RECOMMIND®

Axcelerate 5.9.1

Release Notes

RECOMMIND is now OPENTEXT™

Published: 2016-Nov-11

Contents

1 Introduction

This document lists the new features and resolved issues of the current and previous versions.

Features and resolved issues may be mentioned more than one time, as some of the previous releases refer only to Axcelerate Cloud or only to Axcelerate 5 onPremise.

2 New Feature Descriptions - Axcelerate 5.9.1

2.1 Business Intelligence Data Reload (Reference ID: AXC-36598)

Behavior before change

Data for the Business Intelligence dashboard was loaded every 4 hours. A manual reload was not available.

Behavior after change

In addition to the scheduled reload every 4 hours, the new **Reload Dashboard Data** button in the toolbar menu for any **Business Intelligence** dashboard can be used to load the latest data into the dashboard.

Known Limitations

None.

2.2 Current Criteria Panel Enhancements (Reference ID: AXC-37114)

Behavior before change

The applied criteria section in the **Search** panel only allowed to lock and delete applied criteria. To display the specific values, and to edit, the panel needed to be maximized. The maximized height took less than 50% of the **Search** panel, and allowed to display <~10 different criteria without scrolling.

Behavior after change

The applied criteria section in the **Search** panel allows to display, lock, edit and delete applied criteria, without having to maximize the panel height. The maximized height has been increased to 50% of the **Search** panel, allowing to display >~10 different criteria without scrolling.

Known Limitations

None.

2.3 "Responsive" Field Renamed (Reference ID: AXC-37119)

Behavior before change

For new matters, the default field to track responsiveness of a document was named **Responsive** with **yes** and **no** values.

Behavior after change

For new matters, the default field to track responsiveness of a document is named **Responsiveness** with **Responsive** and **Not responsive** values.

Changes to Default Configuration

The default system template is now using the new display name **Responsiveness** for the `ax_responsive` field, and loads the default values **Responsive** and **Not responsive** from a new OWL file.
The OWL file with **yes** and **no** values still exists, but is no longer used by the current template.

Known Limitations

None.

2.4 U Quick Tag (Reference ID: AXC-32949)

Behavior before change

The **L** button applied the last complete tagging.

Behavior after change

The **L** button continues to apply the last complete tagging.
A new **U** button can be used to apply only the last changes to tagging.

Changes to Default Configuration

None.

Known Limitations

None.

2.5 New (NAS) Storage Handling (Reference ID: CORE-16209)

Behavior before change

NAS based projects used storage handlers, called structured storage handler, to store and access documents, e.g. natives or images, from directory based locations.

Behavior after change

There is now a new (NAS-) storage handler, called structured storage V2, enabling compression and case-wide deduplication and thus enhancing the existing structured storage handler. It truly uses the SQL-managed storage locations, including billing support. The

handler also allows for multiple locations, e.g. drives or NASs, in order to add storage space in case of full disks. This handler is also the default for newly created projects. Furthermore, a workflow is supplied that allows the user to convert an existing NAS storage into a new NAS storage with different configuration parameters for storage database, compression, and deduplication.

Changes to Default Configuration

The structured storage V2 is the default storage for new projects.
The V2-handler comes with its own set of configuration options, most of them are equal to other handlers (i.e. structured storage), except for the storage roots. These storage roots are used to configure the directories where the objects are saved (known as storage location prefix for the structured storage), and it is possible to give more than one directory where one of these has to be designated as "master".
In order to enable case-wide deduplication, the handler will make use of the case name and client id as configured in the application configuration (or master service as fallback).

Known Limitations

The new setup does not support "matter storage".

2.6 Performance Measures (Reference ID: CORE-16714)

Behavior before change

Aggregated performance figures were not collected or reported.

Behavior after change

Performance measures are collected, aggregated and periodically reported to log file, a separate csv file, and via event notification to the Axcelerate 5 reporting database. This provides the ability to pro-actively detect high load situations and provides insight into root causes for high load.

Changes to Default Configuration

Four parameters have been added to the data model configuration under **Common > Event Notification and KPI Reporting** to control the behavior of the collection and generated reports. The added parameters are:

- Active-checkbox (default: on)
- Interval for aggregation and reporting (default: 5 minutes)
- CSV-filename (default: "kpi.csv" in engine-directory)
- Checkbox for reporting database (default: on)

Known Limitations

None.

2.7 Transym OCR Support (Reference ID: CORE-15448)

Behavior before change

Latest supported Transym version was 4.0.0.9.

Behavior after change

Latest supported Transym version is 4.0.0.14.

Known Limitations

None.

3 Issues Resolved in Axcelerate 5.9.1

3.1 AXC-34630 (Foldering panel)

Summary	New Behavior
Partially corrupted field values caused the retrieval of folders in the **Foldering** panel to fail, rendering the complete foldering feature unusable.	The corrupted part of the field values is ignored and the field value is listed in the **Foldering** panel. Functionality is not impacted by the missing part. The incident, including the ID and name of the corrupted field value, is logged in the backend.

3.2 AXC-36584 (Review Workflows page)

Summary	New Behavior
Review Workflows page was not accessible if internal field value properties were missing.	**Review Workflows** page is accessible as the internal field value and its properties are inserted in the same transaction.

3.3 AXC-36619 (Printing)

Summary	New Behavior
Printing with the **Document Overlay** option could not be started, despite having a valid text overlay.	Printing with the **Document Overlay** option and a valid text overlay now works.

3.4 AXC-36706 (Permission synchronization)

Summary	New Behavior
Group permissions could not be synchronized when there was a locked application.	Group permissions can be synchronized if there is a locked application.

3.5 AXC-36811 (Bloomberg chat)

Summary	New Behavior
For Bloomberg chat documents containing invitations with messages, the invitations with messages were not shown in **Text** view.	**Text** view for Bloomberg chat documents now displays both regular messages as well as invitations with messages.

3.6 AXC-37033 (Production workflows)

Summary	New Behavior
Values were not appearing in the **Applied Rule** Smart Filter in production workflow preview mode, even though the **Applied Rule** column of the **Results** list displayed the values.	Values appear as expected in the **Applied Rule** Smart Filter while in production workflow preview mode.

3.7 AXC-37061 (Business Intelligence)

Summary	New Behavior
Billing metrics were not updated and new Business Intelligence dashboards could not be created when there was a meta-engine project in the pod.	Billing metrics are updated and Business Intelligence dashboards can be created even when there is a meta-engine project in the pod.

3.8 AXC-37090 (Pods with high number of applications)

Summary	New Behavior
Reviewer efficiency, creation of new Business Intelligence dashboards, adaptive batching, automatic set creation, automatic iteration creation and pre-conversion rules did not work for some projects when there were more than 50 applications in the pod.	Reviewer efficiency, creation of new Business Intelligence dashboards, adaptive batching, automatic set creation, automatic iteration creation and pre-conversion rules should work with any number of applications.

3.9 AXC-37264 (Metadata panel)

Summary	New Behavior
Fields with a hyphen in the ID (e.g. **Message-Id)** were not parsed correctly and could not be used in the **Metadata** panel.	Fields with a hyphen in the ID (e.g. **Message-Id**) are parsed correctly and can be used in the **Metadata** panel.

3.10 AXC-37270 (Performance)

Summary	New Behavior
The system tried to get display names for all fields that were listed in the **View - modifiable fields** table of the document model in CORE Administration, even for fields that were not needed for Axcelerate 5, thereby decreasing performance.	Now only fields that are used in either field view group named `axcelerate.tagging` or a`xcelerate.foldering` are synchronized (plus service tier extra list of other important Axcelerate 5 fields). The performance of delete calls and update calls is increased.

3.11 AXC-37397 (Tagging rules)

Summary	New Behavior
The backend servers maintain a local copy of all tagging rules. When one server deleted a rule from the database, the other servers did not adapt their local copies and continued to use the outdated rule for display as well as validation.	The backend servers still maintain a local copy of all tagging rules. The synchronization process from the database to these local copies now supports deletion.

3.12 AXC-37398 (Pre-Conversion)

Summary	New Behavior
A pre-conversion rule was running in continuous loop for a corrupted document and lead to a newly created job every 5 minutes.	A pre-conversion rule is no longer running in continuous loop for a corrupted document.

3.13 AXC-37582 (Production)

Summary	New Behavior
Documents were not tagged as production exceptions in case of a production error (e.g. Conversion pages was not set).	Now documents are tagged as production exceptions in case of a production error (e.g. Conversion pages was not set).

3.14 AXC-37597 (Performance)

Summary	New Behavior
Concurrent tagging requests from multiple reviewers caused the reload of temporarily cached tagging rules to be triggered multiple times (once for each tagging request). As these processes were queued and not allowed to be run in parallel, some reviewers had to wait very long (up to timeouts in the browser) until their individual process was run. Because this process is tied to the tagging action itself, their review performance decreased.	The reload processes are still queued in a way, but subsequent processes will now detect that a previous process already reloaded the tagging rules and then stop immediately. Thus, while all reviewers will still have to wait a short time for one (i.e. the first) process to complete, they don't have to wait for multiple, redundant reloads anymore.

3.15 CORE-15304 (Redaction file references)

Summary	New Behavior
When re-publishing a redacted document with images, the redaction files were deleted but still referenced in the xml.	When re-publishing a redacted document with images, not only the redaction files are deleted but also the references in the xml.

3.16 CORE-15992 (Property Post Processors)

Summary	New Behavior
When the `recipient_counter` was not last on the Property Post Processors (PPP) list for the data source, the recipient counter field was not populated with the correct count.	The `recipient_counter` is now last on the Property Post Processors (PPP) list for the data source. When the script `updateLegacyApplication.bat` runs as part of the upgrade procedure, it makes sure `recipient_counter` is last on the PPP list.

3.17 CORE-16389 (Documentum crawls)

Summary	New Behavior
Documentum crawls with custom query but without folder or cabinet in URI were not supported.	Documentum crawls with custom query but without folder or cabinet in URI are supported.

3.18 CORE-16439 (Wildcard expansion)

Summary	New Behavior
The data structure used for wildcard expansion was not always updated properly causing missing entries.	The sorted set data structure is now updated properly in all cases. Affected engines are automatically repaired during engine startup. An E2 message is logged if missing entries are found.

3.19 CORE-16868 (Storage utilization statistics)

Summary	New Behavior
CORE storage collection for PDF and EXPORT used hard-coded directories and thus reported incorrect storage utilization statistics.	COREore storage collection for PDF and EXPORT now use configured directories and storage utilization statistics are correctly reported.

3.20 CORE-16975 (Box connector)

Summary	New Behavior
When a document in Box had versions and no existing creator, no documents were indexed for the versions of that document.	Versions with no existing creator in Box are now indexed.

3.21 CORE-17050 (OCR jobs)

Summary	New Behavior
Canceling OCR jobs resulted in lots of temp files (.tif) in the temp folder on the engine host.	Temp files produced by canceled OCR jobs are now cleaned up properly.

3.22 CORE-17106 (Bloomberg data)

Summary	New Behavior
When multiple Bloomberg TXT files were found underneath a TXT directory only the first one was used. In addition, TXT files with an infix after the .[I]B.startdate-enddate section and between the TXT file suffix were skipped.	Bloomberg TXT files with an arbitrary infix after the .[I]B.startdate-enddate section are now considered. In addition, support has been added for subdirectories underneath a TXT parent directory with more than a single TXT file. Subfolders containing mixed Bloomberg messages and Bloomberg instant messages with separate TXT files are also considered.

3.23 CORE-17124 (Documentum crawls)

Summary	New Behavior
When crawling a Documentum source, the native storage failed for documents without content description.	When crawling a Documentum source, native storage succeeds for documents without content description.

3.24 CORE-17133 (Performance)

Summary	New Behavior
An internal warning message was displayed too often causing the engine to become unresponsive.	The logging frequency of the internal warning message has been reduced, thereby fixing the performance issue.

3.25 CORE-17284 (Performance)

Summary	New Behavior
A meta engine unnecessarily computed regex based highlighting expressions, leading to unnecessary memory consumption that could lead to out of memory situations.	A meta engine does not compute regex based highlighting expressions as they are not needed in this case.

3.26 CORE-17312 (OCR jobs)

Summary	New Behavior
Results of an OCR job were stored locally in a processing unit specific temp directory that was cleared when the respective processor crashed, resulting in a FileNotFound Exception for the respective OCR document.	Results of an OCR job are now stored in a temp directory of the respective engine such that a crash of the OCR processor does not wipe out existing results.

3.27 CORE-17316 (Stored page count)

Summary	New Behavior
Changing a textual tagging field (either in the user interface or using a CSV merge) caused the stored page count to be removed.	The stored page count is not affected anymore when other fields are changed.

3.28 CORE-17382 (Conversion)

Summary	New Behavior
Documents with failed native conversion were not always tagged as an exception and not shown in the Conversion Smart Filter.	Documents with failed native conversion are now always tagged as an exception.

3.29 CORE-17458 (Wordmap)

Summary	New Behavior
Very long CJK field values could cause a buffer overflow in the WordMap potentially causing the engine process to crash.	The internal buffer in the WordMap has been increased to prevent it from overflowing. In addition, the maximum length of a field value has been reduced.

3.30 CORE-17567 (Fetching files)

Summary	New Behavior
The system couldn't perform all file fetch operations on faulty projects which contain compressed and uncompressed data.	The system can fetch both compressed and uncompressed files within a faulty project.

3.31 CORE-17592 (Bloomberg)

Summary	New Behavior
Attachments in Bloomberg chat text files that contained any of the file ID delimiters "[" or "]" in their display name were not processed correctly and caused the document to be tagged with an **Archive entry processing error**	Attachments in Bloomberg chat text files that contain any of the file ID delimiters "[" or "]" in their display name are now properly processed.

3.32 CORE-17597 (Bloomberg)

Summary	New Behavior
Duplicate user information in the header of Bloomberg text was causing documents to be tagged with the **Miscellaneous parser error - already seen mail address** error.	A duplicate user information (with partially incomplete data) will now be mapped to its more complete counterpart during parsing of a Bloomberg text header. The error **Miscellaneous parser error - already seen mail address** will still be tagged to documents which are associated to Bloomberg text files that contain more than one user for the same email address.

3.33 CORE-17598 (Bloomberg)

Summary	New Behavior
When the pipe ("|") character was part of a user data attribute in a Bloomberg text file the corresponding chat actions could not be resolved and caused the document to be tagged with **Miscellaneous parser error - unable to resolve user** error.	Bloomberg chat with users which contain the pipe ("|") character in their name is properly parsed now.

3.34 CORE-17639 (Tagging)

Summary	New Behavior
Executing a bulk tagging while a crawler was running could cause a deadlock in the index engine.	There is no deadlock when executing a bulk tagging while a crawler is running.

3.35 CORE-17683 (Memory)

Summary	New Behavior
There was a memory leak when executing the same search multiple times with the same session.	The memory leak when executing the same search multiple times with the same session is fixed.

3.36 CORE-17757 (Batch server status detection)

Summary	New Behavior
Jobs were sent to batch servers that were in some undefined offline state. The resulting exception caused the job to fail but did not refresh the online/offline status.	The corresponding exception causes the host to be treated as offline. Furthermore, logging is improved to enable better tracking of similar issues in the future.

3.37 CORE-17781 (Index engine write lock)

Summary	New Behavior
An exception thrown during an index engine save could cause subsequent write lock attempts to fail.	Internal save counter is reduced in all cases now, even if an exception occurs, and should not cause write lock attempts to fail.

3.38 CORE-17788 (Categorization state display)

Summary	New Behavior
For hierarchical fields the displayed categorization state (**Agreed**, **Disagreed**, **Suggested**) was wrong.	The categorization state is correctly shown for both flat and hierarchical fields.

3.39 CORE-18013 (Document history)

Summary	New Behavior
If text field values were changed using certain scripts, the respective entry in the **Document history** caused an exception in Axcelerate 5.	If text field values are changed using scripts, the respective entry in the **Document history** is displayed.

3.40 FOUND-7160 (Log messages)

Summary	New Behavior
Some log messages were formatted with log4j (WARN/FATAL) priority.	Log messages are formatted without log4j priority.

3.41 FOUND-9536 (Application removal)

Summary	New Behavior
The removal of an application tree with a meta engine lead to error messages in the MasterService log (without any further impact).	The removal of an application tree containing a meta engine has been revised. Error messages with no impact are not created in the Master Service log.

3.42 FOUND-9588 (0 byte file handling)

Summary	New Behavior
NIST filter was applied for empty files.	NIST filter is not applied for empty files.

3.43 FOUND-9644 (PDF date parsing)

Summary	New Behavior
Some formats of the modification date of comments in PDF documents could not be parsed.	Modification dates of comments in PDF documents get parsed correctly.

3.44 FOUND-9712 (Performance)

Summary	New Behavior
For the ECA/Ingestion application, the **Matters** tab loading speed was low with huge numbers of matters.	For the ECA/Ingestion application it is now possible to filter for last updated matters in order to increase the loading speed of the **Matters** tab.

3.45 FOUND-9951 (Performance)

Summary	New Behavior
A large number (> 250k) of temp files from a terminated process caused an out of memory error in the launcher service while deleting those files.	The launcher service is able to clean up an arbitrary number of temp files after a process has terminated.

3.46 FOUND-10039 (Performance)

Summary	New Behavior
Oracle OutsideIn parser service created and left temp files in user temp directory.	Oracle OutsideIn parser service creates temp files in data source temp directory which is deleted automatically.

3.47 FOUND-10073 (Kerberos and Active Directory login)

Summary	New Behavior
The Kerberos and ActiveDirectory login modules used different formats for server lists, leading to errors when copying and pasting from one module to another.	The Kerberos and ActiveDirectory login modules now use the same server list format so that a value can be copied and pasted from one module to another.

3.48 FOUND-10099 (Data source start)

Summary	New Behavior
Data sources could no longer be started if they had been automatically resumed a couple of times before.	Data sources can be started independently of the number of previous automatic resumes.

3.49 FOUND-10104 (Application case name)

Summary	New Behavior
New Application wizard truncated *Case Name* property at the display name's forward slash.	**New Application** wizard no longer truncates *Case Name* property at the display name's forward slash.

3.50 FOUND-10177 (Configuration access)

Summary	New Behavior
In CORE Administration concurrent access by different processes to default configuration could fail.	In CORE Administration concurrent access by different processes to default configurations is now supported.

3.51 FOUND-10214 (SSO login)

Summary	New Behavior
Login via SSO could cause the user session to terminate.	Login via SSO works as expected.

3.52 FOUND-10291 (Exchange connector)

Summary	New Behavior
When using the Microsoft Exchange connector, invalid characters in Microsoft Exchange webservice responses could cause a crawl failure.	Invalid characters in Microsoft Exchange webservice responses are filtered and do not cause a crawl failure.

3.53 FOUND-10308 (Monitoring)

Summary	New Behavior
In the **Monitoring** section of CORE Administration, the **HOST ID** column on the **Data Sources** tab showed the host's display name.	The **HOST ID** column on the **Data Sources** tab shows the host ID.

3.54 FOUND-10466 (Box data sources)

Summary	New Behavior
During Box data source crawls memory consumption for connections to Box was proportional to the overall number of users to be crawled. For high number of users this could lead to Out-Of-Memory issues.	Memory consumption for connections is now proportional to the number of users that are currently being processed. This number should be smaller or equal to the number of threads used for crawling. For crawls in which the overall number of users is high, this results in a considerably smaller memory footprint, strongly decreasing the chance for a Out-Of-Memory issues.

4 New Feature Descriptions for Axcelerate 5.9.0

4.1 Business Intelligence Data Reload (Reference ID: AXC-36598)

Behavior before change

Data for the Business Intelligence dashboard was loaded every 4 hours. A manual reload was not available.

Behavior after change

In addition to the scheduled reload every 4 hours, the new **Reload Dashboard Data** button in the toolbar menu for any **Business Intelligence** dashboard can be used to load the latest data into the dashboard.

Known Limitations

None.

4.2 Current Criteria Panel Enhancements (Reference ID: AXC-37114)

Behavior before change

The applied criteria section in the **Search** panel only allowed to lock and delete applied criteria. To display the specific values, and to edit, the panel needed to be maximized. The maximized height took less than 50% of the **Search** panel, and allowed to display <~10 different criteria without scrolling.

Behavior after change

The applied criteria section in the **Search** panel allows to display, lock, edit and delete applied criteria, without having to maximize the panel height. The maximized height has been increased to 50% of the **Search** panel, allowing to display >~10 different criteria without scrolling.

Known Limitations

None.

4.3 "Responsive" Field Renamed (Reference ID: AXC-37119)

Behavior before change

For new matters, the default field to track responsiveness of a document was named **Responsive** with **yes** and **no** values.

Behavior after change

For new matters, the default field to track responsiveness of a document is named **Responsiveness** with **Responsive** and **Not responsive** values.

Changes to Default Configuration

The default system template is now using the new display name **Responsiveness** for the `ax_responsive` field, and loads the default values **Responsive** and **Not responsive** from a new OWL file.
The OWL file with **yes** and **no** values still exists, but is no longer used by the current template.

Known Limitations

None.

4.4 U Quick Tag (Reference ID: AXC-32949)

Behavior before change

The **L** button applied the last complete tagging.

Behavior after change

The **L** button continues to apply the last complete tagging.
A new **U** button can be used to apply only the last changes to tagging.

Changes to Default Configuration

None.

Known Limitations

None.

4.5 Performance Measures (Reference ID: CORE-16714)

Behavior before change

Aggregated performance figures were not collected or reported.

Behavior after change

Performance measures are collected, aggregated and periodically reported to log file, a separate csv file, and via event notification to the Axcelerate 5 reporting database. This provides the ability to pro-actively detect high load situations and provides insight into root causes for high load.

Changes to Default Configuration

Four parameters have been added to the data model configuration under **Common > Event Notification and KPI Reporting** to control the behavior of the collection and generated reports. The added parameters are:

- Active-checkbox (default: on)
- Interval for aggregation and reporting (default: 5 minutes)
- CSV-filename (default: "kpi.csv" in engine-directory)
- Checkbox for reporting database (default: on)

Known Limitations

None.

4.6 Transym OCR Support (Reference ID: CORE-15448)

Behavior before change

Latest supported Transym version was 4.0.0.9.

Behavior after change

Latest supported Transym version is 4.0.0.14.

Known Limitations

None.

4.7 SAML 1.1 AuthN in Axcelerate 5 (Reference ID: Found-9480)

Behavior before change

SAML authentication was not supported for Axcelerate 5.

Behavior after change

SAML 1.1 authentication is now supported for Axcelerate 5.

Changes to Default Configuration

None.

Known Limitations

None.

5 Issues Resolved in Axcelerate 5.9.0

5.1 AXC-30193 (Production workflows)

Summary	New Behavior
When adding documents to a production workflow while many taggings are still pending (long coding queue), the user receives a timeout error though the actual operation is still running in the background and typically completes successfully at a later time.	The "add to production workflow" action finishes immediately and the actual operation continues to run in the background.

5.2 AXC-34630 (Foldering panel)

Summary	New Behavior
Partially corrupted field values caused the retrieval of folders in the **Foldering** panel to fail, rendering the complete foldering feature unusable.	The corrupted part of the field values is ignored and the field value is listed in the **Foldering** panel. Functionality is not impacted by the missing part. The incident, including the ID and name of the corrupted field value, is logged in the backend.

5.3 AXC-35769 (Review workflows)

Summary	New Behavior
Calls to retrieve document counts on **Review Workflow** pages for large review workflows caused timeouts and heavy server loads.	Sequential calls to retrieve documents counts on **Review Workflow** pages do not cause timeouts and heavy server loads.

5.4 AXC-36156 (Navigation page)

Summary	New Behavior
When paging to a page other than the first one of the **Navigation** page (matter list), the user was not able to log in to a matter.	When paging to another page than the first one of the **Navigation** page (matter list), the user will see all available matters and is able to log in to each matter.

5.5 AXC-36439 (Highlighting)

Summary	New Behavior
When reviewing documents in two-screen mode, the highlight summary button was missing.	The highlight summary button is shown in both single-screen and two-screen mode on the **Analysis** and **Review** pages.

5.6 AXC-36584 (Review workflows)

Summary	New Behavior
Review Workflow page was not accessible if internal field value properties were missing.	**Review Workflow** page is accessible as the internal field value and its properties are inserted in the same transaction.

5.7 AXC-36619 (Printing)

Summary	New Behavior
Printing with the **Document Overlay** option could not be started, despite having a valid text overlay.	Printing with the **Document Overlay** option and a valid text overlay now works.

5.8 AXC-36681 (Document access)

Summary	New Behavior
When a filter was applied which invalidated a batch restriction, a reviewer could potentially access restricted documents.	A reviewer can only see documents assigned to their currently selected batch.

5.9 AXC-36786 (Jobs page performance)

Summary	New Behavior
When many jobs were present on the **Jobs** page, it could potentially take too long to load and become inaccessible to the user.	**Jobs** page loads approximately 5-10 times faster and does not become inaccessible when many jobs are present.

5.10 AXC-36806 (Navigation page)

Summary	New Behavior
With large amounts of groups and/or applications and with multiple service tiers, timeouts and optimistic locking exceptions could occur. As a consequence, updates to the **Navigation** page matter list were not executed.	Even with large amount of groups and applications, the sync will succeed and update the **Navigation** page matter list. Multiple service tiers instances will not sync in parallel and thus no optimistic locking exceptions are logged.

5.11 AXC-36811 (Bloomberg chat)

Summary	New Behavior
For Bloomberg chat documents containing invitations with messages, the invitations with messages were not shown in **Text** view.	**Text** view for Bloomberg chat documents now displays both regular messages as well as invitations with messages.

5.12 AXC-36899 (Security vulnerability in 7zip 9.20)

Summary	New Behavior
7zip version shipped with Axcelerate 5 was 9.20, which had known security vulnerabilities.	7zip version is 16.0, which fixes the known security vulnerabilities.

5.13 AXC-37033 (Production workflows)

Summary	New Behavior
Values were not appearing in the **Applied Rule** Smart Filter in production workflow preview mode, even though the **Applied Rule** column of the **Results** list displayed the values.	Values appear as expected in the **Applied Rule** Smart Filter while in production workflow preview mode.

5.14 AXC-37061 (Business Intelligence)

Summary	New Behavior
Billing metrics were not updated and new Business Intelligence dashboards could not be created when there was a meta-engine project in the pod.	Billing metrics are updated and Business Intelligence dashboards can be created even when there is a meta-engine project in the pod.

5.15 AXC-37095 (Review batches)

Summary	New Behavior
Marking batches as complete calculated the review count for all batches in the project instead of only the selected ones, thereby reducing the execution speed of the operation and resulting in browser timeouts.	Review count is now calculated only for the selected batches rather than all batches, significantly improving the execution speed of the operation and avoiding unnecessary browser timeouts.

5.16 AXC-37264 (Metadata panel)

Summary	New Behavior
Fields with a hyphen in the ID (e.g. **Message-Id)** were not parsed correctly and could not be used in the **Metadata** panel.	Fields with a hyphen in the ID (e.g. **Message-Id**) are parsed correctly and can be used in the **Metadata** panel.

5.17 AXC-37397 (Tagging rules)

Summary	New Behavior
The backend servers maintain a local copy of all tagging rules. When one server deleted a rule from the database, the other servers did not adapt their local copies and continued to use the outdated rule for display as well as validation.	The backend servers still maintain a local copy of all tagging rules. The synchronization process from the database to these local copies now supports deletion.

5.18 CORE-15070 (Engine starts)

Summary	New Behavior
A defect in the 3rd party software Liquibase could result in a failure to release SQL locks which blocked all engines during the startup procedure.	The schemas are now updated during the software update, avoiding blocked engines during the startup procedure.

5.19 CORE-15749 (Bloomberg email thread detection)

Summary	New Behavior
Thread detection of larger Bloomberg emails could cause data sources and engines to run out of memory.	Thread detection of larger Bloomberg emails does not cause data sources and engines to run out of memory.

5.20 CORE-15992 (Property Post Processors)

Summary	New Behavior
When the `recipient_counter` was not last on the Property Post Processors (PPP) list for the data source, the recipient counter field was not populated with the correct count.	The `recipient_counter` is now last on the Property Post Processors (PPP) list for the data source. When the script `updateLegacyApplication.bat` runs as part of the upgrade procedure, it makes sure `recipient_counter` is last on the PPP list.

5.21 CORE-16008 (Engine save)

Summary	New Behavior
A `rules.xml` containing rules referencing non-existing field values could cause a deadlock during an engine save.	A `rules.xml` containing rules referencing non-existing field values no longer causes a deadlock during an engine save.

5.22 CORE-16439 (Wildcard expansion)

Summary	New Behavior
The data structure used for wildcard expansion was not always updated properly causing missing entries.	The sorted set data structure is now updated properly in all cases. Affected engines are automatically repaired during engine startup. An E2 message is logged if missing entries are found.

5.23 CORE-16460 (Engine save)

Summary	New Behavior
Triggering both a manual shutdown and an emergency shutdown could lead to a failing save job.	Triggering both a manual shutdown and an emergency shutdown no longer leads to a failing save job.

5.24 CORE-16576 (Performance)

Summary	New Behavior
Update of engine statistics could become stuck if engine was very busy.	The updating of engine statistics performance under load has been improved.

5.25 CORE-16613 (Field value deletion)

Summary	New Behavior
Deleting a huge amount of field values in bulk caused a delay in engine startup.	Deleting a huge amount of field values in bulk no longer causes a delay in engine startup.

5.26 CORE-16620 (Search)

Summary	New Behavior
Serializing huge search requests to index partitions resulted in out-of-memory situations since the requests did not fit into the smaller java heap.	Search requests are only sent to the index partitions if there is enough java heap left; otherwise, they are held back until another search returns.

5.27 CORE-16681 (Junk detection)

Summary	New Behavior
Junk detection occasionally produced invalid document XML while replacing junk with replacement text.	New documents always have valid XML, and old documents with invalid XML are repaired during publish and document display.

5.28 CORE-16754 (Training data jobs)

Summary	New Behavior
The flush post processing job could be blocked by potentially long running training data job. As a consequence, concurrent read operations (i.e., document fetching and searches) could be blocked as well.	The flush post processing job is not blocked anymore by training data jobs. In addition, the training data jobs are running much faster now if the *small flush* feature is active.

5.29 CORE-16755 (Coding queue performance)

Summary	New Behavior
Rules processing had a large RMI overhead, causing a negative impact on the coding queue performance.	The RMI overhead of rules processing is improved, reducing the negative impact on the coding queue performance.

5.30 CORE-16765 (CSV merge)

Summary	New Behavior
When using CSV merge on a project that is prepared for storage migration but not yet migrated, if the merge replaced the existing images with fewer images, some old images remained in the storage.	CSV merge no longer leads to orphaned images in storage.

5.31 CORE-16777 (Performance)

Summary	New Behavior
Some index calls involving the field value existence checks in merging meta engines were slow.	Performance of the field value existence checks in merging meta engines is improved.

5.32 CORE-16868 (Storage utilization statistics)

Summary	New Behavior
CORE storage collection for PDF and EXPORT used hard-coded directories and thus reported incorrect storage utilization statistics.	COREore storage collection for PDF and EXPORT now use configured directories and storage utilization statistics are correctly reported.

5.33 CORE-16932 (Performance)

Summary	New Behavior
Conversions with multiple binary records (especially image views) had long latency if binary records were served from S3.	Multiple streams are now downloaded concurrently, thereby reducing the latency.

5.34 CORE-16975 (Box connector)

Summary	New Behavior
When a document in Box had versions and no existing creator, no documents were indexed for the versions of that document.	Versions with no existing creator in Box are now indexed.

5.35 CORE-17050 (OCR jobs)

Summary	New Behavior
Canceling OCR jobs resulted in lots of temp files (.tif) in the temp folder on the engine host.	Temp files produced by canceled OCR jobs are now cleaned up properly.

5.36 CORE-17106 (Bloomberg data)

Summary	New Behavior
When multiple Bloomberg TXT files were found underneath a TXT directory only the first one was used. In addition, TXT files with an infix after the .[I]B.startdate-enddate section and between the TXT file suffix were skipped.	Bloomberg TXT files with an arbitrary infix after the .[I]B.startdate-enddate section are now considered. In addition, support has been added for subdirectories underneath a TXT parent directory with more than a single TXT file. Subfolders containing mixed Bloomberg messages and Bloomberg instant messages with separate TXT files are also considered.

5.37 CORE-17133 (Performance)

Summary	New Behavior
An internal warning message was displayed too often causing the engine to become unresponsive.	The logging frequency of the internal warning message has been reduced, thereby fixing the performance issue.

5.38 CORE-17187 (Storage size calculation)

Summary	New Behavior
Projects with storage compression populated the `rm_nativestoragesize` field with the compressed sizes instead of the file size. This happened during crawl operations.	`rm_nativestoragesize` is populated with the "managed size," i.e., normal file size. For existing projects, this is automatically corrected at the next engine restart.

5.39 CORE-17220 (Native files)

Summary	New Behavior
The publish wizard configured new publish engines incorrectly for SaaS-based S3 projects, resulting in inaccessible published natives.	The publish engines are configured correctly and the published natives are accessible.

5.40 CORE-17284 (Performance)

Summary	New Behavior
A meta engine unnecessarily computed regex based highlighting expressions, leading to unnecessary memory consumption that could lead to out of memory situations.	A meta engine does not compute regex based highlighting expressions as they are not needed in this case.

5.41 CORE-17312 (OCR jobs)

Summary	New Behavior
Results of an OCR job were stored locally in a processing unit specific temp directory that was cleared when the respective processor crashed, resulting in a FileNotFound Exception for the respective OCR document.	Results of an OCR job are now stored in a temp directory of the respective engine such that a crash of the OCR processor does not wipe out existing results.

5.42 CORE-17316 (Stored page count)

Summary	New Behavior
Changing a textual tagging field (either in the user interface or using a CSV merge) caused the stored page count to be removed.	The stored page count is not affected anymore when other fields are changed.

5.43 CORE-17449 (Storage SQL feature)

Summary	New Behavior
The storage SQL feature could be disabled by the master switch even if the project was not empty.	The storage SQL feature cannot be disabled by the master switch if the project is not empty.

5.44 FOUND-7160 (Log messages)

Summary	New Behavior
Some log messages were formatted with log4j (WARN/FATAL) priority.	Log messages are formatted without log4j priority.

5.45 FOUND-9588 (0 byte file handling)

Summary	New Behavior
NIST filter was applied for empty files.	NIST filter is not applied for empty files.

5.46 FOUND-9644 (PDF date parsing)

Summary	New Behavior
Some formats of the modification date of comments in PDF documents could not be parsed.	Modification dates of comments in PDF documents get parsed correctly.

5.47 FOUND-9712 (Performance)

Summary	New Behavior
For the ECA/Ingestion application, the **Matters** tab loading speed was low with huge numbers of matters.	For the ECA/Ingestion application it is now possible to filter for last updated matters in order to increase the loading speed of the **Matters** tab.

5.48 FOUND-9951 (Performance)

Summary	New Behavior
A large number (> 250k) of temp files from a terminated process caused an out of memory error in the launcher service while deleting those files.	The launcher service is able to clean up an arbitrary number of temp files after a process has terminated.

5.49 FOUND-10039 (Performance)

Summary	New Behavior
Oracle OutsideIn parser service created and left temp files in user temp directory.	Oracle OutsideIn parser service creates temp files in data source temp directory which is deleted automatically.

5.50 FOUND-10102 (Performance)

Summary	New Behavior
Large patch files deployed for a large number of processes caused the master service to run out of memory when a patch report was created.	Memory utilization in the master service is optimized while generating a large patch report.

5.51 FOUND-10104 (Application case name)

Summary	New Behavior
New Application wizard truncated *Case Name* property at the display name's forward slash.	**New Application** wizard no longer truncates *Case Name* property at the display name's forward slash.

6 New Feature Descriptions in version 5.7.0 Update 1

6.1 JAVA update to JDK 8 Update 92 (Reference ID: Found-9744)

Behavior before change

JDK bundled with the CORE package was a JAVA previous version.

Behavior after change

JDK version 8 Update 92 is bundled with the CORE package. Please reference the Oracle release notes for further information:[http://www.oracle.com/technetwork/java/javase/8u92-relnotes-2949471.html].

Changes to Default Configuration

None.

Known Limitations

None.

6.2 SAML 1.1 AuthN in Axcelerate 5 (Reference ID: Found-9480)

Behavior before change

SAML authentication was not supported for Axcelerate 5.

Behavior after change

SAML 1.1 authentication is now supported for Axcelerate 5.

Changes to Default Configuration

None.

Known Limitations

None.

6.3 Unlimited Number of Fields per Type (Reference ID: CORE-16491)

Behavior before change

Number of fields in any given engine was limited to 256.

Behavior after change

An index engine where the number of fields is enlarged is now automatically upgraded at engine start. If the number of fields changes across a power of 2, the upgrade process may take some time, since a number of data structures need to be converted.

Changes to Default Configuration

None, existing automatic update to data structure now works up to 65,000+ instead of up to 256.

Known Limitations

None.

6.4 Box Connector v2 (Reference ID: CORE-14333)

Behavior before change

Box connector was running on deprecated Box API.

Behavior after change

IBox connector was completely refactored and is now running on the most current version of the Box API (as of May 2016).

Changes to Default Configuration

None.

Known Limitations

Known limitations are described in the Box connector documentation.

7 Issues Resolved in 5.7.0 Update 1

7.1 AXC-34980 (Viewer)

Summary	New Behavior
Viewer sometimes did not load with the error: "Unable to verify the integrity of the configuration data", due to an encryption issue when checking the signature of the config xml.	Retry logic has been implemented which generates a new signed configuration in the case of a decryption failure, thereby avoiding that the error occurs a second time.

7.2 AXC-35230 (Printing)

Summary	New Behavior
Printing in Internet Explorer 11 did not work because the image download responses had no content type header.	The content type header is now set in the BravaServerProxy.

7.3 AXC-36357 (Production rules)

Summary	New Behavior
When accessing the the rules tab on the production workflows page for workflows with a high number of documents, a request timeout could occur.	Only documents where the rule tagging has changed are updated when accessing the rules tab, thereby reducing tagging operations and the risk of timeouts.

7.4 AXC-36379 (Business Intelligence)

Summary	New Behavior
Deleting a user from the database and logging into the application again as that user leads to the generation of a new UUID which prevented the user from accessing their Qlik sheets. In addition, the entire user directory was synchronized rather than a specific user.	A user will get the same UUID even if the user attributes are deleted from the database. In addition, only the current user who is a case manager will be synchronized in Qlik.

7.5 AXC-36464 (Report preview)

Summary	New Behavior
Downloading complex reports could cause an Out Of Memory Error on the App Server.	The downloading of complex reports is now streamed and the preview is limited to 100 values per axis to prevent Out Of Memory Errors.

7.6 AXC-36557 (Comments field)

Summary	New Behavior
The comments text box in the tagging panel did not wrap in newer Chrome versions.	The comments text box in the tagging panel now wraps in newer Chrome versions.

7.7 AXC-36590 (Production)

Summary	New Behavior
Production folder could not be retrieved if a production workflow was renamed and a production for this workflow was started from the Classic UI.	Production groups belonging to production workflows in AXC 5 are not displayed anymore in the Classic UI.

7.8 AXC-36648 (Associated Results)

Summary	New Behavior
The associated results view could highlight the parent document of a selected document, rather than the selected document, if the associated results list exceeded the 500 documents limit.	If the associated results list exceeds the 500 documents limit, and the selected document is outside the viewable range, the selected document is not highlighted and the user sees the error message: "Selected document is outside of viewable range."

7.9 AXC-36676 (Tagging panel)

Summary	New Behavior
UI does not indicate to the user if the Tagging panel is loading.	UI blocks the review panel and indicates status to the user if the Tagging panel is loading.

7.10 AXC-36682 (Filters on Assignments page)

Summary	New Behavior
The **Target Review State** Smart Filter on the **Assignments** page shows the total number of documents in the project to all users.	The **Target Review State** Smart fFlter and the **Assignee** Smart Filter on the **Assignments** page now show an **Apply** link rather than document counts.

7.11 AXC-36688 (Matter access)

Summary	New Behavior
If no pod base URL could be retrieved or was not set by the sys admin, the user was not able to access her matter.	If no pod base URL could be retrieved or was not set by the sys admin, the user will access the matter in the same tab.

7.12 AXC-36899 (Security vulnerability in 7zip 9.20)

Summary	New Behavior
7zip version shipped with Axcelerate 5 was 9.20, which had known security vulnerabilities.	7zip version is 16.0, which fixes the known security vulnerabilities.

7.13 CORE-16577 (Tagging processing)

Summary	New Behavior
The coding queue growth was caused by slow processing of change queue entries. This was caused by lock contention on the index engine due to a mixture of *many* read requests and *long* read requests. Furthermore, the duration of the periodically executed asynchronous Save - where no coding can be processed - was quite long. Additionally we observed reviewer impact of the long queue which was mainly caused by conversion tagging.	The processing speed of the save is significantly improved, minimizing the pause in potential processing of tagging (coding). In order to reduce the impact of the coding queue on the reviewers several changes are now in effect: » The data structures for accessing pending coding information are reconfigured to allow for faster access. » The conversion tagging are processed without blocking write locks. » The locking of the last result vectors is optimized. » The usage of in-memory structures is broadened.

7.14 CORE-16579 (Junk detection)

Summary	New Behavior
When base64 junk was distributed over multiple <text><p> tag sequences, the term based junk detection did not engage.	XML tags between potential junk terms are now skipped during junk detection so junk is recognized across tags.

7.15 CORE-16712, AXC-36811 (Bloomberg invitations)

Summary	New Behavior
Chats containing invites were always tagged with a single value even if they contained messages.	Two separate values, "Invite With Message Sent" and "Invite Without Message Sent", are now used for documents with invites.

7.16 FOUND-9171 (Inconsistent configurations)

Summary	New Behavior
Due to a template relation condition, it was possible to create inconsistent configurations.	The template relation condition has been fixed and it is no longer possible to create inconsistent configurations. Additionally a script is provided to manually fix broken projects by inheriting specified template values.

7.17 FOUND-9568 (Service tier)

Summary	New Behavior
The service tier startup blocked for over an hour when the jgroups cluster join process was stuck.	When the jgroups cluster join process is stuck, the service tier startup block is limited to 10 minutes and terminates after this time, freeing the startup synchronizer lock.

7.18 FOUND-9633 (Large cache)

Summary	New Behavior
Cache of presentation objects was growing over time.	Cache is now self-cleaning.

7.19 FOUND-9646 (Cleanup after crawl)

Summary	New Behavior
Data sources always wrote checkpoint information which could lead to time consuming cleanup at end of crawl.	Data sources write checkpoint information only if started with enabled "Automatic resume" option.

7.20 FOUND-9670 (Templates)

Summary	New Behavior
Extended template was not used in data source creation via web service.	Extended template is correctly used in data source creation via web service.

7.21 FOUND-9682 (Oracle OutsideIn operations)

Summary	New Behavior
Oracle OutsideIn operations (parsing, MIME type detection, conversion) failed if running as Windows service.	Oracle OutsideIn operations running as Windows service do not fail.

7.22 FOUND-9750 (Stored search)

Summary	New Behavior
Error message regarding a stored search without search parameters did not reference the stored search.	Error message regarding a stored search without search parameters now includes name of the stored search.

7.23 FOUND-9763 (Engine user session)

Summary	New Behavior
Engine user session was left untouched during long running count operation, causing user to get logged out.	Engine user session is touched during long running count operation so user does not get logged out.

8 New Feature Descriptions for Axcelerate 5.8.0

8.1 Production Option "Extracted text vs. Slip-sheet text (Reference ID: AXC-10527)

Behavior before change

The production text for documents produced as a slip-sheet matched the slip-sheet text. When doing a production of documents lacking a native, but matching a rule requesting **Native and Image** output, only the image was provided. (While a rule with only **Native** output would result in the creation of a native placeholder for these documents.)

Behavior after change

The user has the option to either use the document text when producing with slip-sheet output or to use the slip-sheet text (default, and used for all legacy rules).
If a rule with document text output will be used for a redacted document, the corresponding production cannot be started, and the user is informed via a notification, with direct links to modify the settings of the rule or production workflow. When doing a production of documents lacking a native, but matching a rule requesting **Native and Image** output, the image is provided along with a native placeholder file. (Same as when using a rule with only **Native** output which also results in the creation of a native placeholder for these documents.)

Changes to Default Configuration

There are no changes to the project-level configuration (in CORE Administration). In the application, the default for new rules is to use the **Text matches slip-sheet option**. This is also used for existing production format rules as this matches the legacy behavior.

Known Limitations

None.

8.2 UI Support to Re-convert Documents (Reference ID: AXC-20497)

Behavior before change

Documents could not be re-converted by end users through the Axcelerate 5 user interface. Command line scripts were needed for this task, which was time consuming and potentially error prone.
The two main use cases when a re-conversion is needed are:

1. Re-try failed conversions: To fix conversion errors, the users needed to fix the cause - e.g. extend a timeout, repair the document's native, or apply a software fix, and then run a commandline script to re-convert the relevant documents, (e.g., identified through a **Workspace** filter value).
2. Apply updated conversion settings: To apply updated conversion settings - such as the display of tracked changes for Microsoft Word documents - the user needed to apply the change of the setting, restart the engine, and then run a commandline script to re-convert for the relevant documents, (e.g., identified through a **Workspace** filter value), making sure that no redacted documents were included as this might cause shifted or lost redactions. The commandline script had dangerous options such as dropping redactions, or applying the old redactions to new representations of the document (which might have changed).

Behavior after change

Documents can be re-converted by end users through the Axcelerate 5 user interface, allowing them to fix temporary errors of the environment or software, or to apply updated conversion settings.
The two main use cases when a re-conversion is needed are:
1. Re-try failed conversions: To fix conversion errors, the users needs to fix the cause - e.g., extend a timeout, repair the document's native, or apply a software fix, and after running a search / selecting the documents, use the **Convert** wizard to re-convert documents with a failed conversion. Successfully converted documents are skipped, and so are ones with a redaction on the native representation.
2. Apply updated conversion settings: To apply updated conversion settings - such as the display of tracked changes for Microsoft Word documents - the user needs to apply the change of the setting, restart the engine, and after running a search / selecting the documents, use the **Convert** wizard to re-convert documents with a failed conversion. Documents with a redaction on the native representation are automatically skipped. Documents with redactions cannot be re-converted through the UI, thus avoiding the risk of loosing redactions or applying redactions to shifted renditions of the document.

Changes to Default Configuration

None.

Known Limitations

Images created via the classic Axcelerate Analysis module cannot be reconverted.

8.3 Document Flagging for Jobs (Reference ID: CORE-15701)

Behavior before change

Documents processed by a job were not flagged and these documents could hence not be retrieved effectively from within the Axcelerate 5 user interface. For some jobs, the effect of a job execution was not visible at all (e.g. when bulk printing documents); for others the effect could be seen (e.g. using the **Conversion** Smart Filter when converting a document). But even in these cases a repeated execution was invisible as well (e.g. converting a documents 1 or 2 times resulted in the same flagging in the **Conversion** Smart Filter).

Behavior after change

Documents processed by specific jobs are tagged into a new **Job Processing** field and can be retrieved from within the Axcelerate 5 using a new **Job Processing** Smart Filter in the **Work Product** Smart Filter group.
Tagging is available for these job types:

- Conversion (only for Bulk, not for on-the-fly jobs)
- Global Redaction
- Redaction rRemoval
- Production
- Printing
- OCR

Changes to Default Configuration

An additional default field is used for the Job Processing field, with the internal id `rmJobStatus`.

Known Limitations

None.

8.4 Axcelerate - Continuous Improvements (Reference ID: AXC-35731)

Behavior before change

The indication of the boundaries between families was sometimes hard to see, especially in cases when only parts of families are shown in the result set (i.e., when the root doc is missing)
In an Associated Results list (e.g. when using the **Review** page in 2-screen mode), the buttons to navigate to the next main result are only shown on mouse-over. Unless the user is aware of this, the existence of this functionality is usually missed.

Behavior after change

The indication of the boundaries between families has been improved by making the boundary within a family lighter, and by replacing the gradient boundary between family by a dashed line.
In an Associated Result list (e.g. when using the **Review** page in 2-screen mode), the buttons to navigate to the next main result are always shown, and making the user aware of this functionality. On mouse-over, the opacity for these buttons is increased to highlight them further.

8.5 Unlimited Number of Fields per Type (Reference ID: CORE-16491)

Behavior before change

Number of fields in any given engine was limited to 256.

Behavior after change

An index engine where the number of fields is enlarged is now automatically upgraded at engine start. If the number of fields changes across a power of 2, the upgrade process may take some time, since a number of data structures need to be converted.

Changes to Default Configuration

None, existing automatic update to data structure now works up to 65,000+ instead of up to 256.

Known Limitations

None.

8.6 Classic user interfaces run inside Tomcat 8 application (Reference ID: FOUND-983)

Behavior before change

The Classic user interfaces (Classic Review & Analysis, Axcelerate Ingestion, ECA, CORE Administration, Case Management Tool, Perceptiv) run in Tomcat 6.

Behavior after change

The Classic user interfaces (Classic Review & Analysis, Axcelerate Ingestion, ECA, CORE Administration, Case Management Tool, Perceptiv) run in Tomcat 8. The Axcelerate 5 user interface continues to run in Tomcat 7.

Changes to Default Configuration

Increased metaspace default is set to 512MB as Tomcat 8 has a higher memory consumption.

Known Limitations

Support for the classic Axcelerate Analysis and Axcelerate Review modules is limited to 10 concurrent reviewers., These interfaces are only used during the migration of cases to the Axcelerate 5 user intercae and not intended for massive document review. .

8.7 Retry mini job execution in case of batch server failure (Reference ID: CORE-14508)

Behavior before change

Job execution is usually done as a sequence of steps that are called *mini jobs*, typically a single processing job for each document in the executed job. For most jobs, such a mini job is sent from a CORE engine server to a batch server for remote execution. A batch server outage resulted in reported exceptions for all mini jobs in the executed job.

Behavior after change

Mini jobs affected by outages are now automatically retried on a different batch server. There is no longer any exception reported since the mini job is now successfully executed. The corresponding batch server is tracked as unavailable until it becomes online again. Affected job types include:

- Conversion
- Production
- Printing
- OCR
- Bulk redaction

Changes to Default Configuration

None.

Known Limitations

None.

9 Issues Resolved in Axcelerate 5.8.0

9.1 AXC-27381 (Bulk OCR/Production time OCR)

Summary	New Behavior
When Transym 4 fails to OCR a page with error code 8, this caused the whole document to be flagged as an exception.	When Transym 4 fails to OCR a page with error code 8, the corresponding page is replaced with this text: "OCR of page failed".

9.2 AXC-32122 (Robustness and stability)

Summary	New Behavior
Cache of UserCaseConfiguration caused OOM due to duplicated strings.	Cache now contains less redundancies.

9.3 AXC-34980 (Viewer)

Summary	New Behavior
Viewer sometimes did not load with the error: "Unable to verify the integrity of the configuration data", due to an encryption issue when checking the signature of the config xml.	Retry logic has been implemented which generates a new signed configuration in the case of a decryption failure, thereby avoiding that the error occurs a second time.

9.4 AXC-36295 (Metadata fetching)

Summary	New Behavior
Metadata fields for a document were fetched from the engine.	Metadata fields for a document are now fetched from the cache, thereby improving performance.

9.5 AXC-36357 (Production rules)

Summary	New Behavior
When accessing the the rules tab on the production workflows page for workflows with a high number of documents, a request timeout could occur.	Only documents where the rule tagging has changed are updated when accessing the rules tab, thereby reducing tagging operations and the risk of timeouts.

9.6 AXC-36379 (Business Intelligence)

Summary	New Behavior
Deleting a user from the database and logging into the application again as that user leads to the generation of a new UUID which prevented the user from accessing their Qlik sheets. In addition, the entire user directory was synchronized rather than a specific user.	A user will get the same UUID even if the user attributes are deleted from the database. In addition, only the current user who is a case manager will be synchronized in Qlik.

9.7 AXC-36461 (Highlighting)

Summary	New Behavior
Metadata highlighting for document search did not work with Regular Expression.	Metadata highlighting works for document search with Regular Expression.

9.8 AXC-36535 (Robustness and stability)

Summary	New Behavior
Arrangement descriptions and associated data were loaded into memory and stored specifically for each user and project. These descriptions were never unloaded from memory. With a growing number of users and projects, the memory consumption grew accordingly, eventually causing out-of-memory situations.	Arrangement descriptions are unloaded from memory when they have been idle (i.e. have not been accessed in some way) for 30 minutes.

9.9 AXC-36557 (Comments field)

Summary	New Behavior
The comments text box in the tagging panel did not wrap in newer Chrome versions.	The comments text box in the tagging panel now wraps in newer Chrome versions.

9.10 AXC-36567 (Viewer)

Summary	New Behavior
Sporadically, view loads broke on loading the markup file.	Improved closing of views before reuse and error handling.

9.11 AXC-36590 (Production)

Summary	New Behavior
Production folder could not be retrieved if a production workflow was renamed and a production for this workflow was started from the Classic UI.	Production groups belonging to production workflows in AXC 5 are not displayed anymore in the Classic UI.

9.12 AXC-36648 (Associated Results)

Summary	New Behavior
The associated results view could highlight the parent document of a selected document, rather than the selected document, if the associated results list exceeded the 500 documents limit.	If the associated results list exceeds the 500 documents limit, and the selected document is outside the viewable range, the selected document is not highlighted and the user sees the error message: "Selected document is outside of viewable range."

9.13 AXC-36676 (Tagging panel)

Summary	New Behavior
UI does not indicate to the user if the Tagging panel is loading.	UI blocks the review panel and indicates status to the user if the Tagging panel is loading.

9.14 AXC-36682 (Filters on Assignments page)

Summary	New Behavior
The **Target Review State** Smart Filter on the **Assignments** page shows the total number of documents in the project to all users.	The **Target Review State** Smart fFIter and the **Assignee** Smart Filter on the **Assignments** page now show an **Apply** link rather than document counts.

9.15 AXC-36688 (Matter access)

Summary	New Behavior
If no pod base URL could be retrieved or was not set by the sys admin, the user was not able to access her matter.	If no pod base URL could be retrieved or was not set by the sys admin, the user will access the matter in the same tab.

9.16 AXC-36712 (Review performance)

Summary	New Behavior
Each click on **Save+Next unreviewed** button triggered a back end call to determine the next unreviewed document in the list.	Before using the back end to determine the next unreviewed document, the list of documents on client side will be used.

9.17 AXC-36732 (Associated results display)

Summary	New Behavior
The document displayed did not match the highlighted document in the associated results view on the **Analysis** page.	The display is now correct.

9.18 AXC-36733 (Review performance)

Summary	New Behavior
Single document tagging resulted in an expansive call to determine the MIME type of that document.	Single document tagging now sends the MIME type of that document with the save request.

9.19 CORE-14974 (Text field used for enrichment during publish)

Summary	New Behavior
When using a text field for enrichment and starting a publish with more than 1000 documents, the engine ran into an IllegalStateException due to a safety measure that blocked enriching text fields for more than 1000 documents.	When using a text field for enrichment and starting a publish with more than 1000 documents, the engine does not cause an exception.

9.20 CORE-15862 (LiveLink connector)

Summary	New Behavior
LiveLink connector option **Ignore MIME type** was not functioning.	LiveLink connector option **Ignore MIME type** functions.

9.21 CORE-16000 (Jobs)

Summary	New Behavior
The system didn't ensure the proper job priority with the following order: 1. "On The Fly Conversion", 2. "Bulk Print", 3. "Bulk Conversion	The system ensures the proper job priority with the following order: 1. "On The Fly Conversion", 2. "Bulk Print", 3. "Bulk Conversion"

9.22 CORE-16009 (Performance)

Summary	New Behavior
Clicking on the last document of a search page caused the index to load unnecessarily high amounts of data into the java heap.	Unnecessary memory consumption has been eliminated and index routines have been made more robust.

9.23 CORE-16071 (Large searches)

Summary	New Behavior
Large boolean searches for numeric identifiers of the form "ID=<value1> AND ID==<value2>..." were not optimized and resulted in high memory consumption.	Large boolean searches for numeric identifiers of the form "ID=<value1> AND ID==<value2>..." are now optimized and result in less memory consumption.

9.24 CORE-16101 (Query exception handling)

Summary	New Behavior
Improper exception handling for very complex queries could cause the engine to crash	Exceptions caused by very complex queries are now handled properly and do not crash the engine.

9.25 CORE-16218 (Field value sorting)

Summary	New Behavior
Sorting field values by property could cause an exception.	Sorting field values by property no longer causes an exception.

9.26 CORE-16303 (SharePoint XML)

Summary	New Behavior
Invalid XML created by SharePoint web-services wasn't handled gracefully.	Invalid XML created by SharePoint web-services is handled more gracefully

9.27 CORE-16371 (Performance)

Summary	New Behavior
A high percentage of Bloomberg documents have a large number of recipients and storing these recipients in the corresponding fields took a significant portion of the crawl time.	A post processor can be activated that triggers for documents with recipient counts above a configurable threshold (default:50). For these documents the meta entries 'rm_to', 'meta_recipientname', and 'meta_toname' are reduced to a single value 'LARGE_RECIPIENT_COUNT'.

9.28 CORE-16419 (Write lock removal logging)

Summary	New Behavior
No log message was created when a user manually removed a write-lock on an engine.	A log message is created when a user manually removes a write-lock on an engine.

9.29 CORE-16447 (Stack trace dumping)

Summary	New Behavior
Mini dump creation could fail due to a naming conflict.	The process ID is now used to make the mini dump file names unique.

9.30 CORE-16492 (Robustness and stability)

Summary	New Behavior
An out-of-memory of physical RAM caused a server downtime.	Memory consumption is lower than before after deactivation of expensive and unnecessary prefix search indices. Furthermore, the memory consumption of searches in an engine with many ongoing changes has been reduced significantly. The risk of system downtime is reduced due to increased stability of the exception handling.

9.31 CORE-16577 (Tagging processing)

Summary	New Behavior
The coding queue growth was caused by slow processing of change queue entries. This was caused by lock contention on the index engine due to a mixture of *many* read requests and *long* read requests. Furthermore, the duration of the periodically executed asynchronous Save - where no coding can be processed - was quite long. Additionally we observed reviewer impact of the long queue which was mainly caused by conversion tagging.	The processing speed of the save is significantly improved, minimizing the pause in potential processing of tagging (coding). In order to reduce the impact of the coding queue on the reviewers several changes are now in effect: » The data structures for accessing pending coding information are reconfigured to allow for faster access. » The conversion tagging are processed without blocking write locks. » The locking of the last result vectors is optimized. » The usage of in-memory structures is broadened.

9.32 CORE-16579 (Junk detection)

Summary	New Behavior
When base64 junk was distributed over multiple <text><p> tag sequences, the term based junk detection did not engage.	XML tags between potential junk terms are now skipped during junk detection so junk is recognized across tags.

9.33 CORE-16614 (Asynchronous saves)

Summary	New Behavior
Asynchronous save job concurring with a category modification operation could cause the engine to run into a deadlocked state when the save job was stopped (either via the jobs tab or due to another thread removing or adding a category).	When an asynchronous save is stopped the internal save gate will now always be properly opened for other threads again.

9.34 CORE-16621 (Conversion performance)

Summary	New Behavior
Streaming of conversion results from batch server to Engine used an unexpectedly small buffer size, resulting in poor throughput rates.	Buffer size has been increased and the streaming operation is about three times faster.

9.35 CORE-16641 (CertificateUtilities)

Summary	New Behavior
CertificateUtilities was not available in ToolBox.	CertificateUtilities is available in ToolBox.

9.36 CORE-16662 (Native conversion)

Summary	New Behavior
S3-based projects failed to convert Microsoft Office documents of recent Microsoft Office versions.	S3-based projects convert as expected; the required file extension is set correctly.

9.37 Found-9087 (MSG/PST files with TNEF attachments)

Summary	New Behavior
SMIME encrypted/signed MSG files/PST entries with TNEF attachments were not supported.	SMIME encrypted/signed MSG files/PST entries with TNEF attachments are supported.

9.38 FOUND-9171 (Inconsistent configurations)

Summary	New Behavior
Due to a template relation condition, it was possible to create inconsistent configurations.	The template relation condition has been fixed and it is no longer possible to create inconsistent configurations. Additionally a script is provided to manually fix broken projects by inheriting specified template values.

9.39 FOUND-9513 (Process control robustness)

Summary	New Behavior
A non-responding DNS resolver could break the process status handling.	The process control is more robust against a non-responding DNS resolver.

9.40 FOUND-9582 (Adobe Indesign Interchange documents)

Summary	New Behavior
Adobe Indesign Interchange documents caused a NullPointerException if Oracle OutsideIn MIME type detection was enabled.	Adobe Indesign Interchange documents no longer cause a NullPointerException.

9.41 FOUND-9612 (Application import)

Summary	New Behavior
Imports failed if any of the MSV files were empty.	Imports do not fail even if an MSV file is empty.

9.42 FOUND-9633 (Large cache)

Summary	New Behavior
Cache of presentation objects was growing over time.	Cache is now self-cleaning.

9.43 FOUND-9646 (Cleanup after crawl)

Summary	New Behavior
Data sources always wrote checkpoint information which could lead to time consuming cleanup at end of crawl.	Data sources write checkpoint information only if started with enabled "Automatic resume" option.

9.44 FOUND-9670 (Templates)

Summary	New Behavior
Extended template was not used in data source creation via web service.	Extended template is correctly used in data source creation via web service.

9.45 FOUND-9728 (Loading performance)

Summary	New Behavior
Timeouts occur due to unnecessary long loading times when clicking on documents or saving tags.	Loading times when clicking on documents or saving tags are faster and do not increase with the number of available batches.

9.46 FOUND-9740 (CSV Load Wizard)

Summary	New Behavior
ArrayIndexOutOfBoundsException was possible in CSV Load Wizard.	Array indices are now handled correctly inside CSV Load Wizard.

9.47 FOUND-9750 (Stored search)

Summary	New Behavior
Error message regarding a stored search without search parameters did not reference the stored search.	Error message regarding a stored search without search parameters now includes name of the stored search.

9.48 FOUND-9763 (Engine user session)

Summary	New Behavior
Engine user session was left untouched during long running count operation, causing user to get logged out.	Engine user session is touched during long running count operation so user does not get logged out.

10 Issues Resolved in Axcelerate 5.7.2

10.1 AXC-35556 (Search Query Editor)

Summary	New Behavior
When running a report for a complex query in SQE, the browser could timeout.	When running a report for a complex query in SQE, the timeout will only occur after 5 minutes.

10.2 AXC-36156 (Login)

Summary	New Behavior
When paging to another page than the first one, the user was not able to login to the matter.	When paging to another page than the first one, the user will see all available cases and is able to login to each matter.

10.3 AXC-36183 (Storage size audit log)

Summary	New Behavior
For billing, in the `audit.core_storage_sizes` table, a storage handler is not reported if there is no storage allocated yet.	There is now an entry for every storage handler in the `audit.core_storage_sizes` table, even if there is no storage allocated yet (with sizes 0 bytes).

10.4 AXC-36187 (Storage size collection)

Summary	New Behavior
For projects that have not been started or have been removed from disk only, a misleading error was reported when collecting their storage sizes (E1).	For projects that have not been started or have been removed from disk only, a warning is reported together with a message containing possible causes.

11 Issues Resolved in Axcelerate 5.7.1

11.1 AXC-35230 (Printing)

Summary	New Behavior
In Internet Explorer 11, when pressing the print button of the Axcelerate 5 viewer panel, a print job was run, but the print result could not be shown.Instead, the user had to cancel the print job message displayed.	Pressing the Print button in the Axcelerate 5 viewer results in correct print output.

12 Contact Us

About Recommind

Recommind provides the most accurate and automated enterprise search, automatic classification, and eDiscovery software available, giving organizations and their users the information they need when they need it.

Visit us at http://www.recommind.com.

Support

For support issues on Recommind products, visit the Recommind Ticketing System at https://rts.recommind.com.

Documentation

Find Recommind product documentation, Knowledge Base articles, and more information at the Recommind Customer Portal at https://supportkb.recommind.com. For login access to the site, contact your product support:

- For : SearchSupport@recommind.com
- For : eDiscoverySupport@recommind.com

The Recommind Documentation team is interested in your feedback.

For comments or questions about Recommind product documentation, contact us at documentation@recommind.com.

13 Terms of Use

Disclaimer

This document, as well as the products and services described in it, is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this document is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Recommind, Inc., including its affiliates and subsidiaries (collectively, "Recommind"). Recommind assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software or services that may be provided in association with this document.

Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Recommind. Information in this document is provided in connection with Recommind's products and services. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document.

EXCEPT AS PROVIDED IN RECOMMIND'S SOFTWARE LICENSE AGREEMENT OR SERVICES AGREEMENT FOR SUCH PRODUCTS OR SERVICES, RECOMMIND ASSUMES NO LIABILITY WHATSOEVER, AND RECOMMIND DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF RECOMMIND PRODUCTS OR SERVICES INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. RECOMMIND MAKES NO WARRANTIES REGARDING THE COMPLETENESS OR ACCURACY OF ANY INFORMATION, NOR THAT THE PRODUCTS OR SERVICES WILL BE ERROR FREE, UNINTERRUPTED, OR SECURE. IN NO EVENT WILL RECOMMIND, THEIR DIRECTORS, EMPLOYEES, SHAREHOLDERS AND LICENSORS, BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR EXEMPLARY DAMAGES INCLUDING, BUT NOT LIMITED TO, LOSS OF ANTICIPATED PROFITS OR BENEFITS.

Recommind may make changes to specifications, and product and service descriptions at any time, without prior notice. Recommind's products may contain design defects or errors known as errata that may cause the product or service to deviate from published specifications. Current characterized errata are available on request. Whilst every effort has been made to ensure that the information and content within this document is accurate, up-to-date and reliable, Recommind cannot be held responsible for inaccuracies or errors. Recommind software, services and documentation have been developed and prepared with the appropriate degree of skill, expertise and care. While every effort has been made to ensure that this documentation contains the most up-to-date and accurate information available, Recommind accepts no responsibility for any damage that may be claimed by any user whatsoever for the specifications, errors or omissions in the use of the products, services and documentation.

Trademarks and Patents

Recommind's underlying technology is patented under *U.S. Patent Nos. 6,687,696, 7,328,216, 7,657,522, 7,747,631, 7,933,859, 8,024,333, 8,103,678, 8,429,159 and 8,489,538*

Recommind, Inc. is the leader in predictive information management and analysis software, delivering business applications that transform the way enterprises, government entities and law firms conduct eDiscovery, enterprise search, and information governance. Recommind, Axcelerate, Axcelerate Cloud, Axcelerate OnDemand, and CORE's name and logo are registered trademarks of Recommind, Inc.

Copyright