

AI Search: Frequently Asked Questions

Setup and Prerequisites

Do I need Edge AI to use AI Search? No.

AI Search is a component of the Edge AI platform, but it does not require an active Edge AI agent or chat configuration to be used. AI Search runs independently once enabled.

Is AI Search available on-premises? Yes.

AI Search supports full on-premises deployment using Docker containers, in addition to SaaS. See [Installation of AI Search for OnPrem](#) for the full procedure.

How do I enable AI Search for a SaaS Innovator instance?

For Aras-hosted Aras Innovator SaaS instances, enabling AI Search requires adding two files to your Work repository and running your deployment pipeline. See [Installation of AI Search for SaaS](#) for the full procedure.

Which Aras Innovator versions are supported?

Aras Innovator 14.0.15 or later is required.

What SQL Server version is required for on-premises deployment?

SQL Server 2025 or later is required for on-premises deployment. AI Search uses the native vector type introduced in SQL Server 2025. Azure SQL Database with vector support is also supported. If your organization is on SQL Server 2019 or 2022, contact Aras for the compatibility upgrade path before starting installation.

What other infrastructure is required before installing AI Search on-premises?

The following must be in place before deploying AI Search on-premises: the Aras InnovatorEdge CDC pipeline (edgeapi-eventing-configurator, edgeapi-eventing-exporter, and a message broker, either RabbitMQ or Azure Service Bus); a Temporal cluster (version 1.24 or later) reachable from the Docker host; and an OpenAI-compatible embeddings API endpoint. AI Search auto-provisions its own durable queues and bindings on the message broker when CDC configuration is applied; no



manual queue setup is needed on the RabbitMQ or Azure Service Bus side. See Installation of AI Search for the complete prerequisites list.

Configuration

Can I add a Relationship ItemType to the indexing scope? Yes.

Relationship ItemTypes, such as a BOM relationship, can be added directly to the indexing scope in the AI Search Configurator. They can also be used as a source of file properties under a regular ItemType in the File Properties section.

What properties can I select for indexing?

The property selector shows a consistent set of Aras Innovator properties regardless of which ItemType is selected. ItemType-specific custom properties, such as `item_number`, `description`, and `classification` on Part, do not appear in the selector. Properties available through the selector include system properties common to all items, plus relationship-level properties for Relationship ItemTypes such as `quantity`, `source_id`, `related_id`, and `sort_order`.

To index a custom property that does not appear in the selector, add it by specifying its internal property name directly in the JSON schema and importing it using Import Schema JSON in the configurator. The indexing engine reads properties from the actual Aras Innovator table for each ItemType, so any valid property name on that table can be indexed.

What is the File Properties section?

File Properties is a section within each ItemType row in the configurator, separate from File Handling Settings. While File Handling Settings selects which file types are included, File Properties controls which file-related fields are used for text extraction, vectorization, and search. Each file property row has an enable/disable switch. Enable at least one file property when file content indexing is required for an ItemType that has file handling configured.

What file types can I include in file handling?

The file type selector shows only formats that the AI Search file parser supports. Supported formats are: PDF, DOCX, PPTX, HTML, Markdown, and image formats (PNG, JPG, GIF, TIFF, BMP, WEBP). The selector is sourced from your Aras Innovator instance's FileType table, so only registered file types that also appear in the supported format list are shown. The list is searchable; type any part of



a format name to filter it. If an expected file type is missing, check whether two types in your database share the same extension, as only the first is shown. The list is cached for one hour when the configurator opens; reopening the configurator refreshes it.

What happens if I change the indexing scope after the instance is already indexed?

A re-indexing warning is shown before the change is applied. In version 1.1.0, applying a configuration change re-indexes only the ItemTypes and content scopes affected by the difference between the saved and new configurations. ItemTypes that were not changed remain fully searchable while the apply runs. Apply returns immediately and completes in the background; progress is visible in the Indexing Status section of the configurator.

Why does Apply fail with "At least one indexable property is required"?

Every ItemType in the indexing scope needs at least one configured property, file-indexable property, or file relationship. Expand the ItemType row in the configurator and add at least one property, enable file handling with file types, and/or enable file properties in the File Properties section before applying.

Can I import and export my configuration? Yes.

The AI Search Configurator supports exporting the current configuration as a JSON schema file and importing a JSON schema file to load it into the editor. The imported configuration is not saved until you select Apply. Use Verify to validate a configuration before applying it.

What does the Failed Items dialog show?

When an indexer row shows a Failed or Partially Completed status, a Failed chip appears on that row in the Indexing overview table. Select the chip to open the Failed Items dialog, which lists each failed item or file by name and failure reason. From the dialog, you can select Retry to re-queue a failed item for indexing immediately, or Resolve to dismiss it without indexing. Both actions support bulk selection: select multiple rows and then select Retry N selected or Resolve N selected. Resolve requires confirmation. Select Close to dismiss the dialog.

What does Apply and Reindex All do?

Apply and Reindex All is the secondary action on the Apply split button in the AI Search Configurator. Select the dropdown arrow to the right of Apply to access it. This action saves the current



configuration and deletes all item embeddings before rebuilding them from source. File embeddings are not deleted or reprocessed by this operation. Use Apply and Reindex All when item embeddings need to be rebuilt from scratch, for example after an embedding model change. For routine configuration changes, the standard Apply handles re-indexing automatically.

Does the initial indexer table in Indexing Status always show? No.

In version 1.1.0, the initial indexer table is hidden by default. Select Show Initial Indexer details to expand it. Select Hide Initial Indexer details to collapse it. The table organizes indexers into two sections: ITEM INDEXERS (one row per ItemType) and FILE INDEXER (a single row for all file indexing activity).

Data Access and Security

Can a user see AI Search results they are not authorized to see in Aras Innovator? No.

AI Search results are filtered by the requesting user's existing Aras Innovator permissions. AI Search does not grant any new access beyond what a user already has.

Is my data used to train AI models? No.

Content indexed by AI Search and queries submitted to it are not used to train, retrain, or improve any general purpose or shared AI model. See the Aras InnovatorEdge Product Specific Terms for full details.

Does AI Search send my data outside Aras Innovator?

During indexing, the AI Search service reads in-scope item and file content from the Aras Innovator database and sends it to an OpenAI-compatible embeddings API to generate vector representations. These vector representations are stored in the `[ais]` schema within the Aras Innovator database; no separate database is created. The embeddings endpoint is configurable: on-premises customers can use an Aras-provided hosted embeddings service by obtaining the endpoint and credentials from Aras, or use any compatible provider. The endpoint is set via the `EMBEDDINGS_ENDPOINT` environment variable in the AI Search configuration.

Using AI Search

Where do I access AI Search?



Type a query into the search field in the Aras Innovator header to open the AI Search tab within the Aras Innovator interface. Administrators additionally see a gear icon in the top-right corner of the AI Search tab that opens the Admin panel.

What is the "Ready to Search" state?

When you first open the AI Search tab without having run a search yet, the results area shows a Ready to Search state with example query chips such as "123 module description," "Body parts with Preliminary state," "Cylinder head manual," and "Latest modified CAD documents." Select any chip to run that search immediately.

Is there a minimum query length? Yes.

Queries must be at least 3 characters long. Submitting a shorter query returns the message: "Search query too short — Please enter at least 3 characters to search. Try a longer keyword or phrase."

What filters are available?

Six filter properties are available: Item Type, Is Current, State, Modified After, Modified Before, and Is Released.

Item Type filtering supports selecting more than one type in a single search. State filtering uses a text-contains match; typing part of a state name is enough and exact text is not required. Modified After and Modified Before accept a date and cannot be set so that Modified Before falls earlier than Modified After when both are active. Is Released and Is Current accept True or False.

Filters follow a staged and applied model. Adding, editing, or removing a filter stages the change without immediately updating results. Staged filters appear as gray chips. Select Apply Filters to send the staged filters to the search service; applied filters appear as blue chips. Entering a new search query or pressing Enter also applies any staged filters.

Can I use more than one filter of the same type? No.

Each filter property can be used once at a time. Once an Item Type filter is active, that property is not available in the Add Filter dropdown until the existing filter is removed.

What does the Load More button do?



When more results are available beyond the first page, a Load More button appears at the bottom of the result list. Selecting it fetches the next page using the same query and applied filters. Load More does not apply staged filter changes.

Why am I getting no results at all?

This can happen for several reasons. If no ItemTypes have been configured for indexing yet, or the index is temporarily unavailable, you may see "The Index Is Not Ready Yet or Is Currently Unavailable" — wait and try again, or contact your Aras Innovator administrator. If AI Search is running normally but your query returned nothing, you will see "No Matches Found — We couldn't find anything matching your query." Rephrase using different or broader terms. If you see "Search query too short," enter a query of at least 3 characters.

What does "Search service is temporarily unavailable" mean?

This message indicates a transient service error distinct from the index-not-ready state. Select Search again to retry.

What does "Invalid search syntax" mean?

The query contains invalid characters or syntax. Remove or correct special characters and try again.

Why don't older versions of an item show up in search?

AI Search only indexes the current version of an item and the current version of any indexed file attached to it. Earlier versions are not included in the index.

Can I see which part of a file matched my search query? Yes.

Expanding a result card shows the text excerpts from indexed files that made the item relevant, grouped by source and annotated with the file name. When a search matches content at the file chunk level, the chunk appears as its own result entry with a reference link to the source file.

How do I inspect item details without leaving the search results?

Select the detail control on a result card to open the Result Detail Drawer. The drawer shows the item header, type, keyed name, and identifier, along with matched excerpts grouped by file source. The results list stays visible behind the drawer. Close the drawer with the Escape key or the close button; the scroll position in the results list is preserved.



My search showed a message asking me to refine my query. What does that mean?

This message appears when a search has been running for 30 seconds without returning a complete result set. Adding filters or using more specific search terms reduces the scope the query must cover. See the Search Filters section of the AI Search User Guide for guidance.

The Is Current filter does not seem to affect my results. Is this expected? Yes.

The Is Current filter is accepted by the UI and can be staged and applied, but it does not currently restrict search results to current or non-current items. This is a known behavior in version 1.1.0. Do not rely on this filter to limit results by version status.

