

Edge AI Release Notes

This page includes all Edge AI Release Notes starting with the most recent.

2026 October 2026 - Edge AI - Incremental Release - Edge AI version 2.2

Aras InnovatorEdge AI 2.2 Is Now Live!

October 13, 2026 - Incremental Release

InnovatorEdge AI 2.2 introduces a redesigned home page, administrator controls for model availability, and a Beta tool for reviewing an agent's data access coverage. Answer accuracy for data retrieval by relationship, history, and aggregation queries has also been improved.

Home Page

The entry point at edge.aras.cloud/ai now opens a structured home page rather than a plain agent list.

The page opens with a greeting that addresses the signed-in user by name and shows the name of their tenant.

A **Start a conversation** section lists the agents available to the signed-in user. Choosing one opens a conversation directly.

Administrators see a **More for you** section above the agent list with a **Manage agents** card.

Administrators can open the Agent Library by selecting **Go to Agent Library**.

Administration



Model Enable and Disable

The Models screen shows an enabled or disabled state for each model configuration. A disabled model is removed from the agent editor's selection list and cannot be assigned to any agent.

Disabling a model that is already assigned to one or more agents produces a warning before the change takes effect. The platform moves each affected agent to Draft and clears its model assignment. Selecting a different model and republishing restores the agent's availability to end users.

Data Access (Beta)

The agent Overview page includes a Data access row in its Configuration section. Opening it shows the schema of the Edge API connections the agent's tools are bound to: item types as nodes, relationships as connecting lines. Selecting a node lists the properties and access levels available to this agent: read, list, add, edit, and delete.

A tabular view of the same information is also available.

Data access is active only when the agent has at least one enabled tool binding with a configured Edge API connection.

The schema shown reflects the Aras Innovator access of the credential used to connect: the stored API key for API Key connections, or your own account for Bearer Passthrough connections. It may not match what end users of this agent can access.

The schema graph is also reachable from the Access Points screen. Select **View Schema** from the actions menu on any connection to open the graph directly.

Bug Fixes



This release includes bug fixes and stability improvements that contribute to a more reliable experience.

Availability

InnovatorEdge AI 2.2 is available as a cloud service hosted in the United States, Europe, and Japan.

2026 September 29 - Edge AI - Incremental Release - Edge AI version 2.1

Aras InnovatorEdge AI 2.1 Is Now Live!

September 29, 2026 - Incremental Release

InnovatorEdge AI 2.1 improves the reliability of answers the Aras AI Assistant returns for where-used, change impact, and data retrieval queries, and resolves issues identified since the 2.0 release. Administrators gain additional controls over the agent lifecycle.

Improved Answer Accuracy

The Aras AI Assistant returns more complete and correct answers.

Where-Used and Change Impact

Where-used queries now return results across both BOM relationships and change orders in a single response. Previously, a question about where a part is used returned BOM context only and required follow-up questions to reach change order relationships. A question such as "What is affected by ECO-00001012?" and a follow-up such as "Where is MP2295 used?" each now return the complete picture in one answer.



Result Counts

When the assistant groups results by attribute, the count shown for each group reflects the actual number of items in that group.

Answer accuracy has been improved across data retrieval by relationship, data retrieval by property, BOM tier queries, full history retrieval, and generation summary scenarios for the Aras AI Assistant. These improvements apply when the Aras AI Assistant is configured with the default Aras-managed model.

Administration

Connecting to the API Management Service

When no Edge API connections have been registered in a tenant, the Access Points screen now includes a direct link to the API Management Service. Administrators can open the API Management Service from that link, configure an API, and return to register it as an Access Point without needing to navigate there separately. The link remains available after connections exist.

Confirmation Before Publish State Changes

Publish, move to Draft, and saving edits to a published agent each now show a confirmation dialog before the agent's availability to end users changes. Each dialog names the agent and states which users will gain or lose access before the change takes effect. Canceling any dialog leaves the agent's current publish state unchanged; for save actions, the configuration form stays open with edits intact.

Agent Preview Before Publishing



Administrators can open a preview of an agent and converse with it before publishing. The preview runs against the agent's last-saved configuration. Unsaved edits in the configuration form are not included; a save is required first. This allows behavioral testing against live data without making the agent available to end users. (IH-23475)

Bug Fixes

Agent with Blank Instructions

An agent whose platform instructions and solution instructions fields were both cleared could fail to run, producing a misleading internal error. An agent with blank instruction fields now runs correctly. (IH-23541)

Chat Panel Layout

The chat panel broke the chat interface layout in certain conditions. The layout is restored. (IH-23183)

URL Redirect

Navigating to the application URL without a trailing slash returned a 404 error. The URL now redirects correctly with or without a trailing slash. (IH-23506)

Availability

InnovatorEdge AI 2.1 is available as a cloud service hosted in the United States, Europe, and Japan.

2026 September 15 - Edge AI - Release Announcement - Edge AI version 2.0



Aras InnovatorEdge AI Foundational Release 2.0 Is Now Available

September 16, 2026 - Foundational Release

Every engineering team carries a standing backlog of questions its PLM system can already answer. Which parts a change affects. Who owns a document. How many Engineering Change Orders (ECOs) a team is still holding open.

The information is in Aras Innovator, but reaching it has meant knowing the schema, knowing where to look, or asking someone who does.

InnovatorEdge AI 2.0 removes that step, and then goes further.

The assistant no longer only reads Innovator data. It can act on it, under mandatory human approval, inside the workflow where the work is already happening.

- **Conversational PLM intelligence.** Answers from Innovator data in plain language, grounded in live records and scoped to the user's own permissions, with a link to the source record every time.
- **Governed action, not just answers.** The assistant can create records, update properties, change lifecycle state, and delete items in Innovator. Every change requires explicit human approval before it is committed. There is no automatic-commit setting.
- **AI where you work.** Embedded directly in the Innovator interface, aware of the record the user has open. A standalone application is available for work outside Innovator.
- **Enterprise administration.** Two consolidated roles under a single Centralized Identity Access Management (CIAM) application, a governed default assistant in every tenant, Skills and MCP Servers in the Agent Management Platform, and two agent reasoning types.

Conversational PLM Intelligence

Ask the question in plain language. Get the answer from your own data, with a reference to the record it came from.

How It Works

Answers are drawn directly from Innovator through Edge API, executed under the signed-in user's



identity, so results are scoped to what that user is permitted to see. When a question is ambiguous, the assistant asks a short clarifying question before running anything. Follow-up questions refine the previous answer without restating context, so a multi-step investigation flows as a conversation rather than a series of new searches.

What It Covers

- Listing and filtering by state, date, owner, or attribute
- Bill of Materials (BOM) structure and where-used queries
- Change analysis across ECOs and change processes
- Version and generation history
- Counts and aggregations by type or classification
- Relationship and data model diagrams, and pie charts, drawn from the result set already returned

Transparency

Reasoning steps are visible inline with each response, so a user can see how an answer was reached and which records it came from, not only what the answer is. Administrators can configure conversation starters to guide first-time users toward the queries that matter most in their organization.

Governed Action: Approval Before Any Change

Version 2.0 is the release where the assistant stops being read-only. It can create records, update properties, change lifecycle state, and delete items in Aras Innovator.

It cannot do any of that on its own.

Human approval is a platform invariant, not a configurable option. The agent reasons and proposes freely, then stops. It presents an approval card inline in the conversation showing the intended action, the target item, and the data payload, with approve and reject controls. No separate screen, no approval queue, no second application. Rejecting makes no change and ends the action cleanly. Approval authorizes the attempt, not the outcome. The write executes through Edge API under the user's own identity, so Innovator permissions, lifecycle rules, and workflow constraints all still apply. If Innovator refuses the operation, nothing is written and the assistant reports why. Changes committed through the assistant appear in Innovator as changes made by that user, not by a service account.



This is what makes agentic PLM deployable in a governed engineering environment rather than only in a demonstration.

AI Where You Work

Instead of switching applications to find an answer, ask the question in a panel that sits inside Innovator, next to the work.

Embedded in Innovator

The assistant runs in a resizable side panel within Innovator, detects which record the user has open, and uses that record as context automatically.

An engineer reviewing a change order who needs to know who submitted it is one question away from the answer.

Context follows the user as they navigate between records, and can be cleared when the next question is unrelated.

Navigation

Every answer links back to the Innovator record it drew from. Selecting that link opens the record without ending the conversation. The loop stays continuous: question, answer, source record, next question.

More Room When You Need It

The assistant expands to a full-screen view with the session intact, and a standalone application is available at its own URL for work outside Innovator, with the same capabilities as the embedded panel.

Working in Your Own Language

Questions can be asked in any language the underlying model supports, and the assistant answers in the language the question was asked in by default. Administrators can steer response language through agent instructions where a team needs to standardize on one.

Open Tool Interoperability

Version 2.0 opens the assistant beyond Innovator data. Edge AI acts as an MCP client, so an agent can be connected to an external MCP server and can discover and call the tools that server exposes,



bringing adjacent systems into the same conversation as PLM data.

Governance travels with it. Tools are granted per agent by an Edge AI Administrator, an agent can only call the tools it has been granted, and grant changes take effect on the agent's next run. Agent-to-agent delegation lets an agent route part of a request to a specialist agent, and where that specialist needs approval for a data change, the approval surfaces in the user's own conversation.

What Changed for Administrators

Role-Based Access

Version 2.0 consolidates the previous per-application access model into two defined roles managed through a single CIAM application: Edge AI User and Edge AI Administrator. A user with no assigned Edge AI role cannot reach the assistant or the Agent Management Platform, so role assignment is a required onboarding and upgrade step.

The Aras AI Assistant

A governed, pre-configured agent now ships with every tenant, delivered as a Deep Agent with high reasoning effort and all Aras-managed skills enabled. An administrator connects it to an Access Point, selects the tools it should use, and publishes it. Its name, agent type, and instructions are locked by Aras and it cannot be deleted. Administrators extend it by adding their own skills, MCP tools, and sub-agents, and can assign a different model where required.

The Agent Management Platform

The Agent Management Platform is where administrators manage agents, models, Access Points, and access.

Version 2.0 adds two areas:

- **Skills:** Reusable instruction bundles that teach agents domain-specific query patterns
- **MCP Servers:** The tool registry, where the Aras Innovator tools and any external tool servers an organization registers are managed together

Aras-managed skills, tools, and model configurations are shared into each tenant as read-only and are marked as Aras-managed, so administrators can see what they inherit and what they own.

Agent Types



Administrators can create agents using two reasoning approaches: the Deep Agent, which plans and executes and suits multi-step investigation, and the ReAct Agent, which reasons and acts iteratively.

Bring Your Own Model

Administrators can register external model providers from the Models screen: Azure OpenAI and Microsoft AI Foundry, OpenAI, Anthropic, and Google, along with OpenAI-compatible endpoints such as Ollama or vLLM through a custom endpoint. Inference can therefore run on infrastructure the customer controls.

Licensing

Edge AI User and Edge AI Administrator counts are defined by your organization's Edge AI subscription. See the [Administrator's Guide](#) for role definitions and the InnovatorEdge service terms for licensing details.

Availability

InnovatorEdge AI 2.0 is generally available from 15 September 2026 as a cloud service, hosted in the United States, Europe, and Japan.

Your Aras Innovator instance can remain on-premises or hybrid; the assistant reaches it through the Edge API layer. Edge API is a prerequisite, and the Aras Innovator AI Assistant Package must be installed on the Innovator instance to enable the embedded panel.

2026 August 11 - Edge AI - Incremental Release - Edge AI version 1.1.3

Aras InnovatorEdge AI Version 1.1.3 Is Now Available

August 11, 2026 - Incremental Release

Introducing product structure navigation in chat, configurable BOM depth, item version comparison, semantic search via AI Search, and European region model hosting.

What's new



Explore product structures and where-used results in the chat

Users can now ask the agent about the composition of an assembly or where a specific part is used, directly in chat, without opening the Structure Browser in Aras Innovator.

- Ask the agent to show the structure of an item and receive a formatted list of components with part number, description, quantity, and unit of measure at each level.
- Ask where a part is used and receive a list of parent assemblies, sorted so production-relevant assemblies (Released lifecycle state) appear first.

Results default to the latest released revision. Specifying a revision or lifecycle state in the question scopes the response accordingly.

Why this matters:

Product structure interrogation is one of the most frequent activities for engineers and change managers. Both BOM and where-used lookups are now available through natural language without a context switch to the Innovator Structure Browser.

Example scenarios:

- A user asks "Show me the BOM of Top Assembly X" and receives a formatted component list with quantities. They follow up with "Show where-used for Part 001" and receive a list of the assemblies it belongs to, sorted by lifecycle state.
- A user asks "Which assemblies use component MP0370-001 in Released state?" and receives a filtered where-used list showing only Released parent assemblies.

BOM and where-used depth now configurable up to 10 levels

Agents can now traverse product structures up to 10 levels deep. When asking about an assembly's structure or where a part is used, the agent navigates the hierarchy to the requested depth. If no depth is specified, the agent applies a default. If the user specifies a depth lower than 10, that value is honored.



Why this matters:

Multi-level assembly interrogation is a common engineering task. The depth cap ensures responses remain usable without overloading the conversation context.

Example scenario:

- A user asks "Show me the BOM of assembly X up to five levels deep" and receives a structure list scoped to five levels. A follow-up asking for level six returns results to that depth, up to the maximum of 10.

Retrieve item versions and compare changes

Users can now retrieve the full version history of an item and compare any two versions side by side through the agent, including non-consecutive versions.

Why this matters:

Tracking what changed between versions is a routine PLM task. It is now available through natural language without navigating to the item's version history in Innovator, and without being limited to comparing only adjacent versions.

Example scenario:

- A user asks "Compare Version 1 and Version 5 of part MP0101" and receives a side-by-side summary of the differences between those two versions.

Semantic search via AI Search

Agents can now answer questions through semantic meaning as well as structured property queries. When a question is better served by meaning than by exact property values, the agent queries the AI Search index and returns a ranked set of relevant items.



Why this matters:

Not every question maps neatly to a known item number or property value. Semantic search surfaces relevant items based on meaning, reducing the need to know the exact structure of the data before asking a question.

Example scenario:

- A user asks "Are there any parts associated with the hydraulics rework?" The agent queries the AI Search index semantically and returns a ranked set of relevant items, then retrieves authoritative detail for the top results.

European region hosting for Aras-provided models

Tenants with regional hosting requirements can now be assigned an Aras-provided model hosted in the European region. Regional model assignment is managed by Aras Customer Success and does not require action from tenant administrators.

Why this matters:

Provides a path to regional data residency compliance for European customers who require LLM inference to remain within the European region.

Example scenario:

- A European tenant's Aras-provided model is assigned to the European region by Aras Customer Success. Administrators see the regionally assigned model in their Models list and can assign it to agents as they would any other Aras-provided model.

Bug fixes

Query count accuracy improved (APR-531, IH-20742)

Count queries now return accurate totals. Previously, certain query patterns returned counts that did not respect the applied filters. Additionally, the Edge API previously returned a count value even when the query filter referenced a non-existent property; invalid filter queries now correctly return an error.

User identity resolution improved (IH-21062)

Queries that filter by owner, creator, or assignment now correctly resolve the user's full display name and identity record. Previously, queries such as "List all parts created by Admin" or "What is assigned to me?" could fail to match because the agent used the login name rather than the full display name or identity record.

Note: This improvement requires the User and Identity item types to be configured in the Edge API with the required properties and relationships.

Non-consecutive version comparison now supported (IH-21307)

Users can now compare any two versions of an item, not only adjacent versions. Previously, comparing non-consecutive versions, such as Version 1 against Version 5, was not supported.

Write · Enabled Badge No Longer Shows Stale Status (IH-20760)

The **Write · Enabled** badge on a connection in Access Points previously persisted after write actions were removed from the connection, until a manual schema refresh was triggered. The badge now reflects the current configuration immediately.



Context mode get operations restored (IH-20752)

Agents using context mode can now correctly perform get operations. Previously these operations failed in context mode.

Chat session stability under high load (IH-20754)

Chat sessions no longer become unusable after the service encounters a high-load condition. Sessions now recover correctly.

Average-value queries use the correct field (IH-20755)

Average-value queries such as "What is the average duration of the Design Review activity?" now apply filters against the correct field. Previously these queries sometimes resolved against the Classification field instead of the Activity Name field.

2026 July 14 - Edge AI - Incremental Release - Edge AI version 1.1.2

July 14, 2026 - Incremental Release

Aras InnovatorEdge AI Version 1.1.2 is now available.

Introducing expanded query coverage for relational and historical data, and broader item type support.

What's new

Queries on complex relational properties now resolve consistently



The AI Assistant now correctly handles queries that filter items through properties that reference a related item type in Innovator. Previously these queries returned no results or failed silently. Identity-type properties (such as the person who created, modified, owns, or manages a record) are the most common example of this pattern.

Why this matters: A wide range of PLM queries depend on filtering through a relational property. These queries now work reliably across the item types and property types where they are configured.

Example scenario:

- A user asks "List all Parts created by Mike Miller" and receives a filtered table of results. The same works with variations such as "Parts owned by", "Parts modified by", or "Parts managed by" the same person, each resolved correctly through the Identity-type relationship.

The Agent Guides You When a Query Term Does Not Match

When the agent cannot find an exact match for a term used in a query, it now acknowledges the gap, surfaces the closest available result, and suggests how to refine the question, rather than returning a silent empty response.

Why this matters: Users no longer need to know the agent's exact vocabulary to get useful answers. Near-miss queries now produce a guided response instead of a dead end.

Example scenario:

- A user asks "Show me assembly parts marked as flagged." The agent cannot validate "flagged" as a data field, so it proceeds with what it can resolve: "Interpreted assembly parts as Parts whose Type contains assembly. I could not validate a metadata-backed field for flagged, so I retried using only the assembly filter. There are 244 matching Parts; the first 25 are available in the current result set." The user receives results and knows exactly what was searched.

Query items related through complex type structures



The agent can now correctly retrieve related items in cases where the related item type is a Polyitem. Previously, relationship queries involving these item types returned incomplete or no results.

Why this matters: Product structures in Aras Innovator often include item types that can resolve to multiple concrete types. Agents can now navigate these structures accurately, broadening the range of relationship queries that return correct results.

Example scenario:

- A user asks "Show me all Parts managed by Jane Smith." The agent correctly resolves the "managed by" field as an Identity-type relationship property and returns the expected results rather than an empty or partial set.

Access item history through the agent

Users can now ask the agent to retrieve the change history of a specific item directly in the chat interface.

Why this matters: Item history is a routine PLM lookup: tracking when a part changed, who changed it, and what changed. It is now available without leaving the AI Assistant.

Example scenario:

- A user asks "Show me the history of Part MP0101" and receives a list of historical changes for that item, including dates and the nature of each change.

Date range and duration calculations now return accurate results

The agent now correctly calculates time spans and duration metrics involving Innovator date properties. Previously, queries requiring the agent to compute elapsed time or compare date values



across properties returned inconsistent or no results. Workflow activity duration is a key example of this pattern.

Why this matters: Many operational PLM questions involve time: how long a step took, how overdue an item is, or the average cycle time for a process. These queries now produce reliable numeric results.

Example scenario:

- A user asks "What is the average number of days for the Design Review activity?" and receives a computed average duration based on the start and completion date properties available in the activity records.

Bug fixes

Connection Status badge truncated at 1024px viewport (APR-1029)

Fixed a layout issue where the Connection Status badge on the Agent Overview page was hidden or truncated when the browser window was resized to 1024px width.

2026 June 23 - Edge AI - Incremental Release - Edge AI version 1.1.1

Aras InnovatorEdge AI Version 1.1.1 Is Now Available

June 23, 2026 - Incremental Release

Introducing a redesigned agent configuration experience, centralized Access Point management, agent testing with step-by-step transparency, and a renamed Models section.

What's new

Redesigned agent configuration



Administrators who create and manage agents now work in a structured configuration editor. The editor keeps the configuration state visible at every step and saves the agent as a Draft. After saving, a read-only Agent Overview becomes available, providing a direct path back to editing.

Why it matters: Reduces errors during agent setup by making required fields explicit and the agent's current state clear throughout the process.

Example scenario:

- An admin creates a new agent, completes the configuration form, and saves. The agent is now in Draft, not published to end users, while the admin continues to refine and test it before publishing.

Agent testing with execution trace

Before publishing an agent, administrators can test it against live data in a full-page Simulation Chat, visible only to admins. Each response includes an expandable Steps trace showing exactly what the agent did, which data it queried and what decisions it made.

Why it matters: Gives administrators confidence in agent behavior before exposing it to end users, without requiring a separate test environment.

Example scenario:

- An admin opens the Simulation Chat for a newly configured agent, sends a test prompt, and expands the Steps trace to verify the agent is querying the correct dataset and returning results in the expected format, all before publishing.

Access Points - centralized connection management

Edge API connections are now managed from a dedicated Access Points screen. Admins can view all registered connections in one place, add new connections, see the health status of each connection, and edit connection details after creation.

Why it matters: Gives administrators a single place to manage and monitor all data connections



across their tenant, with visibility into connection health before assigning a connection to an agent.

Example scenario:

- An admin navigates to Access Points, sees that one connection is showing an error status, updates the credentials in the connection detail view, and verifies the connection is healthy before the agent that depends on it is republished.

Connection health monitoring on Agent Overview

If a previously working Edge API connection becomes unreachable, the Agent Overview flags the issue and blocks the agent from being republished until the connection is restored or removed.

Why it matters: Prevents an agent from being republished in a broken state, protecting end users from connection-related failures.

Example scenario:

- An API key expires overnight. The following morning, the Agent Overview shows the connection as unhealthy. The admin updates the key in Access Points and the agent can be republished.

Models - updated registry label

The model registry, previously labelled LLM Connections, is now labelled Models. Selectors within agent configuration now distinguish between Reasoning Models and Embedding Models, making model selection clearer for administrators.

Why it matters: The updated labels make it easier to identify and assign the right model type when configuring an agent.

Permission level visible on Agent Library



Each agent on the Agent Library page now displays its permission level. Administrators can assess agent access at a glance and filter the list by permission level without opening individual agent records.

Why it matters: Reduces the steps required to audit agent access across a tenant.

Example scenario:

- A tenant admin performs a quarterly access review by filtering the Agent Library by permission level, identifying agents with broader access than expected, and adjusting their configuration accordingly.

Important notes

Data write actions not available in this release

Agents cannot create or edit data in Aras Innovator in this release. This capability is planned for a future release.

Bug fixes

Blank page after canceling agent creation (APR-872)

Fixed an issue where canceling the creation of a new agent displayed a blank page instead of returning the administrator to the Agent Library.

2026 May 12 - Edge AI - Incremental Release - Edge AI version 1.1.0

Edge AI Commercial Release 1.1.0 Is Now Available

May 12, 2026 - Incremental Release

Introducing conversational BOM and Where-Used queries, bring-your-own LLM support, agent-level model assignment, and multiregional deployment.



What's new

Query-on-the-fly: BOM and Where-Used

The Natural Language Analytics (NLA) Agent now executes live queries across item datasets for Bill of Materials (BOM) and Where-Used analysis. BOM results are rendered as indented tables directly in the chat interface.

Why it matters: Provides immediate visibility into assembly hierarchies and parent-child relationships without requiring precomputed reports or manual navigation to the Structure Browser.

Example scenario:

- A user asks "Show me the BOM of Assembly X" and receives an indented table showing the first two levels of the hierarchy, including part numbers and quantities.

Bring your own LLM (BYOLLM)

Admins can register and manage external model providers through a centralized LLM Connections interface.

Why it matters: Enables the use of existing model subscriptions.

Example scenario:

- A Tenant Admin configures a specific Azure OpenAI endpoint to ensure agent traffic remains within a designated regional boundary.

Agent-level model assignment

Admins can assign a specific registered LLM to an individual agent via a governed catalog dropdown



in the agent configuration menu. A global default model can also be set at the tenant level, ensuring every new agent is preconfigured with an approved model.

Why it matters: Every agent might have different needs in terms of the model to be used.

Example scenario:

- An admin designates either a high-capacity or a speed-optimized model as the tenant default so that any new agent is automatically pre-selected with that model.

Multiregional deployment

The Edge AI Platform is now available for distribution across US and EU regions. All Edge services prominently display the active organization or tenant name in the user menu.

Why it matters: Improves service performance across geographical regions.

Bug fixes

Technical Documentation application support [APR-473]

Resolved an issue where the AI Agent was unable to process requests or retrieve structured content associated with Technical Documentation application ItemTypes.

2026 April 9 - Edge AI - Incremental Release - Edge AI version 1.0.1

Aras InnovatorEdge AI Version 1.0.1 Is Now Available

April 9, 2026 - Incremental Release

What's new



Agent lifecycle management

Agents now move through formal **Draft**, **> Published**, and **Unpublished** states. Administrators see real-time status indicators in both the agent list and detail view. Published agents are locked for editing until explicitly unpublished. Only published agents are visible to end users in the agent selection list.

Why it matters: Gives administrators full control over what end users can access:

- Agents under construction stay hidden until they are ready
- Live agents cannot be accidentally half-edited

Example scenarios:

- Create and test a new agent privately in Draft before ever exposing it to users via the preview chat.
- Unpublish a live agent to safely revise its system instructions, then republish. Users see no in-between state.

API connection timestamps

The administration interface now shows **Last Modified** and **Updated On** timestamps for every configured API connection, so you can tell when the cached schema was last synchronized with Aras Innovator.

Why it matters: Eliminates guesswork around stale metadata.

You can now check whether a schema refresh is needed, rather than discovering a discrepancy after an agent misbehaves.

Example scenario:

- After adding a new property to an ItemType in Innovator, check the timestamp to confirm the Edge AI platform has picked up the change. If not, trigger a manual refresh via a connection



test or description update.

Bug Fixes

Inconsistent chart output for date-based queries (APR-472)

Date-driven queries were producing incorrect or inconsistent chart visualizations. Resolved with stricter chart generation prompts, improved data integrity guardrails, and better OData query validation.

Help button missing in standalone mode (APR-603)

The Help button was not appearing in the Edge AI chat header when running in standalone mode. Restored to the expected location in the application header.

2026 April 2 - Edge AI - Release Announcement - Edge AI version 1.0.0

Aras InnovatorEdge AI foundational release 1.0.0 Is Now Available

April 2, 2026 - Foundational Release

Overview

This release announces general availability of **Aras InnovatorEdge AI**, the AI intelligence layer for the Aras Innovator digital thread.

InnovatorEdge AI activates the intelligence already embedded in your Aras Innovator system of record, making product lifecycle data accessible to every user through natural language - governed by the same permissions, policies, and data model your organization already owns and controls.

What Is Included in This Release

Natural Language Analytics (NLA) Agent



The first production agent in the InnovatorEdge AI ecosystem enables any Innovator user, regardless of PLM expertise, to query structured product data through plain language conversation. Engineers, managers, and executives can ask questions about parts, BOMs, engineering changes, lifecycle states, and project status, and receive results as tables, charts, and Key Performance Indicators (KPIs) with direct links back to source records in Innovator. No AML knowledge. No filter configuration. No administrator dependency.

Aras Innovator AI Assistant

The NLA Agent is delivered through the **AI Assistant**. It is available as an embedded panel within Aras Innovator and as a standalone Edge AI application. The embedded panel is context-aware: it automatically detects the item currently being viewed and scopes queries accordingly.

InnovatorEdge AI Platform

This release establishes the foundational platform on which all current and future Edge AI agents operate:

- **Tenant isolation:** secure separation of agent configurations, API connections, and session data across all tenants
- **Identity propagation:** full CIAM integration ensures every agent interaction respects the individual user's Innovator permissions; no data governance bypass is possible
- **Edge API binding:** all data access flows exclusively through the governed Edge API layer; agents cannot access data outside their authorized scope
- **Agent configuration framework:** administrators configure agents through a no-code UI; instructions, data connections, and starting questions are all configurable without engineering involvement

Why This Matters



PLM data is often the most interconnected product knowledge in the enterprise.

The barrier has not been the data it has been an abundance of data often siloed across the same Innovator account.

InnovatorEdge AI removes that barrier by providing a secure gateway to that information, broadening access while keeping every interaction within existing enterprise governance and user authorization. Every query executes under the user's own Innovator session. Every result is drawn from authorized data. Every interaction is ephemeral; no customer data is used to train or improve AI models.

Deployment

InnovatorEdge AI is available for Full Software as a Service (SaaS) and Hybrid deployments. The Aras Innovator instance can remain on-premises; Edge AI connects through the Edge API layer. This release requires an active Aras Innovator subscription and an active InnovatorEdge API subscription.

What Is Coming Next

The NLA Agent is the first in a broader ecosystem. Subsequent releases will introduce:

- **InnovatorEdge AI Task Agents:** Purpose-built agents for Requirements Ingestion, BOM Import, and Smart Variants workflows
 - **Knowledge Base Discovery Agent:** RAG-powered semantic retrieval over structured and unstructured PLM content
 - **Semantic AI Search:** context-aware, relevance-ranked search across all Innovator items
-

