.

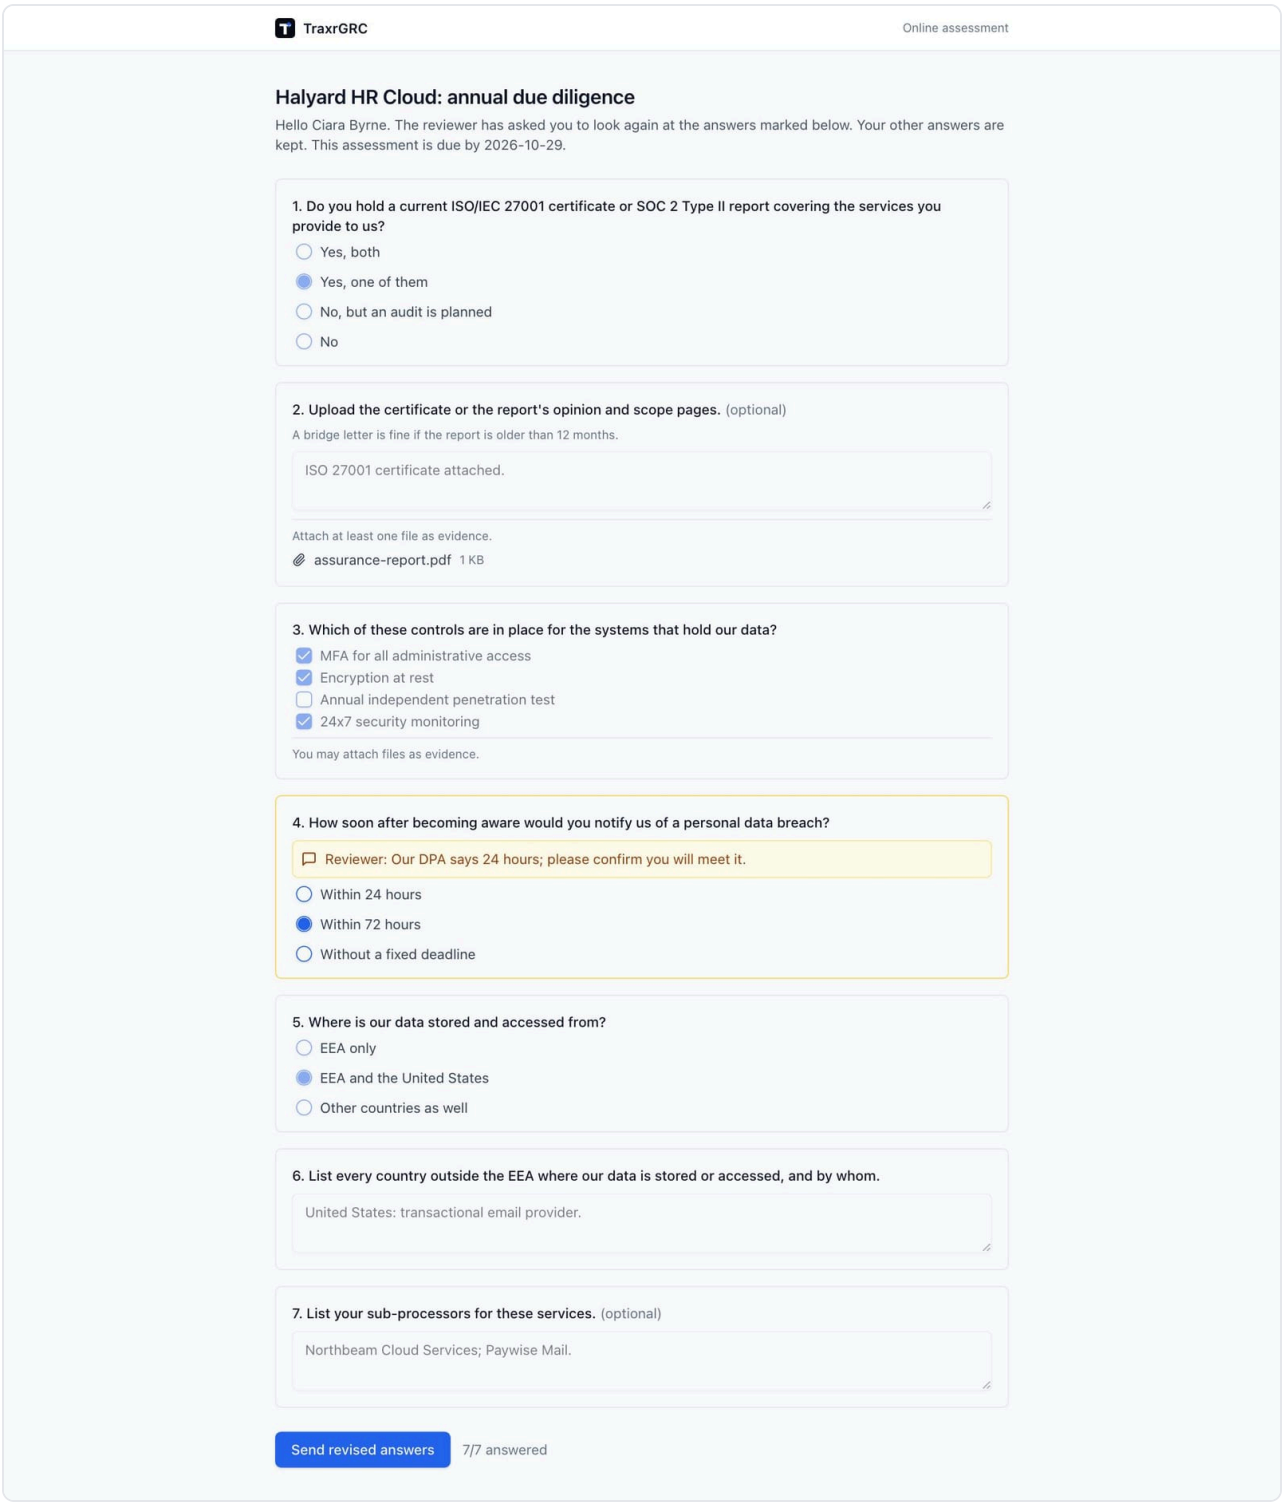


Figure 14.11. The public page for the returned Halyard HR Cloud assessment, with the reviewer's comment on the flagged breach notification answer.

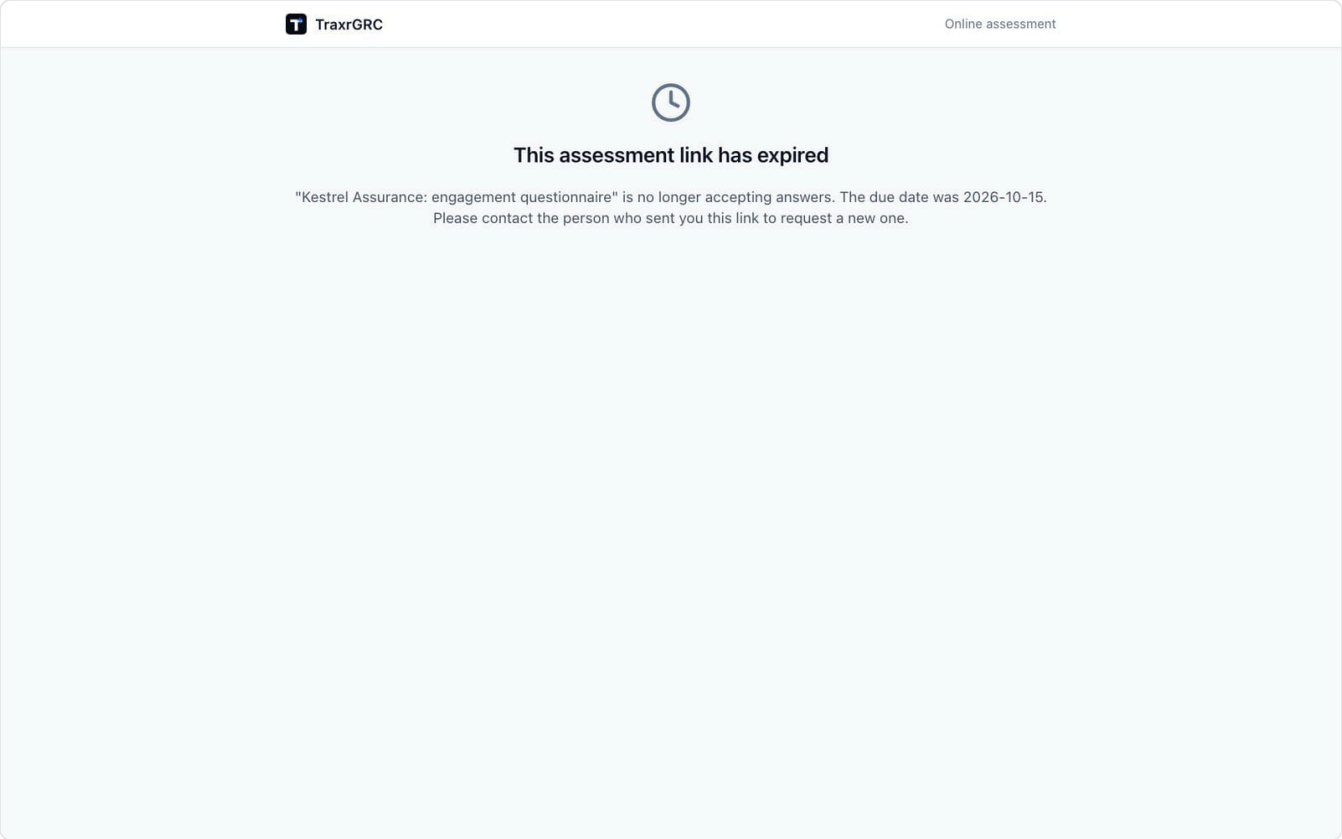


Figure 14.12. The public page for the Kestrel Assurance questionnaire after its due date, showing that the link has expired.

Evidence uploads and upload links. Evidence on the public assessment page is uploaded by the external contact and stays attached to the answer. The separate **Upload files** page (`/upload/...`) belongs to the AI assistant feature: it attaches files to a record for a signed-in user whose assistant issued the link, and it is never sent to external contacts. See Administration & Settings, AI assistants.

14.7 Notifications and audit trail

When a contact submits (or resubmits after rework), the request's owner receives an **email notification** with the title, contact and score. Where TraxrGRC emails the contact, they receive the link (`assessment.sent`), a reminder a week before the due date (`assessment.reminder`, once per due date, never within a day of sending) and the rework notice (`assessment.returned`). A reissued link is sent with the `assessment.sent` template. If a request's due date passes unanswered, its owner is told once (`assessment.overdue`), within a week of the date. The templates are under Settings, Notifications, Templates. Emails to the contact carry the link, so the delivery log (visible to notification administrators) shows it too.

The audit log records:

- the submission, and each upload or removal of a file, with the contact's email as the actor (the submitter holds no TraxrGRC account);

- every review decision, return, finished review and raised risk;
- each link reissue (without the new link);
- repeats sent by the sweep (as "system");
- all template and request changes.

14.8 Who can do what

Action	Permission
See the Assessments page, templates and requests, and download the PDF	`assessments:read` (requests are additionally scoped: at the `own` level you see only requests you own)
Create templates and requests	`assessments:create`
Edit templates and questions; see, copy and reissue the public link; edit or expire requests; change repeats; review answers, send answers back, finish reviews	`assessments:update` (at the `own` level, on records you own)
Raise a third-party risk from a flagged answer	`assessments:update` on the request, `risks:create`, and read on the third party
Delete templates and unsubmitted requests	`assessments:delete`

The public assessment page bypasses all of this by design: it is unauthenticated, and the link is the only credential.

CHAPTER 15

Report Center

The Report Center turns the data you already keep in TraxrGRC into seven cross-module reports: a risk register, a compliance rollup, an exception inventory, a review schedule, awareness training completion, an incident listing and policy acknowledgements. Managers use it to see the state of the program at a glance, auditors use it to collect evidence, and GRC staff use it to find work that is overdue. Every report applies your own permissions to the underlying module, so two people opening the same report can see different rows. You can filter each report, download it as a CSV file for spreadsheets or evidence packs, and have it emailed to colleagues on a schedule.

The **Reports** entry in the sidebar has two pages: **Report Center** (the reports themselves) and **Scheduled** (reports sent by email, described under Scheduled reports below).

15.1 The landing page

Open **Reports → Report Center**. You see a card per report, each with a module badge, the title and a one-line description. Click a card to open the report.

A card appears only if you hold the `reports:read` permission **and** read access to at least one module the report draws on. If your permissions cover none of them, the page says "No reports available for your permissions."

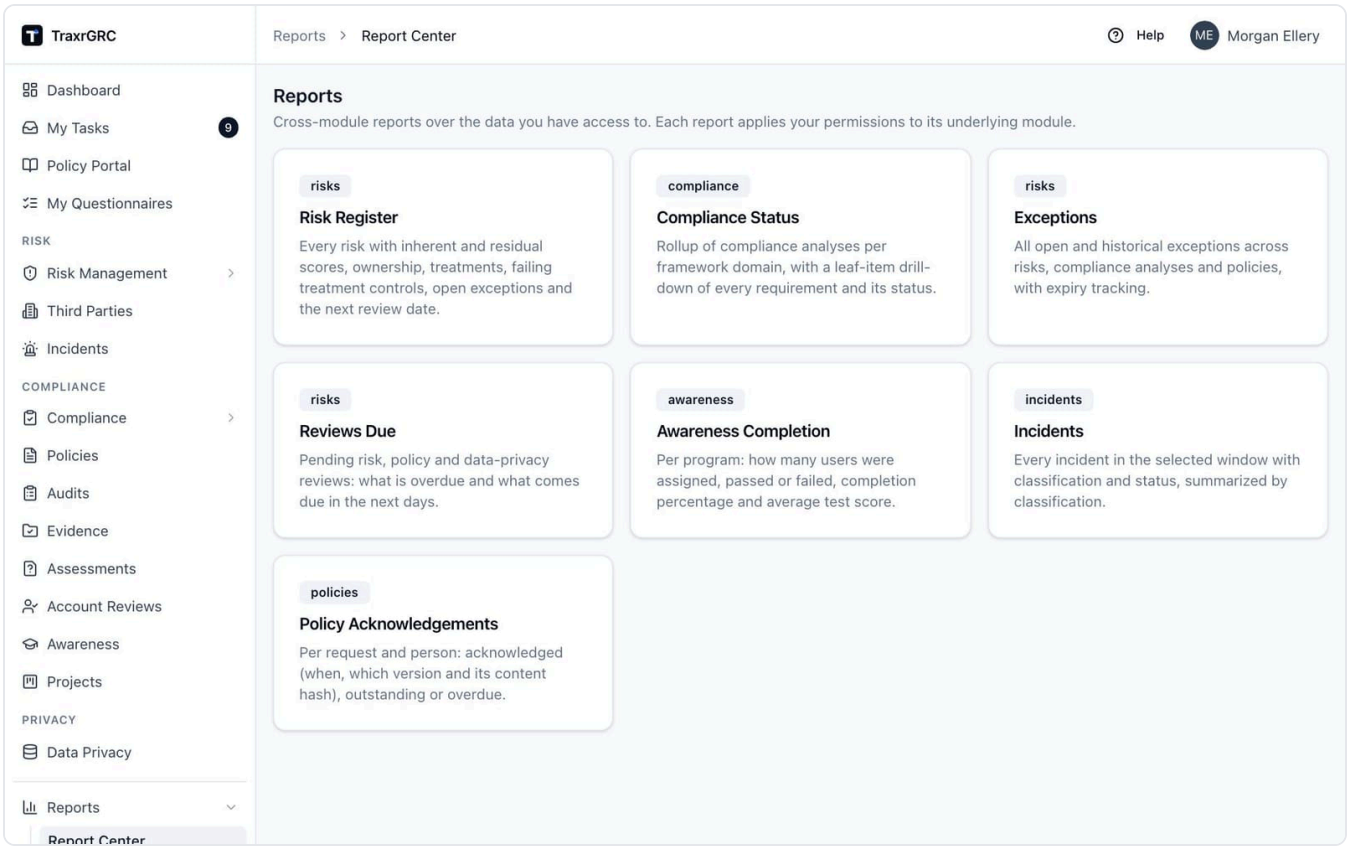


Figure 15.1. The Report Center landing page, with one card per report showing its module badge, title and description.

15.2 Anatomy of a report

Every report page is laid out the same way:

1. **Summary cards** at the top show the headline numbers, computed from the same rows as the table. A card's number renders in red when it flags a problem (details per report below).
2. **Filters** sit below the cards. Their values are carried in the page address, so a filtered view can be bookmarked or shared. Select filters always offer **All**. Click **Apply filters** to run them; **Reset** returns to the unfiltered report.
3. The **results table** lists the rows behind the current filters. Cells that need attention (overdue dates, expired items) render in red. When nothing matches, the table says "No results for the selected filters."
4. **Download CSV** (top right) downloads the report with the **filters you currently have applied**, in the same columns as the on-screen table. The file is named `report-<report>-<date>.csv` and opens correctly in Excel.
5. **Schedule** (top right, shown when you hold `reports:create`) opens a new scheduled report for this report, keeping the filters you have applied.

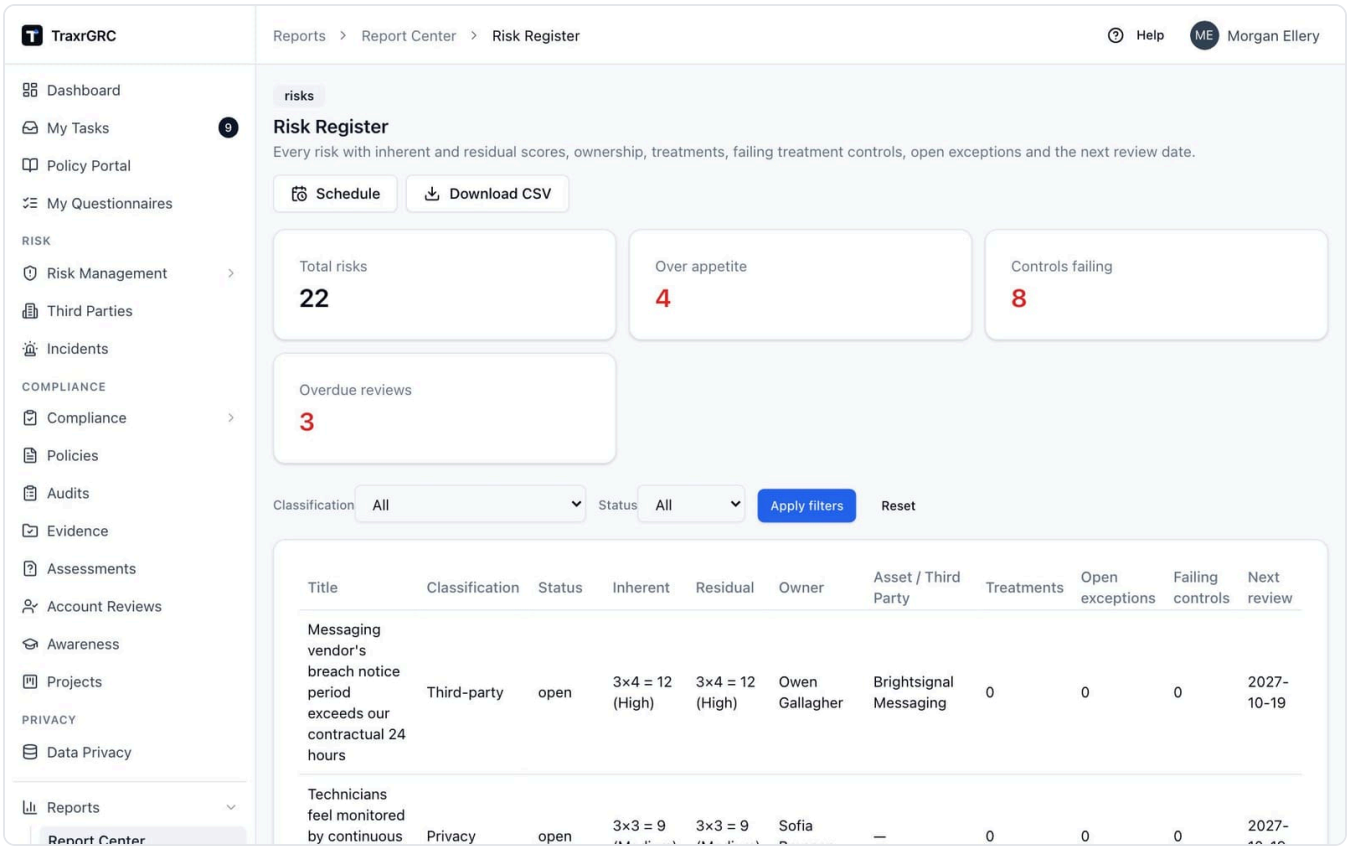


Figure 15.2. The Risk Register report: summary cards with red counts for risks over appetite and failing controls, the Classification and Status filters, and the results table.

15.3 Risk Register

Lists every risk you can read, with its scores and housekeeping data. Requires `reports:read` plus `risks:read`. With `risks:read` at the `own` level the register contains only risks you or your groups own.

Filter	Values	Address parameter
Classification	All, or one risk classification	`classification`
Status	All, Open, Closed, Archived	`status`

Column	Meaning
Title	The risk title.
Classification	The risk classification.
Status	Open, closed or archived.
Inherent, Residual	The full score cell, for example `4×5 = 20 (High)`: likelihood × impact = score, with the score's band.
Owner	The risk owner.
Asset / Third Party	The asset or third party the risk is scoped to.
Treatments	Number of treatments.
Open exceptions	Exceptions still requested or approved.
Failing controls	Treatment controls whose latest test failed or is overdue (see The Risks list), counted whether or not you can read them. Renders red when an open risk is flagged **Controls failing**.
Next review	The next review date; red when it is in the past.

Summary card	Counts	Red when
Total risks	Every row in the report	Never
Over appetite	Risks whose residual score is above your organization's appetite	Above zero
Controls failing	Open risks flagged **Controls failing**	Above zero
Overdue reviews	Risks whose next review date has passed	Above zero

15.4 Compliance Status

Rolls compliance analyses up per framework domain, with a drill-down of every requirement. Requires `reports:read` plus `compliance:read`.

Filter	Values	Address parameter
Framework package	All, or one compliance package	`package`

The main table has one row per domain: **Code**, **Domain**, **Not reviewed**, **Compliant**, **Partially compliant**, **Not compliant**, **N/A**, and **% compliant**, the percentage of in-scope requirement items (everything except N/A) that are Compliant.

The summary cards count every requirement by status: **Not reviewed**, **Compliant**, **Partially compliant**, **Not compliant** (red when above zero) and **N/A**. Only requirement items count, in the cards, the domain rows and the details table: headings and additional items do not, so the figures match the dashboard's compliance donut and the package page.

Below the main table, a **Requirement details** table lists every leaf requirement with its **Code**, **Requirement**, **Domain**, **Status**, **Owner**, linked **Control** and **Open exceptions**. The CSV download contains the per-domain table.

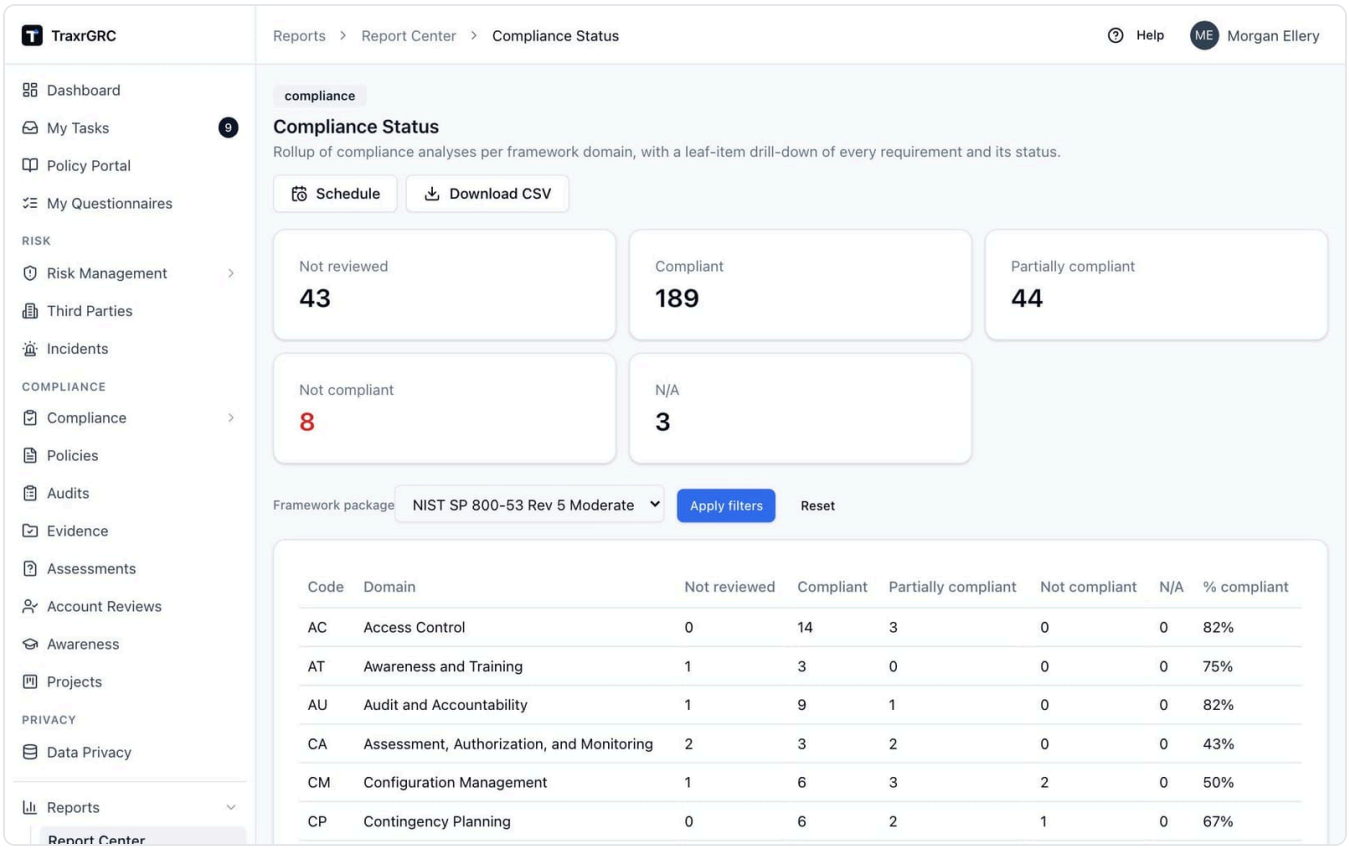


Figure 15.3. The Compliance Status report filtered to the NIST SP 800-53 Rev 5 Moderate package, with status counts per domain and the percentage compliant.

15.5 Exceptions

One table for all exceptions across risks, compliance analyses and policies, with expiry tracking. Requires `reports:read` plus read on at least one of `risks`, `compliance` and `policies`; the table includes only the sources you can read.

Filter	Values	Address parameter
Type	All, Risk, Compliance, Policy	`type`
Status	All, Requested, Approved, Rejected, Expired	`status`

Columns: **Type**, **Title**, **Parent** (the risk, requirement or policy the exception belongs to), **Status**, **Owner**, **Expires**. An approved exception whose expiry date has passed displays as **Expired**, and its Expires cell renders red. Rows sort by expiry date, soonest first.

Summary card	Counts
Awaiting decision	Exceptions in Requested status
Approved, in force	Approved exceptions not yet past their expiry date
Expiring in 14 days	Approved exceptions that expire within 14 days
Expired	Approved exceptions past their expiry date (a rejected request does not count)

Each card except **Approved, in force** renders red when above zero.

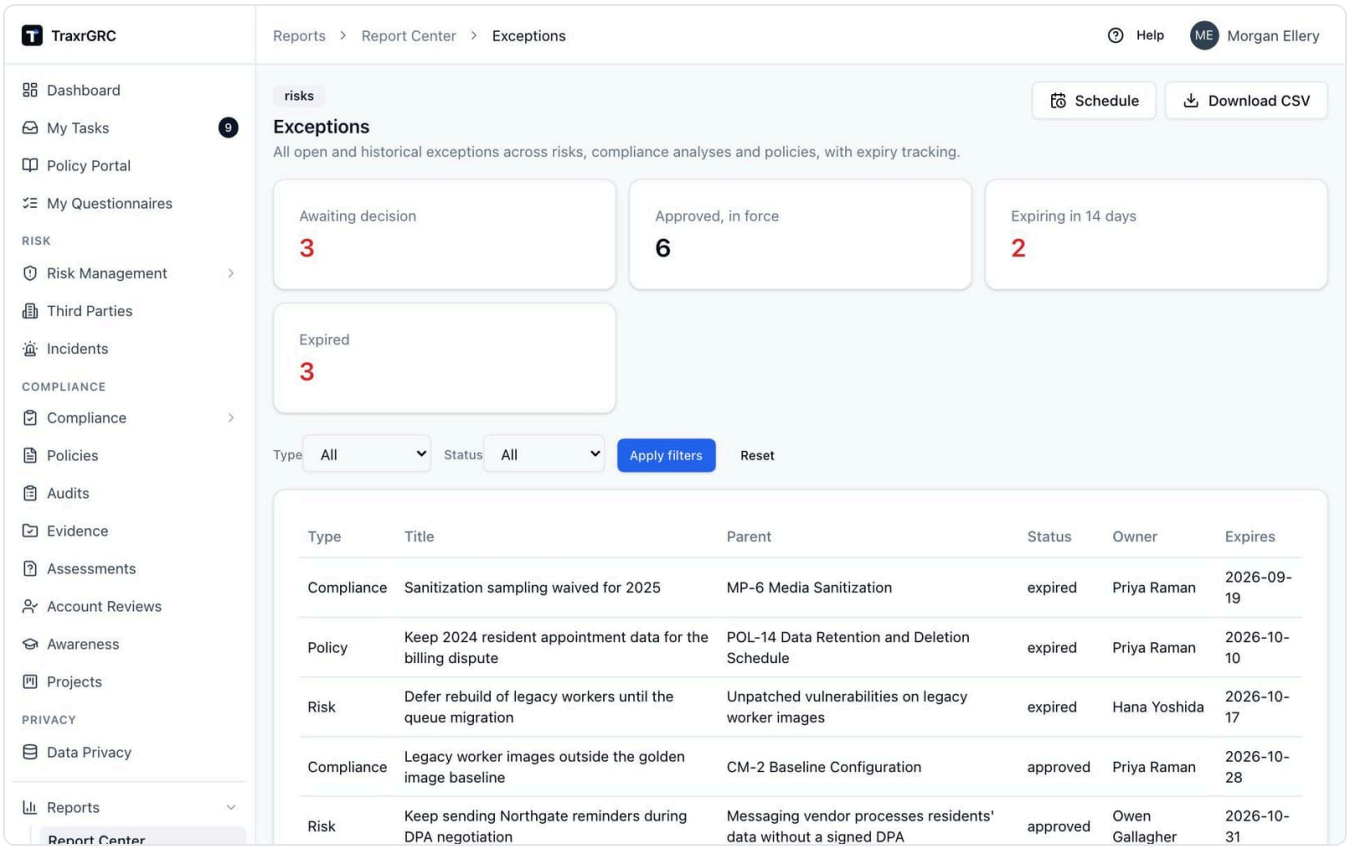


Figure 15.4. The Exceptions report listing risk, compliance and policy exceptions by expiry date, with red counts for open, expiring and expired exceptions.

15.6 Reviews Due

The consolidated review schedule for risks, policies and data privacy (data flow) records. Requires `reports:read` plus read on at least one of `risks`, `policies` and `dataPrivacy`; each source is limited by your ownership level in that module.

Filter	Values	Address parameter
Module	All, Risks, Policies, Data Privacy	`module`
Next N days	Number of days to look ahead (default 30)	`days`
Overdue only	Checkbox; when ticked, only past-due reviews are listed	`overdue` (value `1`)

Only reviews that have not been completed are listed. Columns: **Module**, **Record**, **Due date** (red when overdue), **Owner**. Rows sort by due date, earliest first. A review is overdue from the day after its due date, by the organization's time zone.

Summary cards: **Overdue** (red when above zero) and **Due in next N days**, matching the window you entered.

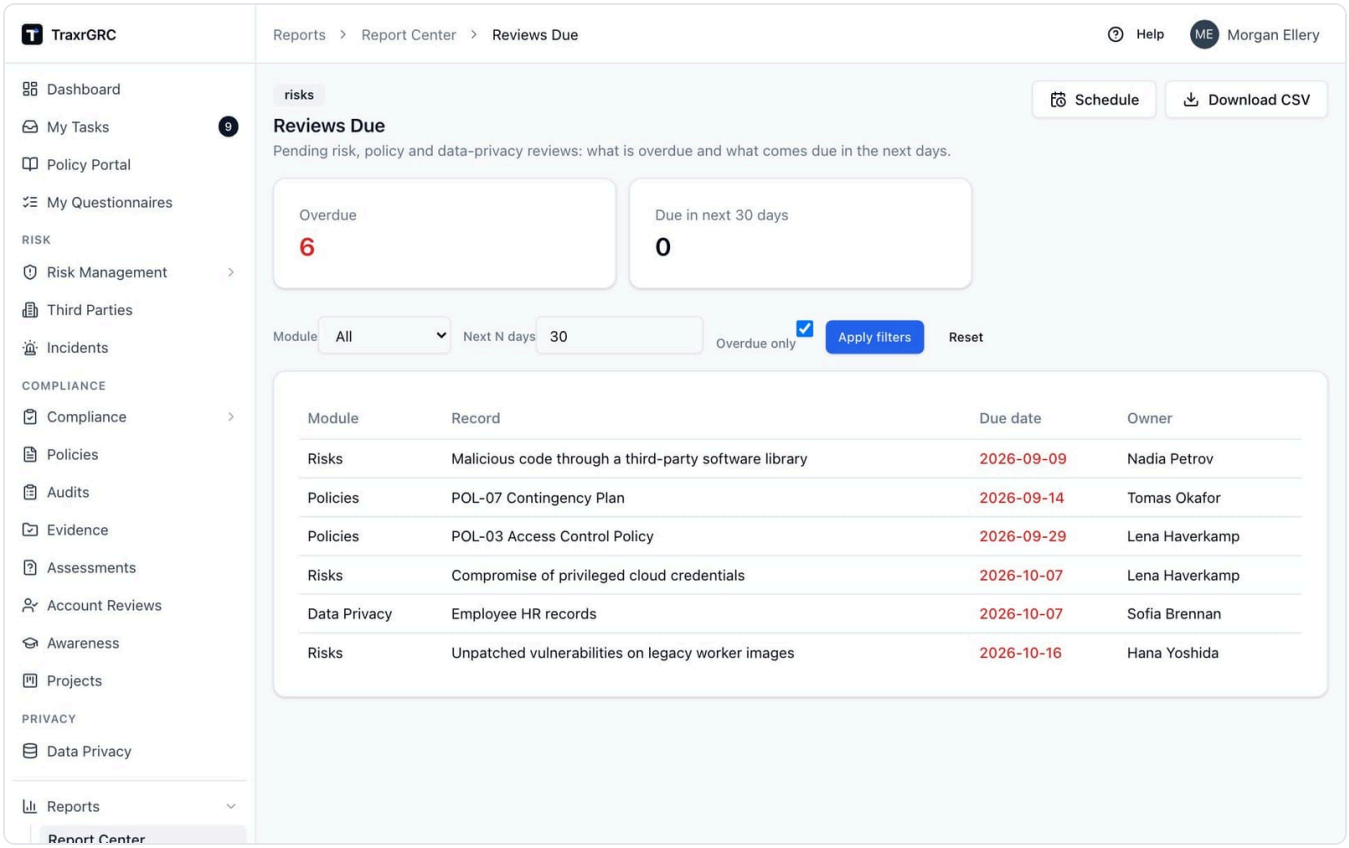


Figure 15.5. The Reviews Due report with Overdue only ticked, listing past-due risk, policy and data privacy reviews with their due dates in red.

15.7 Awareness Completion

Training completion per awareness program. Requires `reports:read` plus `awareness:read`. Only **active** programs are listed.

Filter	Values	Address parameter
Program	All, or one awareness program	`program`

Column	Meaning
Program, Status	The program and its status.
Assigned	The number of active users in TraxrGRC.
Passed, Failed	Recorded test results that passed or failed.
Completion %	Passed results as a share of assigned users.
Average score	The mean test score across all recorded results.

Summary cards: **Active programs**, **Assigned users**, **Completion** (the passed share across all listed programs) and **Average score** (the mean of the per-program averages, or a dash when no results exist).

15.8 Incidents

Every incident in the selected window. Requires `reports:read` plus `incidents:read`. With `incidents:read` at the `own` level you see only incidents you or your groups own.

Filter	Values	Address parameter
Classification	All, or one incident classification	`classification`
Opened from	Date; incidents opened on or after	`from`
Opened to	Date; incidents opened on or before	`to`

Columns: **Opened** (date), **Title**, **Classification** ("Unclassified" when none is set), **Status**, **Owner**. Rows sort newest first.

Summary cards: **Total incidents**, then one card per classification present in the results, with its count.

15.9 Policy Acknowledgements

One row per person per acknowledgement request (see Policy portal and acknowledgements). Requires `reports:read` plus `policies:read`. With `policies:read` at the `own` level you see only requests on policies you or your groups own.

Filter	Values	Address parameter
Policy	All, or one policy with requests	`policy`
Status	All, Acknowledged, Outstanding, Overdue	`status`
Request	All, Active, Superseded, Closed	`state`

Column	Meaning
Policy, Version	The policy and the version the request asked people to acknowledge.
Due	The request's due date.
Request	The request's state: Active, Superseded or Closed.
Name, Email	The person.
Active	Whether the person's account is still active.
In audience	Whether the person is still in the request's audience.
Status	Acknowledged, Outstanding or Overdue (overdue in red).
Acknowledged	When the person acknowledged: in your time zone on screen, UTC in the CSV.
Content hash	SHA-256 of the version text acknowledged.

The report lists everyone in each request's audience now and everyone who acknowledged it, including people since deactivated or moved out of the audience.

Summary cards: **Requests**, **Acknowledged**, **Outstanding** and **Overdue** (red when above zero). Outstanding and Overdue count only people still in the audience.

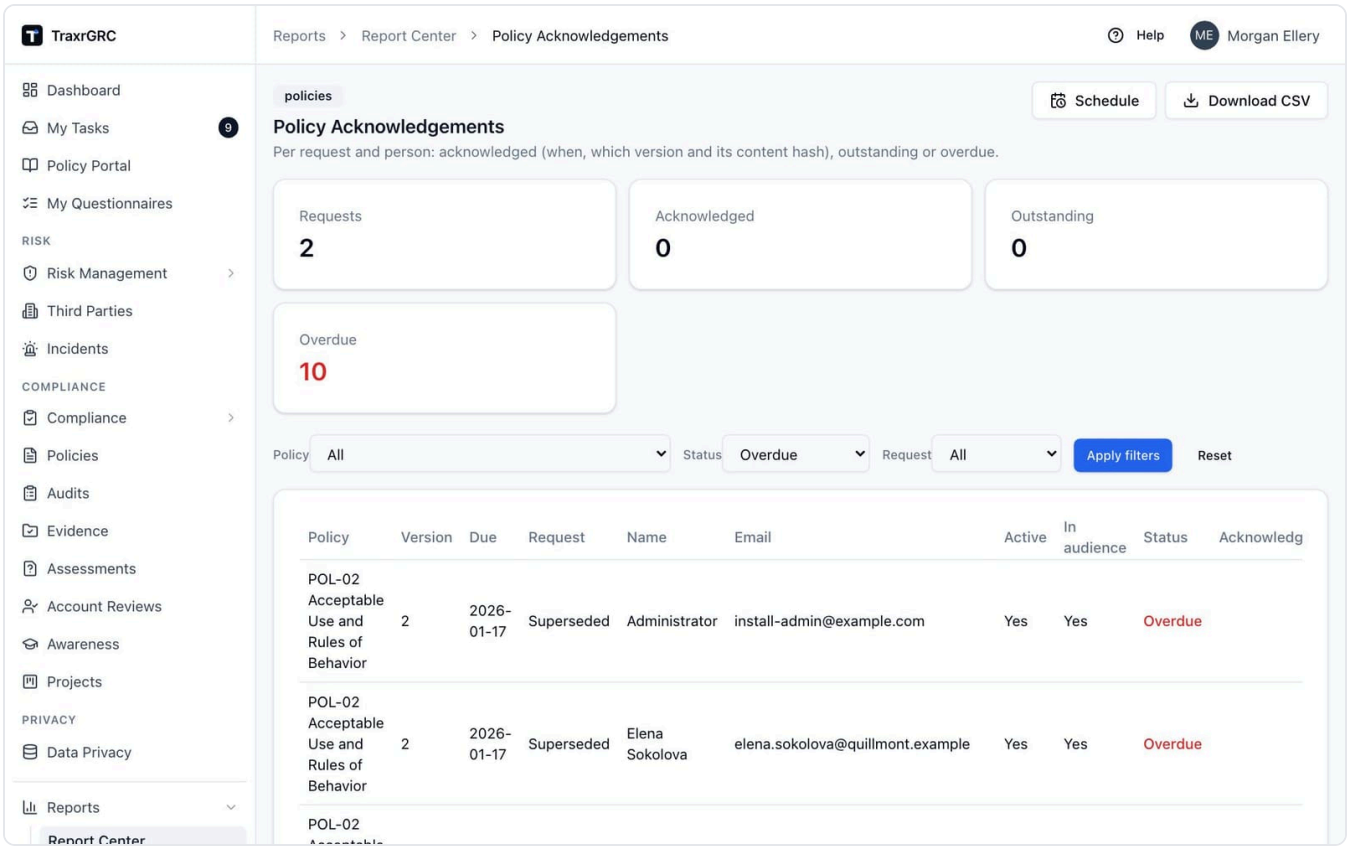


Figure 15.6. The Policy Acknowledgements report filtered to overdue people, showing each person's policy, version, due date and status.

15.10 Scheduled reports

A report can be emailed to people in this installation on a schedule, for example the compliance gaps to management on the first of every month. Open **Reports → Scheduled**, or click **Schedule** on any Report Center report (the schedule keeps the filters you have applied) or **Schedule by email** on **Compliance → Gaps**.

What can be scheduled, and in which format:

Report	Attached as
Any Report Center report	CSV, with the filters chosen
Compliance package report	PDF, for one package
Compliance gaps (POA&M)	CSV, optionally narrowed to a package, a gap status, a coverage, or gaps with overdue projects

15.10.1 The schedules list

The **Scheduled reports** page lists the schedules for reports you can open yourself; schedules for reports you cannot open are hidden.

Column	Meaning
Name	The schedule's name, with a **Paused** badge when it is not active.
Report	The report and its format, for example "Risk Register (CSV)".
When	The timing in words, for example "Weekly on Monday at 07:30 America/New_York".
Next run	When it is next sent, in your own time zone; a dash while paused.
Recipients	The users, groups (prefixed "group") and roles (prefixed "role") it is sent to, and how many people have unsubscribed.

The buttons at the end of each row are:

Button	Action	Needs
Paper plane	**Send now**: queue a run straight away	Update on Reports; schedule active
Pause / play	Pause or resume the schedule	Update on Reports
Pencil	Edit the schedule	Update on Reports
Clock	Show the schedule's recent deliveries	Any schedule right
Trash can	Delete the schedule	Delete on Reports

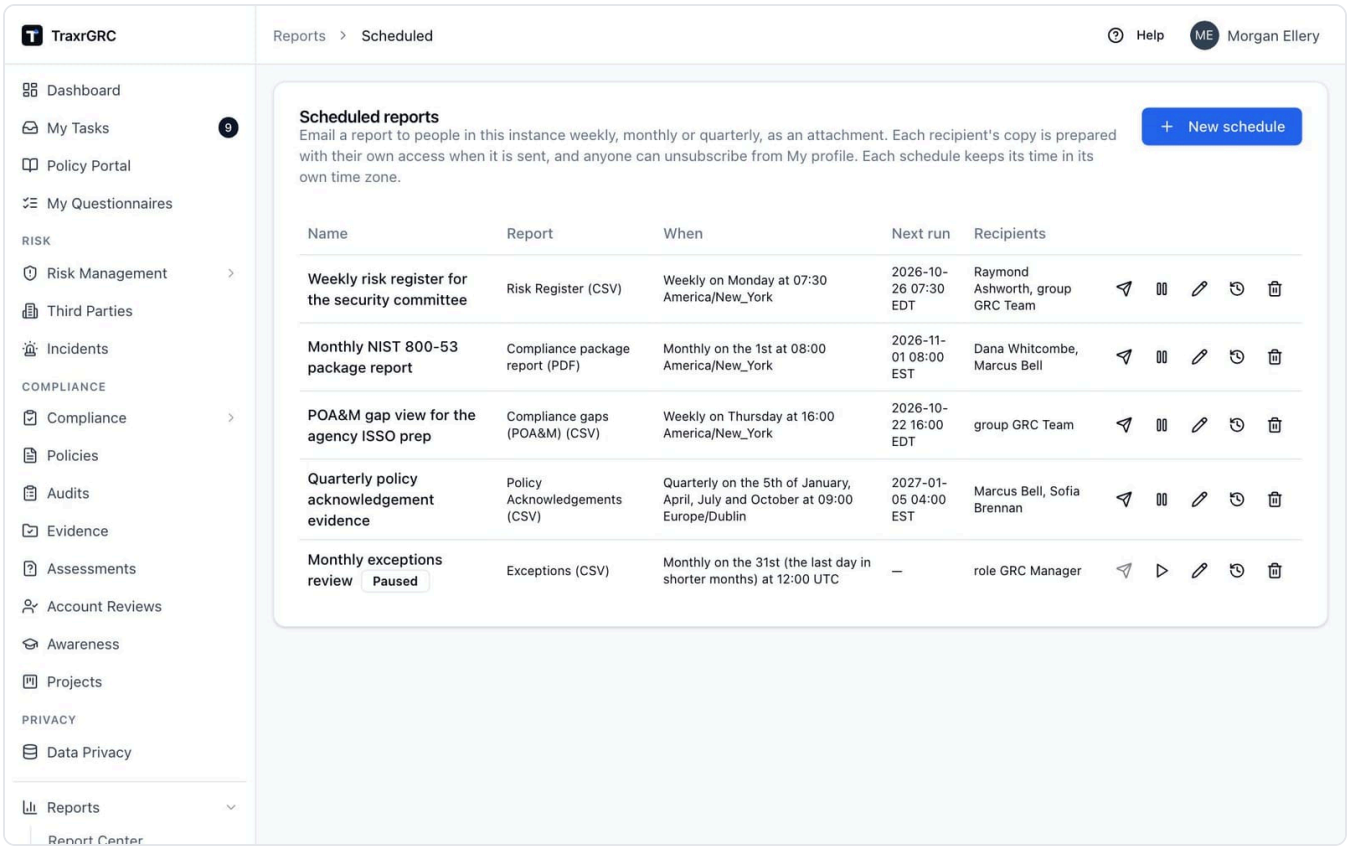


Figure 15.7. The Scheduled reports page listing the demo schedules, including a paused monthly exceptions review, with their timing, next run and recipients.

15.10.2 Creating a schedule

1. Click **New schedule**, or **Schedule** on a report page. The **New scheduled report** dialog opens.
2. Enter a **Name** and choose the **Report**. Each option shows its format, for example "Compliance package report (PDF)".
3. Set the report's options:
    - For a Report Center report, **Filters** shows the filters carried over from the report page. To schedule a filtered report, apply the filters on the report and use its **Schedule** button; **Clear filters** sends the whole report.
    - For the compliance package report, choose the **Package**.
    - For the compliance gaps, optionally choose a **Package** (or **All packages**), a **Status** (**Both gap statuses**, Partially compliant or Not compliant), a **Coverage** (**Any coverage**, No project, In remediation or Projects completed), and tick **Only gaps with overdue projects**.
4. Set the **Frequency**, the **Day** or **Day of the month**, the **Time** and the **Time zone** (see the field reference below).
5. Tick the **Recipients** under **Users**, **Groups** and **Roles**.
6. Leave **Active (off pauses it)** on to start sending, then click **Create schedule**.

To change a schedule, click its pencil, edit the fields and click **Save schedule**.

Field	What it means
Name	Shown in the email subject, the delivery log and each recipient's **My profile** page.
Report	What is sent, with its filters or package.
Frequency	**Weekly** on a day of the week, **Monthly** on a day of the month, or **Quarterly** on a day of the month in January, April, July and October. A day past the end of a shorter month (the 31st in April, the 29th to 31st in February) sends on that month's last day.
Time and Time zone	The time of day it is sent, in the chosen zone, which starts as your own (UTC if you have none). The time stays the same in that zone when daylight saving starts or ends: 09:00 New York is 09:00 in both summer and winter. A time the clocks skip (02:30 on the night they go forward) sends when they jump past it, at 03:30; a time that happens twice (01:30 on the night they go back) sends once, the first time. Schedules made before time zones existed are in UTC.
Recipients	Named users, the members of chosen groups, and the members of every group carrying a chosen role. Active users of this installation only; no outside address can be entered.
Active	Off pauses the schedule.

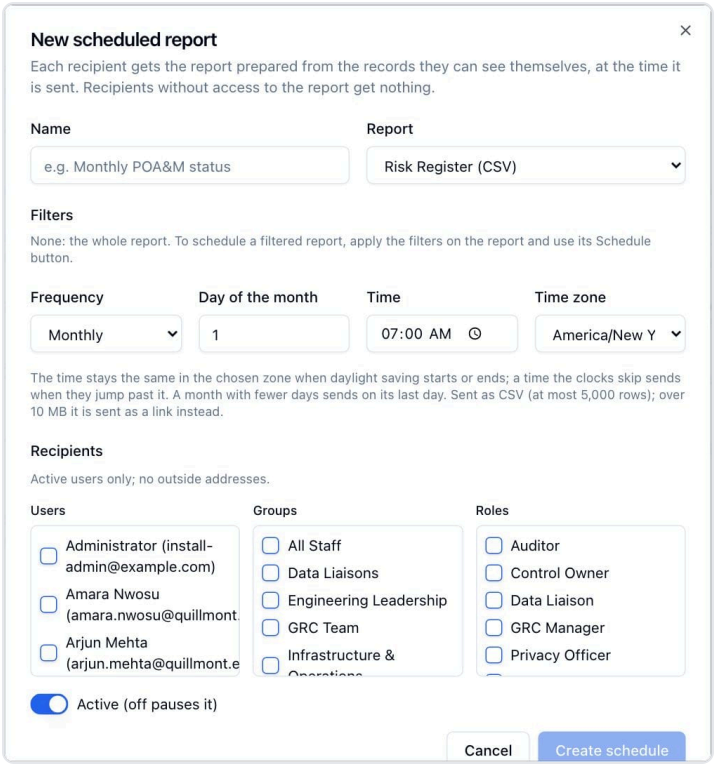


Figure 15.8. The New scheduled report dialog with the report, frequency, time, time zone and recipient fields.

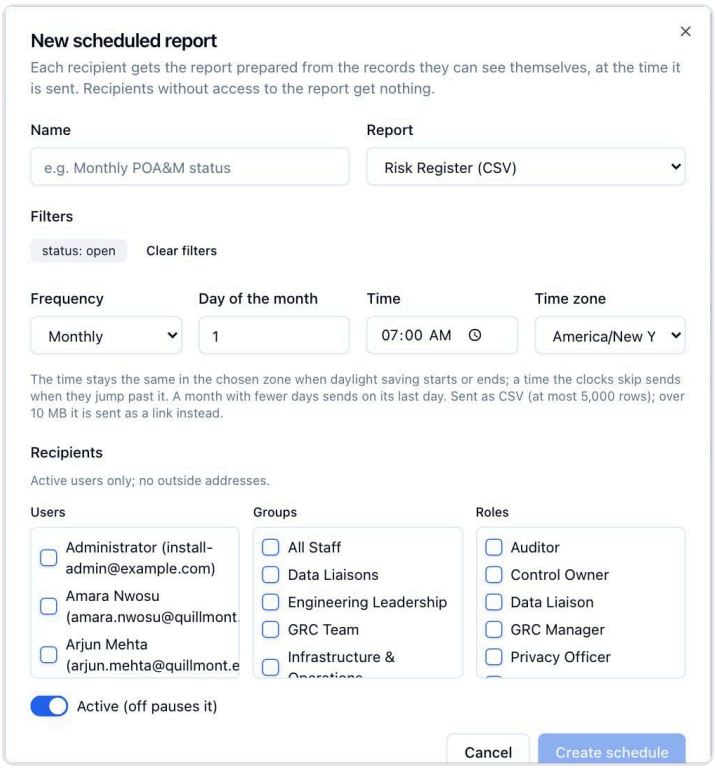


Figure 15.9. The schedule dialog opened from the Risk Register report's Schedule button, carrying the Status filter applied on the report.

15.10.3 How a scheduled report is sent

- **Each recipient gets their own copy.** The report is prepared separately for every recipient, with their permissions and record visibility at the moment it is sent, exactly as if they had opened it themselves. Two recipients can get different files. A recipient who cannot open the report (no `reports:read` or no read on its module for a Report Center report, no Compliance read for the gaps, Compliance read at "all" for the package PDF) is skipped and the skip is logged.
- CSV files hold at most 5,000 rows, like every other export; the email says so when the report had more. Copies are prepared one at a time by the background worker, so a large schedule never slows the application.
- A file over 10 MB is not attached: the email carries a link to the report in the app instead, and the delivery log notes it.
- If the worker was stopped when a schedule was due, the schedule runs once when it is back, not once for every slot it missed, and then waits for its next slot. Pausing and resuming, or changing a schedule's timing, also starts from the next slot after now.
- **Send now** queues a run straight away, for testing. A message confirms how many emails were queued and how many were skipped. It can be used once every five minutes per schedule, and not while the schedule is paused.
- The email states when the report was prepared, in the recipient's own time zone.
- Every email says why the person receives it and links to **My profile**, where they can unsubscribe.

15.10.4 Delivery history

Click a schedule's clock icon to open **Recent deliveries**. It has one row per recipient and run, with **Queued** (when), **To** (the address), **Status** and **Attachment** (the file sent, or "too large, sent as a link"). A suppressed row names the reason:

Reason	Meaning
`no-access`	The recipient may not open the report.
`opted-out`	The recipient unsubscribed on **My profile**.
`recipient-inactive`	The recipient's account is deactivated.
`template-disabled`	The scheduled-report email template is turned off under Settings → Notifications.
`schedule-paused`, `schedule-deleted`	The schedule was paused or deleted before a queued email was sent.
`report-unavailable`	The report could not be prepared, for example because its package was deleted.

Administrators also see every delivery under **Settings → Notifications → Delivery log**.

15.10.5 Deleting a schedule

Click the trash can and confirm **Delete this schedule?**. The schedule stops sending, including emails queued and not yet sent. Its past deliveries stay in the log.

15.10.6 Unsubscribing

Every recipient sees the schedules sent to them on **My profile**, under **Scheduled reports**, with the report and its timing. Turn a schedule's switch off to stop receiving it, or on again to resume. The schedules list shows how many people have unsubscribed from each schedule.

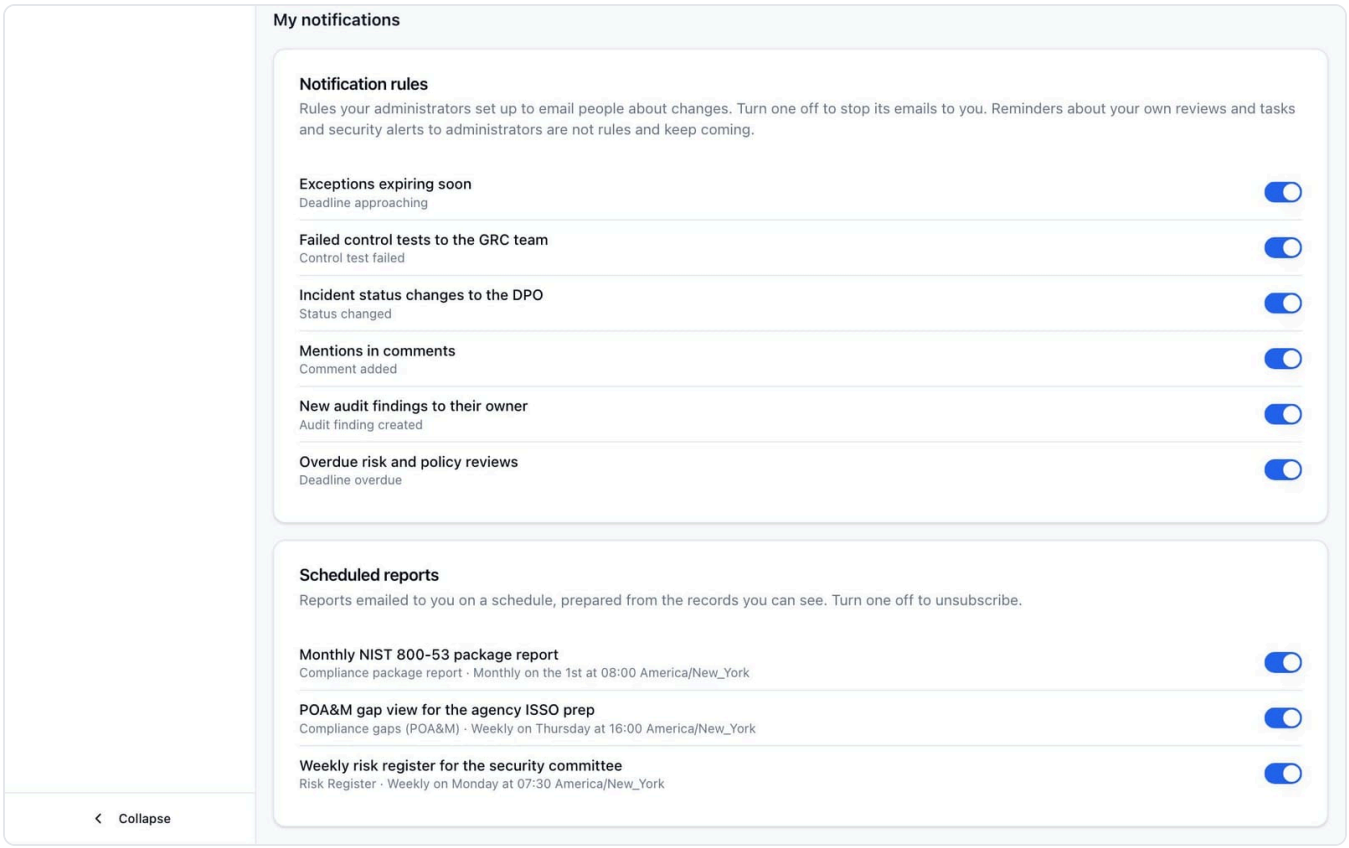


Figure 15.10. The Scheduled reports card on My profile, listing the reports sent to the user with a switch to unsubscribe from each.

15.10.7 Who can manage schedules

Action	Needs
Open **Reports → Scheduled**	Create, Update or Delete on Reports
Create a schedule	Create on Reports
Edit, pause, resume, Send now	Update on Reports
Delete a schedule	Delete on Reports

In every case you must be able to open the report yourself. Read on Reports alone lets someone receive reports but not schedule them. On a new installation the GRC Manager role carries all three rights; on an installation upgraded from an earlier version only administrators can schedule until those rights are added to a role (Settings → Roles).

Every schedule created, changed, paused, resumed, deleted or sent now is recorded in the Audit Log. A new or upgraded installation has no schedules, so nothing is sent until someone creates one. Emails need SMTP configured (Settings → Integrations).

15.11 Who can see reports

- Opening the Report Center at all requires the `reports:read` permission.
- Each report card additionally requires `read` on at least one of its underlying modules; cards you do not qualify for are hidden.
- Within a report, rows follow your record-level access: reports over risks, policies, data privacy and incidents respect the `own` level (you see records you or your groups own), while the compliance and awareness reports are all-or-nothing for that module's readers.
- **Download CSV** applies exactly the same permission checks and filters as the on-screen report, so an export never contains more than the table shows.
- Administrators bypass all of the above.

CHAPTER 16

REST API

The REST API lets scripts and other systems read and maintain TraxrGRC records without the browser: loading risks from a spreadsheet, exporting compliance status to a dashboard, or creating incidents from a ticketing tool. It is used by developers and administrators who build such integrations. This chapter explains how access works and what is available; the machine-readable reference is the OpenAPI document, served at `/api/v1/openapi` on your installation. That address needs no token, so you can load it into an OpenAPI viewer or a client generator directly.

The API covers the main records of each module. Some screens have no API equivalent, among them Settings, the Report Center and its CSV downloads, and the Dashboard.

16.1 Authentication

API requests use bearer tokens:

1. A user with Create on API tokens (an administrator, by default) creates a token in **Settings → API Tokens** (see Administration & Settings). Give it a name, an optional expiry date, and the module scopes it may use.
2. The full token, starting with `aat_`, is shown **once**, at creation. Copy it somewhere safe. TraxrGRC stores only a hash of it.
3. Send it on every request in the `Authorization` header:

```
curl -H "Authorization: Bearer aat_your_token_here" \
  https://grc.example.com/api/v1/risks
```

A request is refused in these cases:

Status	Body	Cause
401	`{ "error": "unauthorized" }`	No token, an unknown token, or a token that is revoked or expired; also a token whose creator has been deleted or deactivated.
403	`{ "error": "forbidden" }`	The token's scopes do not include the module and right the request needs, or its creator no longer holds that right.

Revoked or expired tokens stop working immediately. Every refused request is recorded in the Audit Log (module `security`, action `permission.denied`) with the reason and the token.

16.2 Conventions

Topic	Behavior
Base URL	`/api/v1` on your installation's address.
Format	JSON request and response bodies.
Success envelope	`{ "data": ... }`. A create returns status 201 with `{ "data": { "id": "..." } }`.
Error envelope	`{ "error": "..." }` with an HTTP status: 400 for invalid input (`invalid-input`), 401 unauthenticated, 403 forbidden, 404 not found, 409 for a conflict with the record's current state, and a short message for other refusals.
Lists	Paginated with `page` (from 1) and `pageSize` (default 25, at most 200), for example `GET /api/v1/risks?page=2&pageSize=50`. The data is `{ rows, total, page, pageSize }`.
Identity	A token acts as the user who created it. That user's permissions apply on top of the token's scopes, including the `own` ownership level, so a list returns only the records the creator could see in the browser.
Scope of side effects	Scopes bound everything a request does. A request that also touches another module (a DPIA creating a risk, a record linked to a vendor) needs that module's right in the token's scopes too, or it is refused (403, or 404 for a link to a record the token cannot read).
Audit	API changes are recorded in the Audit Log like any other change, naming both the token and its creator.
Notifications	Changes made through the API send no notification rule emails unless an administrator turns this on (see Notifications in Administration & Settings).

16.3 What is available

Each operation in the OpenAPI document states the scope it requires, for example "Requires the API token scope `risks:read`". The scope names match the role permission names: reading needs Read on the module, and changes need Create, Update, Delete or Review as described in the module's chapter.

Area	Paths under `/api/v1`	Covers
Risks	`/risks`, `/risk-classifications`, `/risk-scale`	Risks with computed scores, treatments, reviews; the classification list and the risk scale (changing those needs Settings rights)
Controls	`/controls`	Controls and their tests
Assets	`/assets`	The asset catalog
Business units	`/business-units`	The business unit hierarchy
Compliance	`/compliance/packages`, `/compliance/analyses`, `/compliance/poam`	Packages and their items, analyses and their exceptions, the POA&M register (read; `?status=`, `?severity=`, `?source=`, `?overdue=true`)
Policies	`/policies`	Policies, content, versions, reviews and approvals recorded outside the app
Third parties	`/third-parties`	Third parties, contacts, linked risks, sub-processors and transfer impact assessments
Incidents	`/incidents`	Incidents, closing, breach details and breach notifications
Audits	`/audits`	Audit plans and findings
Projects	`/projects`	Projects
Awareness	`/awareness/programs`, `/awareness/results`	Programs, test submissions and results
Data privacy	`/data-privacy`	Records of processing, transfers, links, reviews, the Article 30 register and DPIAs
Account reviews	`/account-reviews`	Review programs, their items and CSV import
Assessments	`/assessments/templates`, `/assessments/requests`	Questionnaire templates and requests
Custom fields	`/custom-fields`	Field definitions and per-record values

Area	Paths under `/api/v1`	Covers
Evidence	`/evidence`	Evidence and reports, their links, files with SHA-256 (list, upload, download with an `X-Content-SHA256` header, remove), locking and unlocking, Verify files, and the auditor package (`/evidence/package`, a zip)

Consult the OpenAPI document for the exact paths, parameters, request bodies and responses.

16.4 Related access paths

- **Online assessments** are answered by the external contact through the tokenized public link (`/assess/<token>`), described in Online Assessments. That page needs no API token.
- **AI assistants** connect through the Model Context Protocol endpoint at `/api/mcp`, using the same API tokens (see AI assistants in Administration & Settings).
- **SCIM** user provisioning has its own tokens and address (see Administration & Settings); API tokens do not work there.

Index

#

72-hour clock 18, 171, 173

A

A third party's files 260
Acknowledgement 29, 65, 143, 144, 146, 148, 294
Acknowledging a policy 145
Adding a single account 249
Adding evidence 257
Adding items to a program 249
Adding or editing a finding 185
After the decision 229
AI assistants 62
An evidence record 258
Analysis statuses 114
Anatomy of a report 287
Annual checks 238
Answering (the Data Liaison's view) 235
API token 39, 45, 56, 58, 62, 74, 75, 306
API Tokens 74
Approval workflow 58, 61, 134, 135, 136, 137, 228
Approving or rejecting 135, 228
Archive 18, 79, 85, 129, 143, 202, 231, 253
Archive and reactivate 143
Archiving, reactivating and deleting 253
Article 30 completeness 220
Article 30 register 39, 61, 125, 213, 220, 224, 243
Asking for acknowledgement 146
Assessing an incident as a breach 171
Asset 32, 34, 81, 98, 167, 168, 223, 304
Assets 98
Attachment 34, 93, 143, 158, 178, 190, 211, 231
Attachments 34
Audit finding 39, 117, 185, 187, 188, 197
Audit Log 70
Audit log 37, 39, 45, 62, 70, 73, 75
Authentication 303
Automatic reminders 209
Awareness Completion 293
Awareness permissions 212

B

Backup 39, 75
Backups 75
Branding 59
Building a questionnaire 234
Bulk updates 113
Business unit 78, 81, 84, 98, 99, 304
Business Units 99

C

Changing status 195
Classification 18, 78, 79, 99, 165, 288, 293
Closing and reopening an incident 170
Comment 35, 93, 143, 178, 190, 211, 231, 251
Comments 35
Comments, attachments and history 93, 178, 190, 198, 251
Comments, Attachments and History 158
Comments, Attachments, History and Custom Fields 143
Comments, attachments, history and custom fields 211, 231
Common controls and inheritance 108
Completion and evidence 148
Compliance requirements remediated 197
Compliance Status 289
Contacts tab 156
Content tab: the training material 205
Control 93, 95, 96, 98, 108, 116, 117, 188
Control testing 24, 65, 93, 95, 96, 98, 99
Control testing settings 98
Controls 93
Conventions 304
Creating a policy 130
Creating a program 201, 246
Creating a project 193
Creating a record 215
Creating a remediation project 187
Creating a request 269
Creating a risk 81
Creating a schedule 297
Creating a template 266
Creating a third party 154
Creating an audit plan 182
Creating an incident 167
Crosswalks between frameworks 110
CSV import 32, 36, 67, 107, 110, 111, 121, 249
Custom field 32, 39, 45, 48, 143, 178, 211, 231
Custom Fields 48

D

Dashboard 12, 18, 23, 24, 45, 138, 245
Data Liaison 18, 45, 232, 235, 238, 243
Data map 18, 45, 213, 232, 241, 243
Data mapping questionnaires 232
Data privacy permissions 243
Data processing agreement 159, 65, 153
Data protection 158
Data tables 31
Deciding items in bulk 251
Deciding one item 250
Deleting a plan 185
Deleting a policy 150
Deleting a program 212
Deleting a project 199
Deleting a record 38, 232
Deleting a schedule 300
Deleting a template 269
Deleting a third party 162
Deleting packages and items 126
Delivery history 300
Delivery log 70
Documenting the record 218
DPIA 61, 224, 225, 226, 227, 228, 229, 243
DPIAs 224
DPO advice 227

E

Editing a plan and changing its status 184
Editing and deleting 178
Evidence 255, 257, 258, 259, 261, 263, 264
Evidence on other records 259
Evidence packages for auditors 263
Exception 32, 67, 89, 91, 124, 129, 141, 290
Exceptions 91, 141, 290
Exceptions on a clause 124
Expiry and reminders 262
Export 31, 32, 36, 39, 70, 75, 108, 121
Exporting a package 108

F

Files on controls, requirements and account reviews 260
Filtering by family 113
Finding status workflow 187
Finding the Settings screens 39
Findings raised by failed control audits 188

G

Group 24, 39, 42, 45, 56, 146, 212
Groups 45

H

Header, help and user menu 16
History 37, 93, 143, 158, 178, 190, 211, 231
History tabs 37
How a scheduled report is sent 299
How scoring works 83

I

Importing a package 107
Importing accounts from CSV 249
Importing from CSV 32
Incidents 293
Integrations 50
Integrity 73

L

Label 34, 48, 79, 93, 98, 129, 131, 185
Labels 34, 48
LDAP 52, 12, 42, 50, 54, 58, 70
Linked assets, risks and controls 223
Locked evidence and file fingerprints 261

M

Managing requests 148
Managing the classification catalog 165
Managing the shared dashboard 24
Mapping controls, policies and risks 116
MCP 34, 62, 74, 280, 306
My notifications 29
My profile 27
My Tasks 24

N

NIST SP 800-53 10, 103, 108, 109, 110, 121
Notification rule 27, 29, 39, 65, 67, 230, 284
Notification rules 67
Notifications 65
Notifications and audit trail 284

O

Open audit findings on mapped controls 117
Open findings 188
Organization-defined values 109
Owner 24, 119, 141, 151, 161, 162, 185, 200

P

Packages 103
Password 12, 27, 28, 42, 51, 52, 54, 70
Password change 28
Permissions 102, 127, 151, 162, 164, 180, 191, 212
Personal data breach 18, 32, 65, 67, 168, 171, 173
Personal data breaches 171
POA&M 10, 103, 117, 119, 121, 127, 188
Policies 61
Policy Acknowledgements 294
Policy portal 18, 29, 129, 131, 144, 146, 151
Policy portal and acknowledgements 144
Privacy 61
Program status: draft, active, archived 202
Provisioning (SCIM) 56
Publishing 136
Putting a policy on the portal 146

Q

Questionnaire 18, 45, 232, 234, 235, 236

R

Record of processing 24, 48, 160, 173, 213, 215, 224
Record status: draft, active, archived 231
Record tabs 119
Recording an approval made outside the app 137
Recording the analysis 115
Records of processing that name this vendor 160
Reissuing the public link 274
Related access paths 306
Remediation project 117, 119, 187, 191, 193, 194, 199
Remediation projects 117
Reminders 240
Repeating assessments 280
Replacing earlier evidence 261
Report schedule 29, 295, 299
Requests 269
Residual risk 18, 79, 81, 83, 84, 99, 228, 288
Retakes 210
Retention 73
Review 18, 89, 139, 229, 230, 245, 260, 292
Review notifications 230
Reviewing answers 236, 275
Reviewing items: keep or disable 250
Reviews 89
Reviews Due 292
Reviews tab 139
Risk appetite 18, 65, 83, 84, 86, 89, 99, 288
Risk Register 288
Risk score 18, 81, 83, 99
Risk Settings 99
Risks tab 157
Roles 45

S

SAML 12, 42, 50, 54, 56, 70, 73, 232
SAML single sign-on 54
Scheduled reports 295
Scheduler 69
Scheduling and recording control tests 96
SCIM 36, 39, 42, 45, 56, 75, 306
Security alerts 73
Sending a questionnaire 235
Sending a reminder 208
Session 14, 16, 39, 42, 54, 56, 58, 70
Sessions 58
Shared UI conventions 30
Sidebar 15
Sign-in security 28
Signing in 12
Signing out and session length 14
SMTP 51, 50, 73
Starting and screening a DPIA 225
Status and actions 85
Status badges and overdue highlighting 33
Sub-processor 158, 161, 218
Sub-processors 161
Submitting for approval 135

T

Taking a test 209
Templates 65
Templates and the question editor 266
Test tab: building the question set 206
The analysis workspace 111
The app layout 14
The Article 30 register 224
The Assessments page 265
The audit plans list 180
The breach card 173
The catalog page 110
The compliance report (PDF) 125
The control page 95
The Dashboard 18
The data map 241
The detail page 156
The draft → approval → publish workflow 134
The Evidence page 255
The full assessment 226
The gap view (POA&M) 119
The incident detail page 168
The incidents list 165
The landing page 286
The linked audit finding 197
The NIST SP 800-53 catalog 109
The package page 105
The plan detail page 183
The POA&M register 121
The policies list 129
The policy detail page 131
The Policy Portal 29
The program detail page 203
The program page 247
The programs list 200, 245
The project detail page 194
The projects list 191
The public assessment page (for external contacts) 280
The Questionnaires tab 232

The record detail page 216
The records list 213
The register 153
The request page 273
The requests list 271
The Reviews tab 229
The risk detail page 84
The risk register and its tabs 78
The Risks list 79
The schedules list 296
The template page 266
Third party 18, 121, 153, 154, 160, 162, 260
Time zone 27, 36, 42, 52, 54, 56, 297
Times and time zones 36
Tracking tab: who has completed the program 207
Transfer impact assessment 65, 153, 161, 221
Transfer impact assessments 161
Transfers 221
Treatment 78, 79, 83, 84, 86, 91, 93, 288
Treatments 86
Two-step verification 12, 28, 42, 54, 58, 65, 70, 73

U

Unpublished changes 138
Unsubscribing 300
Uploading a reissued document 138
User 16, 42, 45, 52, 54, 56, 70
Users 42

V

Version 103, 131, 136, 138, 139, 144, 148
Versions tab 139

W

What a decision records 251
What is available 304
What is recorded 148
What the portal shows 144
What you can manage 10
When a page is off limits 38
Who can be a Data Liaison 232
Who can do what 151, 162, 164, 180, 191, 254, 285
Who can do what (Compliance Management) 127
Who can do what (Risk Management) 102
Who can do what with evidence 264
Who can manage schedules 301
Who can see reports 302
Working on an item 114
Working with findings 185
Writing content (Content tab) 132

Y

Your own dashboard 23