opentext™

OpenText Core Data Discovery & Risk Insights

Application Version 26.3.0

Processing Agent Version 26.3.100

Release Notes

Document Release Date: August 2026
Software Release Date: August 2026

Legal notices

Contents

OpenText Core Data Discovery & Risk Insights Release Notes

Application version: 26.3.0

Processing Agent version: 26.3.100

Publication date: August 2026

This document is an overview of the changes made to the OpenText Core Data Discovery & Risk Insights application and processing agent.

Support matrix

For information about the installation requirements and compatibility with other products, see the OpenText Core Data Discovery & Risk Insights Support Matrix. The support matrix may be updated between releases so it is available only from the Support portal.

About OpenText Core Data Discovery & Risk Insights

OpenText Core Data Discovery & Risk Insights lets you find, protect, and secure sensitive and high-value data within on-premises and cloud data platforms across your enterprise. Identify, collect, and organize content to ensure discovery of sensitive data. Configure how structured and unstructured sources and datasets are processed and categorized with Connect. Analyze your data under management with Analyze. Organize, review, and take action on documents and unstructured data items with Manage.

About this release

This product release and release notes cover Application 26.3.0 and Processing Agent 26.3.100. These release notes reference changes included in this release only. For information about changes made in prior releases, see the release notes for the desired release.

While the prior release's Processing Agent remains compatible with this application's version for existing functionality, product enhancements and fixes may require an upgrade to the latest version of the agent. In order to facilitate transition to the new version, the existing processing agent an be used for a transitional period of four (4) weeks after upgrade of the application to this release. Beyond this grace period, out of date agents will be automatically disconnected and an upgrade will be mandatory to reconnect and process data.

What's new

The following features and enhancements are included in this release.

General

The following features and enhancements apply to more than one component in OpenText Core Data Discovery & Risk Insights 26.3.0.

- Updated the application to use a single port for the UI and the API Gateway.
- All references to Extended ECM have been updated to reflect the new product name, Content Management.
- When an administration user sends an email verification message from Administration or a user sends an email verification to their own address from Analyze, Connect, or Manage, the UI now displays the message, "The verification email can only be sent once per minute." This rate limiting function provides added security.
- The following grammar updates are included in this release.
  - The PII TIN grammar now includes the 12 Digit Universal Account Number (UAN) for India (in). The landmarks are available in both English and Hindi scripts.
  - The PII address grammar for Bulgarian (bg) addresses now has improved recall for internal apartment details. This update improves the handling of different orderings and common format variations, to capture additional commonly used Bulgarian address writing styles. For example, apartment internal unit patterns can now be in any order.
  - The PII banking grammar IBAN matching for now accepts IBANs with any space and separator combinations, as long as the correct characters and total length are present. For example, GB82WEST12345698765432, GB82 WEST 1234 5698 7654 32, and GB82 WEST123456987654 32 all match.
- *Preview capability for cloud implementations*: Added an option to use AI to review the surrounding context of a matched grammar value to determine whether the match is a false-positive.

  The validation occurs during the dataset scan and therefore does not require additional post-processing. When the validation identifies a false-positive, the application follows the same workflow as manually marking a grammar value as a false-positive—recording the match, the Aviator verdict, and a system-generated reason, leaving an audit trail. The grammar value remains discoverable but is marked as an AI-assessed false positive.

  Aviator for grammar value validation is available for cloud implementations only and can be turned on for individual tenants. If activated, turn on **Use Aviator AI-assisted validation** for specific individual built-in grammar types. Then turn on this same option at the dataset level to use Aviator to validate matched values for the selected grammar types when scanning the selected datasets.

Processing Agent

The following features and enhancements apply specifically to the OpenText Core Data Discovery & Risk Insights 26.3.100 processing agent.

- Improved the performance of optical character recognition (OCR) processing.
  - For compound Office documents, only the item requiring OCR is processed. For example, you request OCR from a workbook that contains Word documents, a few of which have embedded images. Using this example, every document in the workbook would have been processed for OCR. With this improvement, only the embedded images are processed and OCR is skipped for documents that do not have embedded images.
  - Reduced the amount of overall data and memory needed to send items to Media Server for OCR. This reduces the pressure on the Media Server environment and possible timeouts during heavy OCR-processing timeframes.
  - Reduced the wait time before starting OCR jobs.
  - Improved log messaging for Media server activity to provide clearer workflow and file-level tracing. This includes improved checks to ensure canceled jobs are stopped earlier.
  - Added an advanced setting in the Agent Administration UI to control the number of concurrent analysis worker threads. During OCR processing, this is the number of workers on the agent host machine that are available to do the work and was previously fixed. "Analysis Worker Count" has been added to the Analysis advanced settings and defaults to 4.
    - Increase the thread count if you are experiencing an OCR backlog and you have sufficient CPU/memory and Media Server capacity.
    - Decrease the thread count if you are experiencing memory issues or you see timeouts and retries increase.

    > **NOTE:** Only adjust when necessary and in line with the capacity for Media Server. Adjust alongside infrastructure capacity and monitor throughput, completion time, and error/retry rates after changes. For further guidance, contact Support.
- Updated the internal connection to Exchange O365 to use Microsoft Graph instead of Exchange Web Services (EWS).

Connect

The following features and enhancements apply specifically to Connect, and are in addition to changes already mentioned in relation to general enhancements.

- Enhanced scheduling for agent cluster and dataset scans with more flexible options.

  You now have more flexibility to define when scheduled scans can and cannot run. Improvements focus on time zone-aware scheduling, flexible recurrence options, consolidated

and improved blackout management, as well as scheduling visibility at the dataset level. These options are available in the dataset template as well.

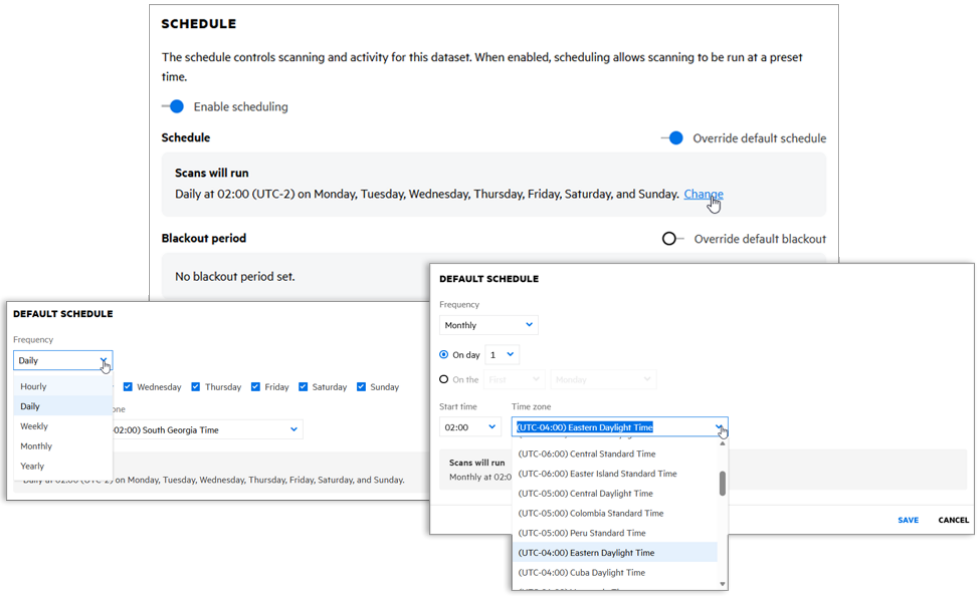

  - Schedule agent cluster or dataset scans to run hourly, daily, weekly, monthly, or yearly.
  - Define an agent cluster or dataset blackout period—when to not run scans—daily, weekly, monthly, or yearly.
  - For applicable frequencies, select the time zone for the schedule or blackout period. Previously, the defined time was relative to the server on which the associated agent is installed. The Time Zone defaults to the time zone of the browser you are using.
  - The option to **Run continuously with delay interval** has been removed. Agent clusters and datasets that used this option will have scheduling deselected after the upgrade to this release. To reinstate a schedule, edit the cluster or dataset and define a new schedule.
  - The blackout period configuration in the agent administration UI has been deprecated. The options exist in this release but have been deactivated. To define processing blackout periods, do so using the expanded options for agent clusters and datasets.

- Take action against multiple datasets more quickly.

  You can now perform scan reanalyze, enable scheduling, and disable scheduling actions against multiple datasets using the new action toolbar. Previously, you had to perform each of these actions against each dataset individually. The in-line action icons for datasets have been moved to an action toolbar that displays above the dataset list when at least one actionable dataset is selected. The action toolbar includes only the actions applicable for the selected datasets. You can perform a single action on up 500 datasets at a time.

  When disabling the schedule for multiple datasets, schedules set at the cluster level and the dataset level are disabled. When enabling the schedule for multiple datasets, keep the following in mind.

    - Datasets that do not have a dataset-level schedule defined use the schedule defined for the associated agent cluster.
    - Datasets that have a defined dataset-level schedule use that schedule, overriding the schedule (if set) for the associated agent cluster.
- Icons have been added to the agent activity list to let you quickly determine the datasets that have a dataset scan currently in progress () and that have grammar sets that have been updated since the last time the dataset was processed ().
- Moved the inline cancel and retry action icons in the agent activity list to an action toolbar above the agent activity list.
- Cancel multiple agent in-progress scan activities.

  For scan in-progress scan activities that can be canceled, you can now cancel multiple activities. Select one or more scan activities that can be canceled and then click the Cancel button in the action toolbar that displays above the agent activity list.
- Added the full file path to the agent activity error details.

Analyze

The following features and enhancements apply specifically to Analyze, and are in addition to changes already mentioned in relation to general enhancements.

- Improved the performance of the Risk Analysis Overview dashboard and the Financial Risk Modeler.

  The data for these dashboards is pre-computed on a schedule and stored. Dashboards then read the stored metrics instead of running live aggregation queries. This improves the retrieval and display of the dashboards.
- Added the ability to refresh the Risk Analysis Overview dashboard and the Financial Risk Modeler.

Manage

The following features and enhancements apply specifically to Manage, and are in addition to changes already mentioned in relation to general enhancements.

- Expanded the capabilities for masking at source from a workbook to protect more documents in their original source location.
    - In addition to file system and SharePoint, you can now mask supported file types from OneDrive datasets in their original OneDrive location.
    - When you protect documents in a workbook and select an internal masking rule (masking at source), you can now choose one of the following options.
        - Select **Mask in native format - unsupported file types will not be masked** to mask only the files that can be masked and support masking in their native format.

        - Select **Mask in native format with fallback to PDF** to mask files that can be masked and support masking in their native format. For files that cannot be masked in their native format and can be converted to PDF, convert to PDF and then mask the PDF.
        - Select **Mask and convert all supported files to PDF** to convert all files (that can be converted) to PDF and then mask the PDF.
    - Added dupport for masking `.json`, `.md`, and log files at source in their native format.
- Expanded the capabilities of sending masked documents to a target to align with the mask at source workflow.

  When you send masked documents from a workbook to a target, the new masking options added to the mask at source action have been added to the file placement selections (folder hierarchy or no folders). A similar pre-execution summary has also been added to let you preview the outcome before taking action.

API updates

The following updates have been made to the OpenText Core Data Discovery & Risk Insights APIs.

- Due to the scan scheduling improvements, the response data for the schedule and blackout properties has changed across the following methods.

  The `schedule` and `blackout` objects have changed for the following APIs.

  `GET /ajp/v1/job`

  `POST /ajp/v1/job`

  `PUT /ajp/v1/job/{jobId}`

  `GET /ajp/v1/agentcluster`

  `POST /ajp/v1/agentcluster`

  `PUT /ajp/v1/agentcluster/{clusterId}`

  `GET /cc/v1/repository-template`

  `POST /cc/v1/repository-template`

  `PUT /cc/v1/repository-template/{templateId}`

    - The legacy `schedule` object cron-based fields— `include`, `exclude`, `continuous`, and `delayInMin`—are no longer returned in API responses.
    - The `schedule` object now includes three structured fields—`recurrence` (a typed object with a type discriminator, such as WEEKLY, DAILY), `time` (HH:mm), and `timezone` (IANA timezone ID, such as America/New_York).
    - For datasets only, the `schedule` object also includes a `type` field (`INHERITED` or `CUSTOM`) that indicates whether the dataset abides the cluster's schedule or its own schedule.
    - The `blackout` property is now a separate top-level object on the response. Previously, blackout windows were expressed as a cron entry in the `exclude` field embedded inside the

`schedule` object.

- The `blackout` object uses the same typed `recurrence` structure as schedule, with additional fields for the blackout window—`startTime`, `endTime` (or `duration` + `durationUnit` for duration-based blackout windows), `timezone`, and optionally `endDayOfWeek` for cross-day windows.
- For datasets only, the `blackout` object also carries a `type` field (`INHERITED` or `CUSTOM`). A `blackout` object with `type: CUSTOM` and no `recurrence` field signals that the dataset explicitly suppresses the cluster's blackout.

For more information about the new schemas, see the descriptions in Swagger UI.

- Due to the scan scheduling improvements, `continuous: true` schedules are no longer supported.

  The following APIs are affected.

  `GET /ajp/v1/job`

  `POST /ajp/v1/job`

  `PUT /ajp/v1/job/{jobId}`

  `GET /ajp/v1/agentcluster`

  `POST /ajp/v1/agentcluster`

  `PUT /ajp/v1/agentcluster/{clusterId}`

  `GET /cc/v1/repository-template`

  `POST /cc/v1/repository-template`

  `PUT /cc/v1/repository-template/{templateId}`

  - Submitting `continuous: true` as a write value using the API is now rejected with HTTP 400.
  - All existing datasets and templates with `continuous: true` schedules are automatically migrated on upgrade. The cron schedule is converted to the equivalent typed `recurrence` (WEEKLY, DAILY), and `scheduleType` is changed to `MANUAL`. These entities will no longer trigger automated scans after upgrade. Review and reconfigure the affected datasets manually as needed.
  - For clusters, the `continuous` flag is dropped during migration. The associated cron schedule is converted to the equivalent typed recurrence and the cluster remains active.

Resolved issues for the application

The following issues have been resolved in this release of the OpenText Core Data Discovery & Risk Insights application.

- Resolved an issue where the on-screen progress percentage for an in-progress agent activity was different than the percentage that showed when you hover over the progress bar.
- [426686] Resolved an issue where a year value (YYYY) was being falsely matched as an Australian driver's license.
- Resolved an issue where the Root + Child Count did not display in the agent activity detail panel for Mask as Source activities.
- [3307509] Resolved an issue where documents labeled with certain Microsoft Purview labels were not marked as protected.
- Resolved an issue where the agent activity incorrectly reported a SharePoint Online activity as partially completed instead of invalid. In this scenario, the connection to SharePoint Online could not be established.
- Resolved an issue where the activity error window triggered from the agent activity detail pane may not show all of the related errors for an ingestion error.
- [427219] Resolved an issue where a "Failed to gather metadata for item" error occurred during ingestion. This occurred when an item was deleted from the original source location while the scan was in progress. This error is now prevented as the item no longer exists at the source.
- Resolved an issue where the status of a canceled mask at source activity was not displayed consistently across all UI pages.
- [427447] Resolved an issue where the Windows account and Exchange account information necessary for Exchange connection tasks were not fully documented.
- Resolved an issue where files with a large amount of context failed to be processed for tagging.
- [427797] Resolved an issue where a corrupted dataset resulted in the entire dataset list not being presented in Connect.
- [427924] Resolved an issue where a change was made to the grammar set configuration of an Exchange dataset, but no changes are made to the documents. For example, upon a rescan, grammar rule matches from the original scan remain despite the grammar set configuration had changed to remove those grammar rules.
- Resolved an issue where Exchange emails restored from an external source displayed the record date instead of the sent date.
- Resolved an issue where SharePoint items from a document library were being returned from an unexpected directory. This occurred when the SharePoint library name or title was defined as the Site Sub-URL instead of the expected Site URL.
- [426715] Resolved an issue where a workbook delete activity completes with errors, the workbook detail pane activity did not show the actual number of successfully deleted

documents out of the total number of documents.

- [428280] Resolved an issue where Microsoft Purview labeling processing would generate an exception if a document has an MIP label which doesn't contain the name field and the label name cannot be looked up in a local cache.
- [428779]Resolved an issue where the error message for an item deleted from the original source location indicated that it failed to be processed because it was modified. The correct message that the file cannot be found displays as expected and now includes the full original file path.
- The following resolved issues for grammars are included in this release.
    - [426686] Resolved an issue where a four digit year (YYYY) was incorrectly matched as an Australian driver license.

Resolved issues for the processing agent

The following issues were resolved in this release of the OpenText Core Data Discovery & Risk Insights processing agent.

- [425993] Resolved an issue where OCR requests where not paused if Media Server was down or otherwise not functioning as expected.
- [427370] Resolved an issue where SharePoint items that had been flagged by antivirus software failed processing but then were incorrectly retried. These items are now skipped during reprocessing.
- Resolved an issue where a workbook send masked activity would not progress if the workbook contained documents that had been deleted from the original source location.
- [427053] Resolved an issue where, during processing of Exchange items, documents that are no longer in their original source location (such as deleted or moved) were skipped and no log line item was generated. In this scenario, the document was still present in the content list pane (Analyze > Research) and Manage > Workspaces > Workbooks) but the body of the email was not viewable in the content view pane.
- [428281] Resolved an issue where the error message that resulted from an empty response from Content Manager was not clear. The error message has been updated as appropriate.

Documentation

OpenText Core Data Discovery & Risk Insightsincludes a single Help Center that is incorporated into each User Interface and is updated with each software release as appropriate.

To view the OpenText Core Data Discovery documentation outside of the product, see the documentation pages.

To navigate to the OpenText Core Data Discovery & Risk Insights documentation

1. 1. From a web browser, navigate to www.microfocus.com/documentation/file-analysis-suite/.
2. 2. On the OpenText Core Data Discovery & Risk Insights Documentation page, select the desired release version.

You must have Adobe® Reader installed to view files in PDF format (*.pdf). To download Adobe Reader, go to the Adobe web site.