EGKITS SaaS V11

by IDSYSEG · Complete module guide · Version 1.0

Updated Nov 2026

AI Knowledge & Training Center and Strategic Advisor

Approved knowledge for the AI agent, and evidence-based advice that people decide on

Approved, cited answers | Evidence-based advice | Tenant isolation

EGKits SaaS V11

A complete reference for business owners, administrators, knowledge owners, auditors and engineers. It explains how the AI Knowledge & Training Center decides what the EGKits AI agent knows, how content from help, guides, web sites and files becomes approved, versioned and cited knowledge, and how the Strategic Advisor finds gaps from the system's own evidence and proposes prioritized recommendations without changing anything by itself.

Field	Value
Document version	1.0
Publication date	23 November 2026
Applies to	EGKits SaaS V11, AI Agent area
Audience	Owners, administrators, knowledge owners, auditors and engineers
Distribution	Public. No tenant data, prompts, credentials or production logs.
English file	`EGKits_AI_Knowledge_Strategic_Advisor_Guide_EN.pdf`
Arabic file	`EGKits_AI_Knowledge_Strategic_Advisor_Guide_AR.pdf`

Enforced control

The AI agent answers only from approved knowledge the person asking may read, never approves its own content, and never sees another workspace's knowledge. The Strategic Advisor only recommends: it never changes code, settings, business rules or data.

Training sources

Built-in help, workflows, guides, PDF guides, the API and the module registry, plus web sites, uploads and articles.

Human approval

New and changed content waits for a person, with differences, conflicts and instruction-like text shown.

Strategic Advisor

Deterministic gap rules, ranked recommendations, human decisions and scheduled assessment reports.

Contents

Select a heading to jump to that section.

1. How to read this guide 3
2. Executive summary 3
- What makes it safe for a business 3
3. Core concepts 4
4. Platform and workspace knowledge 4
5. Security and isolation 4
6. Training sources 5
- Website crawling 5
7. Synchronization 6
8. Processing pipeline 6
9. Validation and approval 7
10. Versions, rollback and retirement 8
11. Coverage and search quality 8
12. Answers and citations 8
13. Strategic Advisor 9
- Evidence basis 9
14. Gap analysis rules 9
15. Readiness and limitations 10
16. Priority and classification 10
17. Findings and decisions 11
18. Assessment reports 11
19. Getting started 12
20. Permissions 12
21. Page reference 13
22. Configuration defaults 14
23. Troubleshooting 14
24. Glossary 15

1. How to read this guide

The guide follows the path of knowledge through the product: where it comes from, how it is processed and approved, how the agent uses it, and how the Strategic Advisor reads the same evidence to find gaps. Each chapter names the page, the permission and the statuses exactly as the product shows them.

Reader	Start with	Then read
Business owner or CEO	2. Executive summary	5, 13, 16, 17 and 18
Administrator	19. Getting started	4, 6, 7, 9, 20, 22 and 23
Knowledge owner	6. Training sources	7 to 12
Auditor or security officer	5. Security and isolation	9, 10, 15, 17 and 20
Engineer or operator	8. Processing pipeline	7, 22 and 23

Four kinds of note appear in the text. They are not interchangeable.

Note	Meaning
Enforced control	A rule the server enforces whatever the user, the content or the model does. Each one has automated tests.
Configuration	A value the workspace or the deployment can change. The default is stated.
Known limit	A documented boundary of the current release.
Tip	A practical recommendation for daily use.

2. Executive summary

The AI Knowledge & Training Center decides what the EGKits AI agent knows. It collects content from training sources, turns it into versioned knowledge items, and holds every new or changed version until a person approves it. The agent answers only from approved versions, cites the items it used, and uses only the items the person asking may read.

The Strategic Advisor reads the same evidence, together with the module registry, the API of the running application and the business workflows, to find gaps between what EGKits does, what it documents and what it should do. It ranks recommendations by impact, value, return, security, risk and effort, and leaves every decision to a person.

What makes it safe for a business

- Approved answers only. New and changed content waits for a person; until then the previous approved version keeps serving.
- Untrusted content stays data. Uploaded files, web pages and documents are never followed as instructions; instruction-like text is flagged and approving it needs a note.
- Permission-filtered retrieval. Every search runs with the asker's permissions, including citations and coverage.
- Workspace isolation. Workspace knowledge, findings, recommendations and reports belong to one workspace. The workspace comes from the signed-in session, never from the request.
- Facts apart from opinions. Verified findings, AI inferences and market observations are labelled and shown separately.
- Nothing changes by itself. Neither module changes code, settings, business rules or business data.

3. Core concepts

The same terms are used on every page, in the history and in the reports.

Concept	What it is
Training source	Where content comes from, with a type, trust level, priority, schedule and exclusions.
Knowledge item	One document from a source, such as a help topic, a page or a file, with its module, language and required permission.
Version	An immutable snapshot of an item's content. Only one version is served at a time.
Passage	A searchable slice of a version, about 1,200 characters with a 150-character overlap.
Synchronization	One run of a source: what it found, added, updated, left unchanged, skipped or failed, and why.
Approval	A person's decision to serve or reject a new version.
Trust level	Untrusted, Community, Trusted or Authoritative. When sources disagree, higher trust and priority win.
Citation	The item an answer relied on. Citations are kept with the answer; passage text is not.
Finding	A gap detected by an advisor rule, with the evidence that was checked.
Recommendation	A proposed improvement for one or more findings, with a priority score and a decision status.
Assessment report	A dated snapshot of findings and recommendations, in PDF or Markdown.

4. Platform and workspace knowledge

Knowledge lives in two separate stores. A question is answered from both together, each filtered by the asker's permissions.

Store	Who changes it	Who reads it
Platform	The platform workspace, with the platform knowledge permission. Holds the built-in help, documentation, PDF guides, API list and module registry.	Every workspace, read-only.
Workspace	People in the workspace with the knowledge and source permissions. Holds the workspace's own sources, uploads and articles.	That workspace only.

Enforced control

Platform administrators get no automatic access to a workspace's knowledge, findings or reports. Platform checks of the Strategic Advisor run only in the platform workspace.

5. Security and isolation

These controls are enforced by the server and covered by automated security and isolation tests. They do not depend on instructions given to an AI model.

Control	What it means
Workspace isolation	Every record carries its workspace and is filtered by the data model. Background runs resolve the workspace again before touching data.
Server-side workspace	The workspace is taken from the signed-in session. A workspace id sent by the browser is never proof of access.
Human approval	Only a person can approve, reject or restore knowledge, and decide on recommendations.
Untrusted content	Files, pages and documents are data, never instructions. Instruction-like text is flagged.

Control	What it means
Secret masking	Keys, passwords and similar secrets are masked before anything is stored. History and exports never contain them.
Provider keys	AI provider keys stay in server-side secret storage and are never sent to the browser.
Traceability	Every source, run, version, approval, rollback and decision is recorded with the person and the time.

Enforced control
Generated AI content is never treated as authoritative business data. Nothing the agent learns from feedback is served until a person approves it.

6. Training sources

A training source says where content comes from and how often it is read. Built-in sources are created by the platform; a workspace adds its own.

Type	Content	Default
In-app help	Help topics shown inside EGKits, in English and Arabic	Built-in, daily at 02:00 UTC
Help workflows	Business workflows and their stages	Built-in, daily at 02:00 UTC
Public documentation	The guides on the public website	Built-in, daily at 02:00 UTC
PDF catalog	The downloadable PDF guides	Built-in, Sundays at 03:00 UTC
API manifest	The API of the running application	Built-in, daily at 02:00 UTC
Module registry	Modules and their capabilities	Built-in, daily at 02:00 UTC
Website	Pages of a site on the allowed domain list	Off until a domain is allowed
Uploaded files	PDF, Word, Excel, PowerPoint, text and Markdown, up to 25 MB each	On upload
Manual entry	Articles written by hand	On save

On a new installation the built-in sources synchronize once shortly after start-up, then follow their schedule. Schedules run at most once an hour, in UTC: by hand only, hourly, daily, weekly or a custom cron expression. Sources created by the platform cannot be deleted; deactivate them instead.

Website crawling

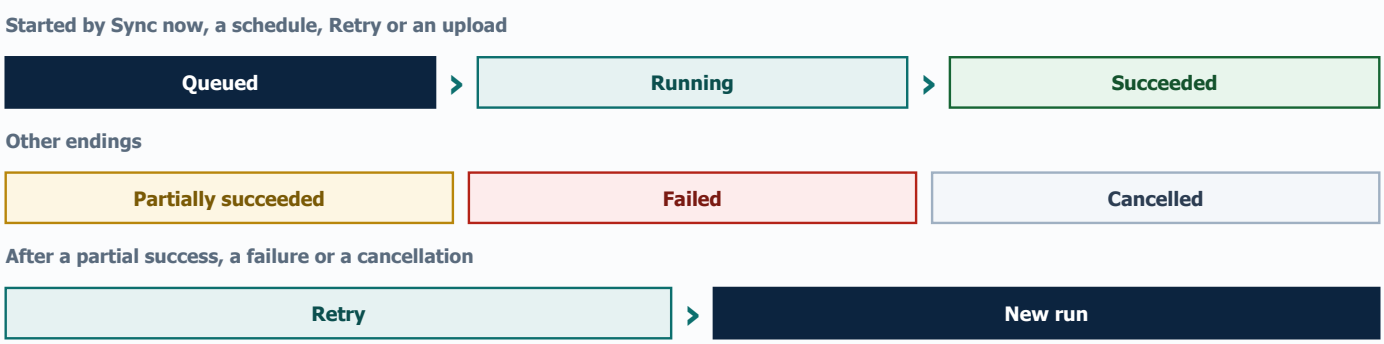
- Only https pages on domains in the workspace's allowed list are fetched.
- robots.txt, noindex and nofollow are honoured, and the crawler waits 500 ms between requests.
- At most three redirects and four levels of links are followed, and the crawl stops at the workspace's page limit.
- Pages already held are requested with their validators, so unchanged pages are not downloaded again.

Tip
Dry run shows what a synchronization would add, update or fail, without storing anything. Use it before the first real run of a new website source.

7. Synchronization

A synchronization run reads one source and records the outcome of every item. Only one run of a source can be queued or running at a time.

Run lifecycle



Item outcome	Meaning
Added	A new item and its first version
Updated	A new version of an existing item
Unchanged	Identical to the served version; no version is created
Missing	The source no longer lists the item; its approved version keeps serving
Skipped or Excluded	Out of scope, or matched an exclusion of the source
Failed	The item stopped at a step; the reason is recorded and the run continues

Known limit
Cancel stops after the item being processed. A run left unfinished, for example by a restart, is closed as failed after 120 minutes so the source can run again.

8. Processing pipeline

Every item goes through thirteen steps. A step that fails stops only that item; the rest of the run continues. The Processing page shows each step, running work and recent problems.

The thirteen steps

1. Discover the item in its source.
2. Extract the text.
3. Normalize it and mask secrets.
4. Detect the language.
5. Classify it by application and module.

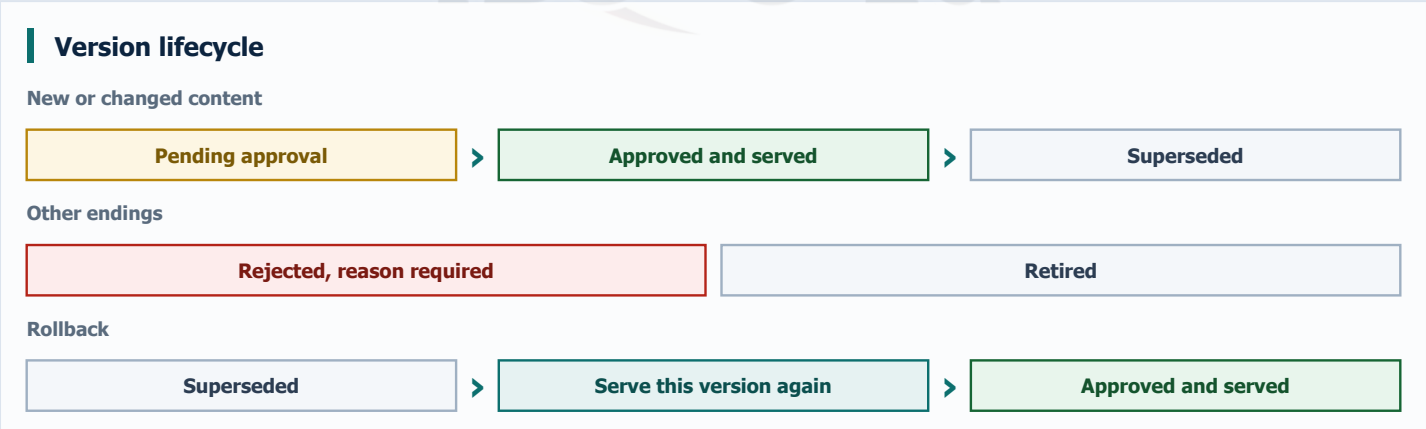
6. Split it into passages.
7. Embed the passages for search by meaning.
8. Relate it to existing knowledge.
9. Remove duplicates. Content identical to the served version is recorded as unchanged.
10. Validate it.
11. Check for conflicts with approved knowledge and for instruction-like text.
12. Hold it for approval when needed.
13. Publish the approved version.

Passages are found by meaning when an embedding model is configured in the AI Assistant settings, and by keywords otherwise. Every six hours a maintenance job closes abandoned runs and embeds versions that were indexed by keywords only.

Configuration
A run handles at most 5,000 items, 512 MB or 30 minutes. A version is split into at most 400 passages.

9. Validation and approval

New and changed content waits on the Validation and approval page. The reviewer sees the differences from the served version, any conflict with approved knowledge and any instruction-like text, and then approves or rejects the version.



- Only an authoritative source may approve its own content automatically, and only when nothing was flagged. Everything else waits for a person.
- Approving a version that conflicts with approved knowledge, or contains instruction-like text, needs a note, and the item must be opened first.
- Until a new version is approved, the previous approved version keeps serving.

Enforced control
The AI agent can never approve knowledge, including content it proposed itself. Approval needs the approval permission and is recorded with the person, the time and the note.

10. Versions, rollback and retirement

Versions are never edited. A change always creates a new version, so every answer can be traced to the exact content that was served when it was given.

Action	Rule
Compare versions	Shows the differences between two versions of an item.
Serve this version again	Only an earlier approved version can be served again. Needs the approval permission.
Retire item	The item is no longer served. Its history is kept.
Rebuild passages	Indexes the served version again, for example after an embedding model is configured.

Status set	Values
Knowledge item	Active, Missing, Retired
Version	Pending approval, Approved, Rejected, Superseded, Retired
Passage index	Pending, Indexed, Keyword only, Unsupported, Failed

Tip
When a source no longer lists an item, the item becomes Missing and its approved version keeps serving until someone retires it. Review Missing items after reorganizing a website or a document library.

11. Coverage and search quality

The Coverage page shows, for each module, whether the agent has approved knowledge in English and in Arabic, together with the health of every source. It helps a knowledge owner find the modules the agent cannot answer about yet.

- Test search shows what a question would find, with the passages and items, using the reviewer's own permissions.
- Search quality runs the English and Arabic evaluation questions through retrieval and reports how many find their expected items, per language.
- Coverage counts only the items the viewer may read, so two people can see different numbers.

12. Answers and citations

When the agent answers from knowledge it lists the items it used. Citations stay with the answer, so they are shown again when the conversation is reopened; the passage text itself is not stored with the conversation.

An answer flagged as wrong can be captured as a feedback candidate for review when the workspace turns on learning from feedback. A candidate is a proposal like any other new content: it is not served until a person approves it.

Enforced control
An item that requires a permission is invisible to people without it, in answers, citations, test search and coverage alike.

13. Strategic Advisor

The Strategic Advisor looks for gaps between what EGKits does, what it documents and what it should do. Its findings come from deterministic rules that check the system's own records: the module registry, the API of the running application, the approved help and guides, the business workflows and the health of the Knowledge Center. The same evidence always gives the same findings.

Evidence basis

Basis	Meaning
Verified	Checked against the system's own records by a deterministic rule.
AI inferred	A lead suggested by AI, not confirmed against an authorized source.
Market observation	An AI-written opinion about the product and its market. Optional, at most five per analysis, and always labelled as not verified.

Enforced control
The advisor never changes code, settings, business rules or data. Accepting a recommendation, or turning it into a task, changes nothing by itself.

Analysis run

Run analysis or a scheduled report

Queued › Collect evidence › Rules A to D › Save and report › Succeeded

Optional and failure paths

Market observations, labelled | Failed

14. Gap analysis rules

Rules A to C run in the platform workspace only. Every workspace gets the D rules for its own sources, index and insights.

Gap type	Rule	Finds
A. Features against capabilities	Capability without API	A module whose capabilities have no matching endpoint in the running API
B. Implementation against documentation	Module without help	A module with no approved help, guide, workflow or PDF
B	Help missing Arabic or English	Help topics that exist in only one language
B	Help module not in the registry	Help describing a module the registry does not list
C. Business process	Stage without screen	Workflow stages that do not name the screen that performs them
C	Stage screen without API	A stage whose screen has no matching API
D. Technical and architectural	Endpoint without permission	An endpoint that requires no permission and is not public
D	Stale source	A source not synchronized for more than 14 days
D	Source never synchronized	A source with no completed run a day after it was created
D	Source failures	A source whose last run failed or had failed items

Gap type	Rule	Finds
D	Keyword-only index	More than half of at least 20 passages without embeddings
D	Approval backlog	Versions waiting for approval for more than 7 days
D	Agent insight	Open warnings raised by the AI agent
E. Product and market	Market observation	Optional AI-written observations, kept apart

15. Readiness and limitations

The Strategic Advisor overview shows how many rules can run and which evidence is ready, waiting for its first synchronization, or missing: the module registry, the API of the running application, help and guides, business workflows, knowledge sources, the knowledge index and AI insights. It is worked out from the evidence available now and saves nothing.

Rules are skipped, with the reason stated in the report, when:

- the API of the running application is not available, so only the generated API list exists;
- the built-in help and documentation sources have not finished their first synchronization;
- no help or guides have been synchronized;
- no business workflows were found.

> **Known limit**
> When evidence is missing the advisor says so instead of reporting gaps that may not exist. A finding that a skipped rule produced earlier stays open; it is not marked resolved because the rule could not run.

16. Priority and classification

Each recommendation gets a priority score from 0 to 100 from six factors rated 1 to 5, less a cost of up to 10% for effort, complexity and dependencies.

Factor	Weight
Business impact	25%
Security impact	20%
Customer value	15%
Expected return	15%
Operational risk	15%
Effort, complexity and dependencies	Up to 10% deducted

Priority	Score
Critical	75 and above, or the highest security impact with high operational risk
High	55 to 74
Medium	35 to 54
Low	Below 35

Recommendations are also classified, for example critical improvement, security and compliance, process automation,

documentation improvement, integration opportunity or new product feature.

17. Findings and decisions

Findings

Detected by an analysis

Open › Resolved, not detected again

A person marks it

Open › Dismissed, not a gap › Reopen

Recommendations

Typical path

Proposed › Accepted › Assigned › Converted to task

Other decisions

Rejected, note required | Postponed to a date | Reopen to Proposed

Decision	Rule
Accept	Records that the recommendation will be acted on.
Reject	A note is required. A rejected recommendation must be reopened before it can be assigned or converted.
Postpone	A future date is required. The recommendation returns when the date passes.
Assign	An assignee is required.
Convert to task	Creates a To-Do task with the problem, the solution and the steps. Converting again opens the same task.
Reopen	Brings a decided recommendation back for a new decision. A recommendation converted to a task cannot change any more.

Explain asks an AI model to explain a recommendation from its evidence. The text is labelled as AI-written and does not change the findings, the score or the status.

> **Enforced control**
> Every decision records the person, the time, the previous and new status and the note. The AI agent cannot decide on a recommendation.

18. Assessment reports

A report is a snapshot of the findings and recommendations at the time it was written, with counts by type and priority, the evidence basis kept apart, and the limitations of the analysis. Reports download as PDF or Markdown.

Period	When
On demand	Each time a person runs an analysis, unless the report is turned off. Written in the reader's language.
Daily	Every day, after the daily check at 04:41 UTC
Weekly	On Mondays

Period	When
Monthly	On the 1st of the month

Configuration
Scheduled reports are off until the workspace turns them on in its Knowledge Center settings, and are written in the workspace language.

19. Getting started

Knowledge Center

1. Assign the Knowledge Center permissions to the right roles.
2. Configure an AI provider and an embedding model in the AI Assistant settings. Without one, passages are searched by keywords.
3. Open Knowledge Center and check the counts. On a new installation the built-in sources are already synchronizing.
4. Open Training sources to add a website on an allowed domain, uploaded files or a manual article.
5. Open Validation and approval and approve the versions the agent should answer from.
6. Open Coverage to find modules without knowledge in English or Arabic, and use Test search.

Strategic Advisor

1. Assign the advisor permissions to the right roles.
2. Wait until the readiness card shows help as ready on a new installation.
3. Open Strategic Advisor and choose Run analysis. Add market observations only if you want AI opinions alongside verified findings.
4. Review Gap analysis, mark findings that are not real gaps, and open the evidence of the rest.
5. Decide on Recommendations and turn the accepted ones into tasks.
6. Turn on the daily, weekly or monthly reports you want.

20. Permissions

Permissions are assigned to roles in the workspace. The server checks them on every request, whatever the page shows.

Permission	Allows
AiKnowledge.View	Read knowledge items
AiKnowledge.Search	Test search
AiKnowledge.Create	Write articles and upload files
AiKnowledge.Update	Change an item's details and who may see it, and rebuild passages
AiKnowledge.Delete	Retire items
AiKnowledge.Export	Export lists to Excel
AiKnowledgeSources.View	See training sources and settings
AiKnowledgeSources.Create	Add training sources
AiKnowledgeSources.Update	Change sources and settings
AiKnowledgeSources.Delete	Delete workspace sources
AiKnowledgeSources.Execute	Start, retry and cancel synchronizations
AiKnowledgeApprovals.View	See the approval queue
AiKnowledgeApprovals.Approve	Approve or reject versions, and serve an earlier version again
AiKnowledgePlatform.Manage	Change platform knowledge, in the platform workspace only
AiAdvisor.View	Overview, readiness, findings, recommendations and reports
AiAdvisor.Generate	Run an analysis, write reports and ask for an AI explanation
AiAdvisor.Manage	Decide on recommendations and dismiss or reopen findings
Todos.Create	Also needed to convert a recommendation to a task

> **Tip**
> The History export shows people's names only to someone who may read the users list.

21. Page reference

Page	Route	What you do there
Knowledge Center	/ai-agent/knowledge-center	Counts, waiting approvals and links to every section
Knowledge repository	/ai-agent/knowledge-center/repository	Find items, compare versions, roll back, retire, upload, write articles
Training sources	/ai-agent/knowledge-center/sources	Add, change, test and synchronize sources
Processing	/ai-agent/knowledge-center/processing	The thirteen steps, running work and recent problems
Sync center	/ai-agent/knowledge-center/sync	Every run with its items, retry and cancel
Validation and approval	/ai-agent/knowledge-center/approvals	Review differences and conflicts, approve or reject
History and logs	/ai-agent/knowledge-center/history	Who changed what and when
Coverage	/ai-agent/knowledge-center/coverage	Coverage by language, source health, test search and search quality
Strategic Advisor	/ai-agent/advisor	Readiness, latest analysis and Run analysis
Gap analysis	/ai-agent/advisor/gaps	Findings with their evidence
Recommendations	/ai-agent/advisor/recommendations	Priorities and decisions
Assessment reports	/ai-agent/advisor/reports	Reports and their downloads

22. Configuration defaults

Every limit has a safe default in code, so nothing has to be added to configuration files. A deployment may override them in the AIKnowledge configuration section. The allowed website domains, the page limit and scheduled reports are workspace settings on the Knowledge Center pages.

Setting	Default
Passage size and overlap	1,200 and 150 characters
Passages per version	400
Upload size	25 MB
Per run	5,000 items, 512 MB, 30 minutes
Abandoned synchronization	Closed after 120 minutes
Abandoned analysis	Closed after 60 minutes
Maintenance	Every 6 hours at minute 23, UTC
Report check	Daily at 04:41 UTC
Stale source	14 days
Approval backlog	7 days
Keyword-only index	More than 50% of at least 20 passages
Market observations	At most 5 per analysis
Web page and PDF size	4 MB and 20 MB
Web request timeout and delay	20 seconds and 500 ms
Redirects and link depth	3 and 4

Known limit

Scheduled synchronizations and reports need the server's background job service. When it is not available, schedules are saved but do not run; use Sync now and Run analysis instead.

23. Troubleshooting

Symptom	What to do
The counts are zero on a new installation	The built-in sources synchronize shortly after start-up. Follow them in the Sync center.
Schedules are saved but do not run	Scheduled work is not available on this server. Use Sync now.
A web source fetches nothing	Check that crawling is allowed for the workspace and the domain is on the allowed list.
An item shows Keyword only	No embedding model was available when it was indexed. Rebuild passages once a model is configured.
A run is Partially succeeded	Open it and filter by Failed to see each item's reason.
An answer ignores a new document	The new version is still waiting in Validation and approval.
A colleague cannot see an item you see	The item requires a permission the colleague does not have.
Fewer advisor rules can run than the total	The readiness card lists the reason. Synchronize the missing sources or wait for the first synchronization.
The API rules are skipped	The running application's endpoint list was not available. Run the analysis again once the server has fully started.

Symptom	What to do
The PDF report has no Arabic font	Download the Markdown report and ask the administrator to install an Arabic-capable font such as Noto Sans Arabic or DejaVu Sans.
No market observations	No AI model is configured, or the option was not chosen. Verified findings are not affected.

24. Glossary

Term	Meaning
Authoritative source	The highest trust level. The only level that may approve its own unflagged content.
Dry run	A synchronization preview that stores nothing.
Embedding	A numeric representation of a passage used to search by meaning.
Instruction-like text	Content that tries to tell the AI what to do. It is flagged and never followed.
Served version	The one approved version of an item that the agent answers from.
Readiness	Which evidence the advisor can use now and how many rules can run.
Not a gap	A person's decision that a finding is not real. It stays closed until reopened.

Related reading on the public website: the AI Knowledge & Training Center, Strategic Advisor and AI Knowledge & Advisor Workflows guides under Documentation.