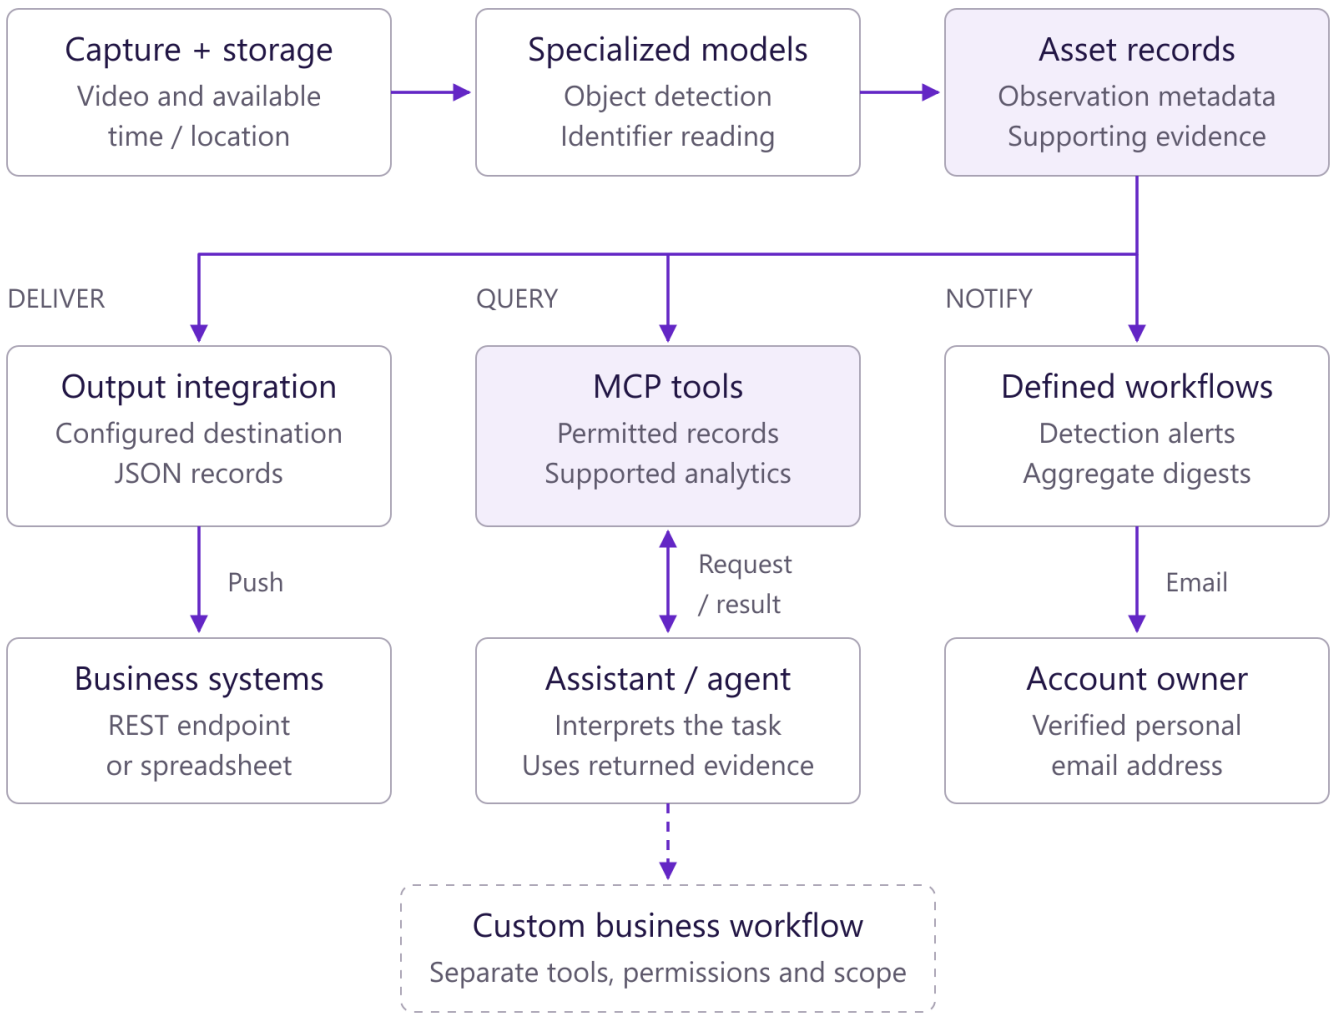
How Intelligence Uses SiteTrax.io Data

Intelligence reads the records SiteTrax.io already stores and returns operational results scoped to the projects your account is permitted to see. The native hub uses your published profile for personal guidance. A connected assistant can retrieve selected published facts and priorities when relevant to a request.

System architecture

SiteTrax.io Intelligence architecture

Logical data flow. Hosting and storage ownership depend on the deployment.

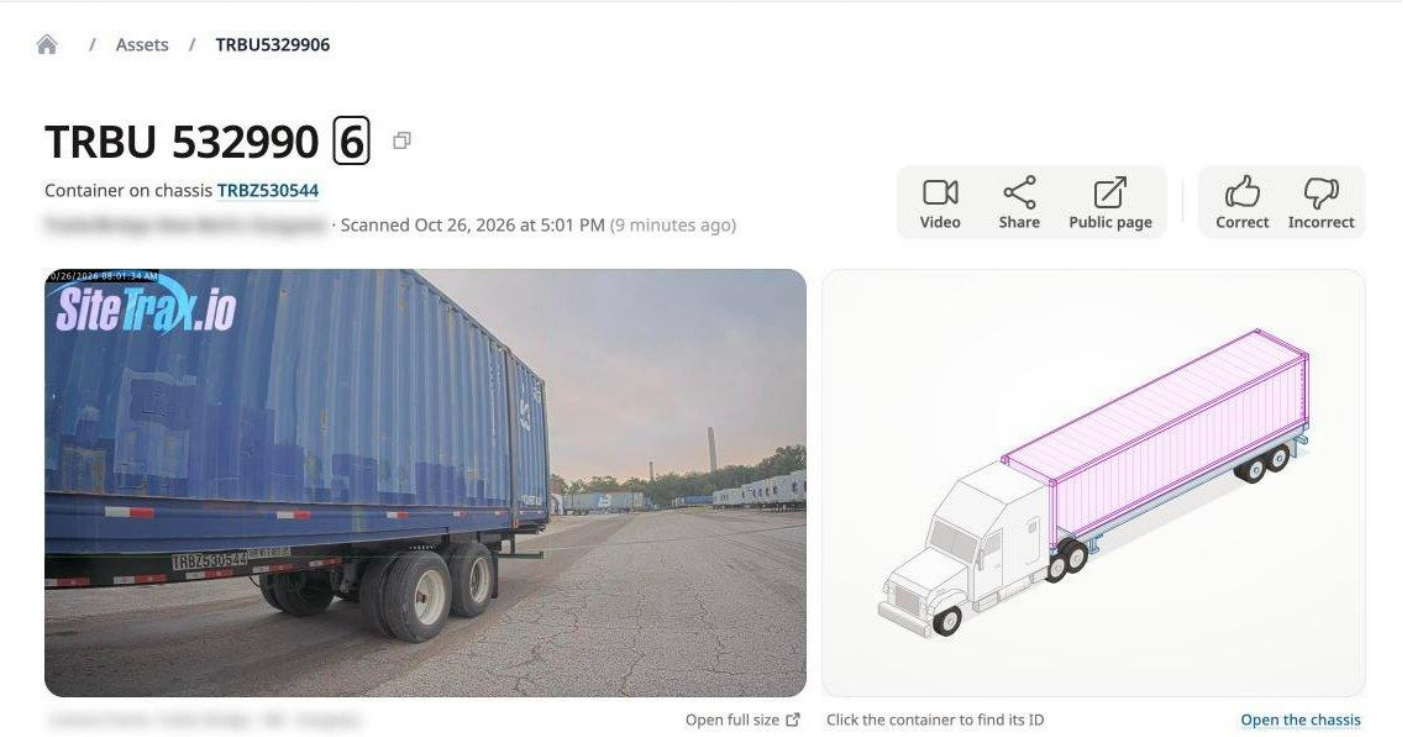


Open the architecture diagram at full size. Solid arrows show documented paths; dashed boxes and arrows show separately scoped workflows. These are logical responsibilities, not a claim about physical hosting.

SiteTrax.io's specialized models extract identifiers from captured imagery. A connected assistant's foundation model interprets the user's task, requests permitted SiteTrax.io results and explains the evidence. Broader actions require their own workflow, tools and permissions.

The chain from camera to answer

1. **Capture.** A camera or capture app records video at a gate, in a yard, or on a route.
2. **Processing.** SiteTrax.io processes the capture and stores asset/result rows with identifier text, emitted type and status code. Some rows represent failed reads or a synthetic missing side of a container/chassis pair, rather than an independently identified unit.
3. **The record.** A record contains stored identifier text, type, project and processing status, with scan time, camera, heading, coordinates, confidence and media links where available. Direction requires the relevant heading or site designation, and retention or redaction can withhold media.
4. **Intelligence.** Aggregations, sequences and exception checks are calculated on the server. An assistant can request published profile context separately when it helps answer your question.
5. **Delivery.** The result is shown in the application, or returned through the MCP server to the assistant provider you connected. Query-time access avoids a separate bulk synchronized index; it still sends the requested results to that provider.



How to read an asset record

For a full equipment history, confirm both the complete identifier and the equipment type. Text search can find partial matches, and the same text can occur in different equipment families.

Reading these records correctly matters more than any single feature, because a misread record produces a confident wrong answer.

- **One record is one stored processing result.** It is not guaranteed to be one physical asset or one direct observation. Where container and chassis processing are combined, SiteTrax.io can add a synthetic row for the paired side.
- **Identifier text is only usable when it is real text.** Blank values and placeholders such as a single dash are not equipment identifiers and must not be treated as one.
- **Status code and asset type must be read together.** A processing status describes the outcome of reading the image. It is not by itself a statement about the physical condition of the unit.
- **Counts mean what they say.** A count of records is not the same as a count of distinct units, and a count of distinct identifier values is not a guarantee of distinct equipment.
- **Direction is configuration-dependent.** It can describe movement derived from a record's heading and stored direction mapping, or the designated role of a project or camera. Ask which interpretation was used. Vertical or missing headings do not establish heading-derived entry or exit.

Common status interpretations

Read status and emitted asset type together. These examples describe processing outcomes and do not independently establish an operational incident.

Code	What it means
A0	An accepted read under the processing rules; not independently verified accuracy.
I1 / I2	A container owner-prefix warning / a failed container check-digit check.
I3	No usable identifier was read. This does not prove physical damage.
I4	A readable chassis or USDOT identifier was absent from the supplied asset list. The identifier is not necessarily invalid.
I5	A synthetic missing container or chassis side of a combined pair, rather than a separately observed unit.
I6	A low-confidence result that otherwise met the acceptance rules.

Code	What it means
I7	A configured text substitution changed the identifier after processing. Legacy display text may misleadingly describe it as low confidence.
-	Generic OCR has no semantic status classification; the dash is not automatically a warning.

The share of rows carrying A0 is a processing-result measure. It should not be labeled OCR accuracy without a defined eligible population and independently verified outcomes.

Dates and evidence

Asset analysis uses scan time when present and falls back to record creation time for older records. Video analysis uses video creation time. Ordinary MCP date filters use UTC-midnight business-day boundaries; a grouping timezone changes the displayed time buckets. Sequence analysis can instead use exact timezone-aware start and end timestamps for local calendar periods. State the requested timezone and check the resolved boundaries in the answer.

Your profile as context

When you publish the profile interview, SiteTrax.io stores an immutable revision and generates a working brief for review in the application. The public MCP profile tool returns selected facts and priorities from your latest published revision. Available sections include organization, operations, inferred user timezone, confirmed tailored answers and the accessible projects from the publication snapshot. It does not return the working brief, custom AI instructions or response-style preferences.

Custom AI instructions and response preferences can be saved in your profile, but saving them does not configure a connected assistant. Ask for the desired tone or level of explanation in that assistant. Profile retrieval and use depend on the assistant and the question. Your published answers remain customer-supplied context, not independently verified operational or contractual evidence. A saved free-time allowance or rate alone does not establish a payable charge.

Publishing also snapshots the projects covered by that revision. Profile retrieval keeps only snapshotted projects you can still access; newly shared projects are not automatically added to the old publication. Reopen and republish when the intended profile scope changes. If no projects remain accessible, or the returned profile list is incomplete, the assistant must make that limitation clear rather than silently querying all shared projects.

Scope and permissions

For an operational question, explicitly name the projects you want analyzed. A profile provides context, while each data request has its own project selection. Omitted project filters normally mean all projects currently shared with the authenticated account.

Every operational request is limited to the projects currently shared with the authenticated account. Connecting an assistant does not widen SiteTrax.io permissions. Projects selected in a profile describe context; they are not a separate OAuth project grant. Current project sharing remains the access limit. See Account access and OAuth.

> The profile belongs to a specific user account. If you publish a profile in the application but connect an assistant using a different SiteTrax.io account, that assistant will not see the profile. Use the same account on both sides.

Revision #6
Created 2026-09-04 15:13:41 UTC by SiteTrax.io Admin
Updated 2026-10-26 12:28:28 UTC by SiteTrax.io Admin