Documentation Migration

Working record of the September 2026 migration to the SiteTrax.io Intelligence documentation set, covering what was created, updated and retired, the conflicts found between the documentation and the live application, the defects surfaced while auditing, and the recommendations arising from it.

This is an internal project artifact rather than end-user documentation. It can be archived or deleted once its recommendations have been actioned.

- SiteTrax.io Intelligence Documentation Migration Report

SiteTrax.io Intelligence Documentation Migration Report

Historical migration record. This page describes an earlier documentation migration. Use current product pages and dedicated provider sections for setup and availability.

> **Historical migration record.** This report records work and observations from 4 September 2026. Its capability, availability and issue descriptions are historical, and the recommendations below are not a current backlog. For current customer guidance, use the SiteTrax.io Intelligence shelf, Account Access and OAuth and Known Issues, reviewed on 8 September 2026.



Date: 4 September 2026
Application version audited: SiteTrax.io 3.0.0 at app.sitetrax.io
Scope: Create the SiteTrax.io Intelligence documentation set and migrate references away from the legacy service.sitetrax.io experience.

Summary

The SiteTrax.io Intelligence shelf existed but held no books. It now holds eight books and 35 pages. Thirty-two pages were written from scratch against the live application, two existing MCP pages moved onto the shelf with their content preserved, and the legacy service portal content was

marked retired rather than deleted so that no existing URL breaks.

The live application was treated as the source of truth throughout. Where documentation and application disagreed, the application won, and each disagreement is recorded below.

Pages created

Book	Pages
Introduction to SiteTrax.io Intelligence	5 — What is SiteTrax.io Intelligence?, Why Intelligence Matters, How Intelligence Uses SiteTrax.io Data, Core Intelligence Features, Intelligence Roadmap Overview
Getting Started	5 — Accessing SiteTrax.io Intelligence, User Requirements, First Login Experience, Navigating the Intelligence Interface, Common Workflows
AI Integrations	6 pages: two assistant connection guides, Understanding MCP, Using MCP with SiteTrax.io, Supported AI Platforms and Troubleshooting Connections.
Operational Use Cases	8 — Container Visibility, Yard Operations, Gate Operations, Detention Analysis, Operational Investigations, Asset Tracking, Customer Service Workflows, Executive Reporting
Administration	5 — User Management, Permissions, Organization Configuration, Enabling Intelligence, Best Practices
FAQ and Troubleshooting	3 — Frequently Asked Questions, Known Issues, Troubleshooting Guide
Documentation Migration	1 — this report

Every page carries a branded feature image in the Intelligence Indigo palette.

Pages moved and rewritten

Two empty assistant-connection placeholder pages were moved into AI Integrations and written using the providers’ published documentation. Current setup and availability now live in each dedicated provider section.

Books moved

The MCP Server book moved to the Intelligence shelf without changing its page URLs. Preserve existing public URLs used by external connector submissions. Provider-specific guides belong in

dedicated provider sections.

Pages updated

Page	Book	Change
SiteTrax.io API - Output (JSON)	SiteTrax.io API	Legacy hostname and portal terminology corrected
SiteTrax.io Capture - Android	SiteTrax.io Mobile	Legacy hostname and portal terminology corrected
Camera Installation Requirements and Guidelines	SiteTrax.io Gate	Legacy hostname and portal terminology corrected
Bring Your Own Camera (BYOC)	SiteTrax.io Gate	Legacy hostname and portal terminology corrected
Tutorial and Introduction to Snap	SiteTrax.io Snap	Legacy hostname and portal terminology corrected
Spreadsheet-to-map integration guide	Dedicated mapping-provider section	Legacy hostname and portal terminology corrected

All `service.sitetrax.io` links in these six pages now point at `app.sitetrax.io`, and references to the service portal now read SiteTrax.io Intelligence.

Pages retired

All 22 pages of the **SiteTrax.io Service Portal** book now carry a retirement notice at the top directing readers to SiteTrax.io Intelligence. The book and shelf descriptions are prefixed **RETIRED**.

Nothing was deleted and nothing was renamed. Renaming the book or shelf would regenerate their slugs and break all 22 page URLs, so the names were left intact and the retirement signalled in the descriptions and on every page. Their original text is preserved as the historical record of the retired portal.

Screenshots

Forty-six images were captured and uploaded to docs.sitetrax.io: 33 branded feature images and 13 screenshots of the live 3.0.0 application.

Screenshot	Shows
Intelligence hub	Profile status, coverage cards, Intelligence in the navigation

Screenshot	Shows
Assistant connections	Assistant connection cards
Dashboard	3.0.0 layout with period, project and totals
Assets gallery and asset detail	Readings, exceptions filter, record with image, payload and feedback controls
Videos list	Recordings with asset counts
Project users and integrations	Access management and outbound destinations
Interview steps 1, 2 and 7, and the published profile	Onboarding, adaptive detention inputs, published revision

The project users screenshot has real addresses replaced with example.com placeholders. No legacy 2.6.1 screenshots were reused in the new set, and the old screenshots remain only inside the retired book where they are correct as history.

Conflicts found between documentation and application

Finding	Evidence	Resolution
At migration time, an enterprise connector guide and the application disagreed on availability. Consult the dedicated provider section for current status.	At migration time, an enterprise connector guide and the application disagreed on availability. Consult the dedicated provider section for current status.	An enterprise connector guide supported an external submission. Current release status and setup limits must be stated in its dedicated provider section.
A Reports section appears in legacy screenshots	Absent from 3.0.0 navigation; `/dashboard/reports` returns not found	Recorded in Known Issues so nobody hunts for it
Legacy screenshots show version 2.6.1	Current application is 3.0.0 with different navigation and dashboard	Fresh screenshots throughout the new set
One legacy screenshot was reused as the feature image on four different pages	Pages 31, 32, 33 and 35 of the retired book shared one image	Not corrected. Those pages are retired; the new set has a distinct image per page

Product defects found while auditing

Defect	How to reproduce	Why it matters

Unhandled error on an unknown asset ID	Open `app.sitetrax.io/dashboard/assets/view/99999999`. The application shows *Unexpected Application Error! Cannot read properties of undefined (reading id)* instead of a not-found page. A valid ID loads correctly with or without a trailing slash.	The MCP link contract tells assistants to return links in exactly this form. Any assistant answer referencing a record the user cannot access lands on a crash screen, which reads as an outage rather than a permissions result.
Intelligence profile does not follow across accounts	A profile published in the browser session was not visible to the MCP connection, which was authenticated as a different account	Correct behaviour, since profiles are per user, but it produces generic answers with no explanation. Documented in Troubleshooting Connections and Known Issues. Consider surfacing the signed-in account in assistant responses.

Quality assurance performed

- All 13 internal links across the shelf resolve.
- All 45 image references across the shelf resolve.
- All 32 new pages lead with a feature image. Both MCP Server pages carry the existing 440x250 branded images in their established position.
- No duplicate page titles were introduced.
- The only remaining `service.sitetrax.io` mentions on the shelf are five deliberate references, in code formatting, explaining that the hostname is retired. None are live links.
- Book and page ordering was checked against the migration specification.

Remaining gaps

Gap	Why	What would close it
No screenshot of the first-run welcome panel	It only appears for an account with no published profile, and the audit account now has one	Capture on a fresh account before its interview is completed
Dashboard projects summary table not documented	The audit account has a single project, so a multi-project dashboard could not be observed	Re-audit the dashboard on an account with several projects
Organization-level administration is thin	No organization settings or roles console was reachable from the audit account. Administration is documented as project-scoped because that is what exists at this access level	Re-audit with an administrator account and extend the Permissions page

Gap	Why	What would close it
Roadmap page is deliberately minimal	Only the explicitly documented provider connection was described as in progress. Other availability was not inferred.	Product input on what may be published
Assistant-side screenshots	External assistant setup screens belong to their providers, change frequently and may expose an individual account.	Optional. The written steps cite vendor documentation, which ages better than vendor screenshots

Recommendations

1. **Fix the unknown-asset-ID error** before promoting assistant deep links widely. A not-found page stating the record does not exist or is not shared with you would remove the sharpest edge in the assistant experience.
2. An enterprise connector guide supported an external submission. Current release status and setup limits must be stated in its dedicated provider section.
3. **Delete the retired Service Portal shelf and book** once you are satisfied nothing external links to it. Check inbound links first; the banners can carry it safely for a quarter in the meantime.
4. **Tidy three empty pages titled New Page** in the SiteTrax.io Gate book, and remove the file `test-upload.png` from the image gallery, which was created while verifying the upload path.
5. **Replace the demo Intelligence profile** on the audit account. It is published as revision 1 and labelled as a documentation demo in its custom instructions, and an edit draft is open against it. Republish with real answers or discard, whichever suits.
6. **Adopt one feature-image position.** The new pages lead with the image; the two MCP pages place it after the opening line. Either is fine, but pick one.
7. **Set a re-audit trigger on the application version.** This set was written against 3.0.0. When the version in the sidebar changes, re-check the navigation, the Intelligence hub and the assistant cards.

Suggested future documentation

Proposed	Why
Alerts and Scheduled Digests	The only capability requiring a write scope, and currently undocumented for end users
Reference: status codes, asset types and capture types	Status codes such as A0 and capture types such as International Intermodal Vertical ID appear throughout the interface with no glossary. This is the single most requested kind of page in operational products.

Proposed	Why
Capture Coverage Planning	The coverage review tells operators to confirm both gate directions and map yard zones, but nothing explains how to do that well
Security and Data Handling	Consolidates OAuth, scopes, redaction, retention and what does and does not leave SiteTrax.io. Procurement asks for this.
Release Notes	Would have made the 2.6.1 to 3.0.0 documentation drift visible as it happened rather than at audit time
Interpreting SiteTrax.io Data	Record semantics are currently repeated as cautions across many pages. One canonical page would reduce duplication and drift.

> This report is a working record, not end-user documentation. Archive or delete it once the recommendations have been actioned.

Phase 2 — site-wide consistency and answer-engine readiness

Date: 4 September 2026. A second pass covering the whole documentation site, not only the Intelligence shelf.

Consistency fixes applied

Finding	Action
The SiteTrax.io Intelligence shelf had no description and no cover, against 650 to 870 characters and a cover on every other shelf	Branded cover created and a 799-character description written
The six new books carried 79 to 147 character descriptions against a house range of 279 to 1,273	All six rewritten to 442 to 615 characters, each opening with a definitional sentence
Two API pages contained stray in-body H1 headings competing with the page title, one of them three times	Demoted to H2. The site now has zero in-body H1 headings
Seven older pages had no feature image against 33 new pages that all did	Assistant connection cards

Answer-engine and generative-engine readiness

The site was audited for how well answer engines and large language models can read and cite it. It is publicly readable and crawlable, which is the precondition for everything else. Four gaps were

found and closed.

Gap	Fix
No meta description on any page. Only an Open Graph description was emitted, and on pages that open with a heading it concatenated words without spaces	A clean description is now built from the first substantial paragraphs and applied to both the meta description and the Open Graph description
No canonical URL	Added per page
No social or preview image, and no Twitter card	The page feature image is now used for both, which is why the feature-image backfill above mattered beyond appearance
No structured data anywhere on the site	Schema.org markup added: a site-level Organization, WebSite and SoftwareApplication graph, plus per-page BreadcrumbList and TechArticle, and FAQPage on qualifying pages. The FAQ page emits 19 question and answer pairs

All of this is delivered through Settings, Customization, Custom HTML Head. The site-level entity graph is static and therefore visible to every crawler. The per-page tags are built in the browser after the page loads.

Known limitation at the time of this migration: per-page metadata was injected with JavaScript. Crawlers that do not render JavaScript see only the static site-level graph. The durable improvement is server-rendered metadata from the page template or proxy.

Still outstanding

- **No sitemap.** `docs.sitetrax.io/sitemap.xml` returns not found, and BookStack does not generate one. This needs server-side work or a generator, and it is the highest-value remaining item for discovery.
- **Server-side rendering of the metadata** described in the limitation above.
- **Three empty pages titled New Page** remain in the SiteTrax.io Gate book. They are the only pages on the site without a feature image, and they should be filled or deleted.
- **A stray** `test-upload.png` remains in the image gallery from verifying the upload path.