Axcelerate 5.11.0

Release Notes

Published: 2017-Mar-03

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 on premise.

2 New Feature Descriptions for Axcelerate 5.11.0

2.1 Internet Explorer 10 No Longer Supported (AXC-38536)

Behavior before change

Internet Explorer 10 and 11 were supported.

Behavior after change

Internet Explorer 10 is no longer supported.
Internet Explorer 11 is still supported.

2.2 Daeja No Longer Available in Axcelerate Classic User Modules (AXC-29736)

Behavior before change

It was possible to conduct the following actions in the classic Axcelerate user interface with the Daeja ViewONE applet:

- Create productions
- View documents using Image view
- View produced documents using Production view
- Convert documents to TIFF
- Create conversion slip-sheets
- Apply redactions and annotations
- Mark and unmark documents **Redaction required**

Behavior after change

The Daeja ViewONE applet is no longer available in the Axcelerate classic user modules. As a result, the following actions are no longer possible in the classic Axcelerate user interface:

- Creating productions
- Viewing documents using Image view
- Viewing produced documents using Production view
- Converting documents to TIFF

- Creating conversion slip-sheets
- Applying redactions and annotations
- Marking and unmarking documents **Redaction required**

Note: These actions remain available in the Axcelerate 5 user interface.

2.3 Axcelerate 5 Application Cloning (AXC-36962)

Behavior before change

For a new Axcelerate 5 matter cloned from an existing one, the following application-sided settings were copied into the new matter:

- Fields and values, including **Review State** field values
- Tagging arrangements, including the corresponding fields and values, tagging rules, quick tags, and the mapping to review states

Behavior after change

For a new Axcelerate 5 matter cloned from an existing one, the following application-sided settings are also copied into the new matter:

- Universes, review workflows and their settings, including document independent searches (scopes, samples and iterations are not copied).
- Batching Templates
- Production Templates

Changes to Default Configuration

The axcelerate.standaloneSystemTemplate-Axc5-v2.config.xml has a new property enabled by default (litigation_ngTemplatable).

Known Limitations

The matter to clone must be in the same pod and must be both running and enabled for cloning (new internal application configuration/enabled by default for axcelerate.standaloneSystemTemplate-Axc5-v2.config). Cloning an offline matter or a matter existing in a different pod is not possible.
All settings are cloned, i.e., you cannot pick and choose which settings are copied to the new matter.
Schedules are not supported by templates.

2.4 Application-Level User Security (AXC-37590)

Behavior before change

Access to Axcelerate 5 features was controlled by assigning user groups to one or multiple of the following roles:

- Reviewer
- Power Reviewer
- Power User
- Case Manager

In addition to the existing role-based security, access to the following feature could be added or removed individually for individual user groups:

- Download native files

Behavior after change

Access to the following features can also now be added or removed individually for individual user groups:

- Manage batches
- Bulk tag documents
- View associated documents outside of batch
- Manage rights of users to pull batches
- Add and modify field values on the tagging and foldering panels
- Manage fields via the Fields & Values page
- Manage arrangements
- Access document history
- Manage batching templates

2.5 Client Agent Supported on Windows 10 (FOUND-9958)

Behavior before change

Agent supported to run on Windows 7.

Behavior after change

Agent supported to run on Windows 10.

2.6 64-bit Crawls configurable (FOUND-10089)

Behavior before change

Data source crawls were run in 32-bit mode. Crawls were always limited to a 2 GB memory. If large files were loaded in parallel, there was a risk of running out of memory.

Behavior after change

For file share and for publishing data sources, administrators can now configure the data source to use Java 64-bit runtime environment, provided that Microsoft Outlook 64-bit is installed on the crawler server. 64-bit crawls are not limited to the 2 GB memory limit, are faster, and they reduce the risk of running out of memory due to large files.
The default crawl mode is still 32-bit mode, which requires Microsoft Outlook 32-bit on crawler servers.

Known Limitations

The Aid4Mail parser is not compatible with the new 64-bit mode.

2.7 Case-wide Sharing for Amazon S3 Storages (CORE-16541)

Behavior before change

Axcelerate Cloud used SQL-managed storage with single-instance storage, i.e., de-duplication per storage handler.

Behavior after change

New system templates are available that allow case-wide sharing of image files and native files, i.e., file sharing between Axcelerate Ingestion and matters created from Axcelerate Ingestion, thus reducing storage space. If case-wide sharing is correctly configured, deleting Axcelerate Ingestion has no effect on published matters.

2.8 User Interface Changes

Behavior before change

The **Status** icon in the following locations was red:

- Navigation and Processing pages, to indicate “Not running” applications
- Fields and Values page, to indicate inactive fields
- Templates page, to indicate inactive templates
- Production Workflows>Production Rules page, to indicate inactive production rules
- Preconversion Rules page, to indicate inactive preconversion rules

The icon used to add a new value to a field on the **Administration > Fields and Values** page was represented by the icon.
In the **Tagging** panel, the **Edit** button was used to select documents.
On the **Analysis** page, if you had rights to add or edit field values in the **Tagging** panel, you could take these actions only after first selecting a document to activate the panel.

Behavior after change

The **Status** icon is now gray .
The icon used to add a new value to a field on the **Administration > Fields and Values** page is represented by the New icon.
In the **Tagging** panel, the **Select** button replaces the **Edit** button for document selection.
On the **Analysis** page, if you have rights to add or edit field values in the **Tagging** panel, you can now take these actions without first selecting a document to activate the panel.

3 Issues Resolved in Axcelerate 5.11.0

3.1 AXC-27093 (Duplicate Status Category Display Name)

Summary	New Behavior
A global **Duplicate Status** category had an incorrect display name, i.e., **Remove**.	The display name of the global **Duplicate Status** category is again named **Duplicate**.

3.2 AXC-27834 (Results List View)

Summary	New Behavior
Changing the **Results** list grid columns while viewing documents in QAM mode, e.g., documents queued for production or documents in a production set, with the **Applied rule** or **New rule** column active, caused a JavaScript error and broke the Results list grid.	Resorting columns when the **Applied rule** or **New rule** column is active no longer causes a JavaScript error.

3.3 AXC-34909 (Results List)

Summary	New Behavior
Values for multi-value fields were not sorted correctly when viewed in the **Results** list.	Values for multi-value fields are now sorted correctly when viewed in the **Results** list.

3.4 AXC-35531 (Efficiency Dashboard)

Summary	New Behavior
The **Efficiency** dashboard was not using the standard resizing behavior for all text, which resulted in text of different sizes.	The **Efficiency** dashboard now uses the same resizing behavior for all text.

3.5 AXC-35704 (Adaptive Batching)

Summary	New Behavior
Adaptive batching was skipped for samples.	Adaptive batching no longer skipped for samples.

3.6 AXC-36127 (User Preferences)

Summary	New Behavior
Saving the user preferences showed a **Save Success** message.	Saving the user preferences now shows a **Preferences saved successfully** message.

3.7 AXC-36675 (Tagging)

Summary	New Behavior
Save tagging and foldering requests were not validated correctly, which allowed manipulation of tagging and text fields.	Save tagging and foldering requests are now validated correctly and tagging and text fields cannot be manipulated.

3.8 AXC-37003 (Apache Commons BeanUtils)

Summary	New Behavior
Older version of Apache Commons BeanUtils do not suppress the class property, which could allow remote attackers to manipulate the **ClassLoader** and execute arbitrary code via the class parameter.	The security issue is fixed by using the newest version of Apache Commons BeanUtils.

3.9 AXC-37116 (Associated Results View)

Summary	New Behavior
Navigation to a previous document in **Associated Results** view caused an error when there were more than 500 documents.	Navigation to a previous document in **Associated Results** view no longer causes an error when there are more than 500 documents.

3.10 AXC-37410 (Jobs)

Summary	New Behavior
A JavaScript error occured when a paused job was restarted on the **Jobs** page.	A JavaScript error no longer occurs when a paused job is restarted on the **Jobs** page.

3.11 AXC-37595 (Mission Control > Review Workflows Page)

Summary	New Behavior
Loading a review workflow page with many target review states and/or many tagging fields could run into a timeout for large universes.	Loading a review workflow page with many target review states and/or many tagging fields no longer causes a timeout for large universes.

3.12 AXC-37667 (Navigation Page)

Summary	New Behavior
Application sync changes were synchronized with one large database transaction, which could cause a timeout and prevent some matters from displaying on the Navigation page.	The Application sync changes are now done with small transactions and will not cause a timeout.

3.13 AXC-37702 (Arrangements)

Summary	New Behavior
When an arrangement with id **axcelerate.tagging.foldering** was changed from type **Folder** to **Tagging** and then a fields-and-values synchronize was triggered, the arrangement was removed from the user interface so it could not be used, edited or restored.	When an arrangement with id **axcelerate.tagging.foldering** is changed to a **Tagging** type, the arrangement no longer disappears from the user interface after a fields-and-values synchronize.

3.14 AXC-37785 (Conversion)

Summary	New Behavior
BravaServer generated a job description that pointed to a non-existant XDL source folder, which resulted in failed conversions.	This issue has been resolved in BravaServer, thereby avoiding failed conversions.

3.15 AXC-37790 (Recent Searches)

Summary	New Behavior
Deletion of recent searches caused the Tagging panel to not load.	Deletion of recent searches is now performed using the database instead of application logic, thereby ensuring the Tagging panel will load.

3.16 AXC-38148 (User Preferences)

Summary	New Behavior
Saving user preferences could lead to an error if some unknown parameters were sent to the server.	Unknown parameters are ignored so saving user preferences does not fail.

3.17 AXC-38305 (Tomcat Update)

Summary	New Behavior
The current Tomcat version for Axcelerate 5 was 7.0.63, which contained known security issues.	The Tomcat version for Axcelerate 5 has been updated to 7.0.73, which contains fixes for these security issues.

3.18 AXC-38426 (User Logout Events)

Summary	New Behavior
A deadlock BPM process caused the workflow engine to not restart the BPM process, resulting in user logout events.	The workflow engine now tries to restart the BPM process up to 10 times, thereby alleviating user logout events.

3.19 AXC-38437 (Search)

Summary	New Behavior
Search requests for associated results were executed twice in parallel when navigating through an associated results list.	Search requests for associated results are always executed one at a time.

3.20 AXC-38440 (Page Navigation)

Summary	New Behavior
Persisting the most recently accessed page for a user could lead to an exception in the database and cause a red error box when switching to another page in the application.	Switching to another page in the application no longer causes a red error box.

3.21 AXC-38448 (404 Responses)

Summary	New Behavior
The wrong internal usage of a REST endpoint led to 404 responses from the service-tier, polluting the service-tier Tomcat access log.	The REST endpoint is now correctly used and 404 responses no longer pollute the service-tier Tomcat access log.

3.22 AXC-38650 (Production Export)

Summary	New Behavior
Exporting of a subset of sub-productions could cause the export job to fail when the sub-productions were not contiguous (for instance: export of Vol1 and Vol3 instead of Vol1, Vol2 and Vol3).	Exporting arbitrary subsets of sub-productions no longer causes the export job to fail.

3.23 CORE-9061 (CSV Crawls)

Summary	New Behavior
CSV crawls could produce nested tags inside a **rm_numeric_identifier** field.	The XML parser is fixed and produces correct data.

3.24 CORE-9390 (Publish)

Summary	New Behavior
Document XML structure could become invalid during publishing.	The internal <rm_id.profile> elements are now always added to the end of the document's XML structure.

3.25 CORE-12697 (Conversion)

Summary	New Behavior
Running a "convert for display" job on a document that had been viewed previously resulted in unnecessary and misleading stack traces and unnecessary zip operations on batch servers.	The offending logging and zip operation has been disabled.

3.26 CORE-14923 (SSL for SQL connections)

Summary	New Behavior
SSL could only be configured if the system had the non-standard configuration "Use no SQL connection pool". This meant that SSL was not possible for SQL-managed storages.	The tool used to configure the SQL connection pool accepts a parameter for SSL and configures the connection pool and its consumers accordingly.

3.27 CORE-15615 (Crawls)

Summary	New Behavior
Content Server connector did not support large collections.	Content Server connector now supports large collections.

3.28 CORE-17045 (Publish)

Summary	New Behavior
AWS S3 connectivity slowness during large publish.	Log error message improved.

3.29 CORE-17085 (CORE Administration)

Summary	New Behavior
Opening the **Jobs** tab for a project with a meta-engine could result in a **NullPointerException** when one of the underlying single-engines was unavailable and thus did not send its job status.	The **Jobs** tab can now be opened for a project with meta-engine even if an underlying single-engine is unavailable. Information on such engine is then missing from the list and the event is logged as an error.

3.30 CORE-17561 (OCR Jobs)

Summary	New Behavior
Due to the resource management in distributed computing environments (OCR resource != conversion resource), the engine job manager would occasionally start all document jobs for a bulk OCR operation at once. Most of those jobs would wait in a blocked state within the engine for their respective resource. Pausing a bulk operation was not possible, because all document jobs were already started and had to be finished first.	The engine prevents starting new document jobs for a macro job when a certain threshold of blocked jobs for that macro job id is exceeded.

3.31 CORE-17729 (CSV Export)

Summary	New Behavior
For meta engine projects, it could happen that the CSV export of some sub-engines failed.	We reduced the likelihood of failure and improved the error handling.

3.32 CORE-17759 (Production Storage Handler)

Summary	New Behavior
It was possible to configure a storage handler of an engine for **Bates Number File Storage** and at the same time provide options for **Structured Storage**, which are neither required nor wanted for **Bates Number File Storage**.	Added a sanity check, which causes an engine (re-)start to fail when **Bates Number File Storage** is misconfigured.

3.33 CORE-17772 (OCR Jobs)

Summary	New Behavior
A canceled OCR job led to orphaned temporary files on batch processing servers.	Temporary files are cleaned up after a job is canceled.

3.34 CORE-17845 (Conversion)

Summary	New Behavior
On-the-fly conversions had a long timeout which caused batch servers to receive and hold too many jobs.	Timeout for on-the-fly jobs is back to two minutes.

3.35 CORE-17891 (CSV Load)

Summary	New Behavior
Bloomberg documents with many recipients, loaded by CSV load, produced illegal values for the email thread detection and led to crawler out-of-memory while email threads were computed.	Validity is ensured by means of additional sanity checks and out-of-memory has been avoided by suitable limits.

3.36 CORE-17984 (Redaction Jobs)

Summary	New Behavior
Redaction jobs left unused temporary files on batch processing servers.	Temporary files of redaction jobs are cleaned up properly.

3.37 CORE-18022 (TIFF conversion, OCR)

Summary	New Behavior
Axcelerate Classic conversion and Bulk OCR did not cover exception documents.	Axcelerate Classic conversion and Bulk OCR cover exception documents. Only documents with these exception classes are excluded from Axcelerate Classic conversion and Bulk OCR: File format, Archive, Password protected/encrypted.

3.38 CORE-18030 (API Function)

Summary	New Behavior
The **changeTextType()** API function could be misused to change taxonomy fields, bypassing folder-level security.	The **changeTextType()** API function can no longer be misused to change taxonomy fields.

3.39 CORE-18147 (Crawls)

Summary	New Behavior
Incremental crawl was slow due to the root resource for deletions being processed multiple times, resulting in redundant deletion resources.	The root resource for deletions is processed only once.

3.40 CORE-18163 (Sortable Fields)

Summary	New Behavior
The **rm_attachmentcount** and **rm_embeddingcount** fields were not marked as sortable in the data model configuration. As a consequence they were not visible in the Axcelerate 5 user interface.	A new data model template is now active with **rm_attachmentcount** and **rm_embeddingcount** being marked as sortable. This change affects only new projects created with Axcelerate 5.11.

3.41 CORE-18164 (Storage Migration)

Summary	New Behavior
When removing entries from the storage locations persistency, the smallest key of the underlying map was sometimes not updated correctly. This led to an exception during storage location exports.	The min and max keys of the storage locations persistency are now updated correctly when removing objects.

3.42 CORE-18231 (Processing Resources)

Summary	New Behavior
A bug in the job grid caused the engine to ignore offline states of bat servers. This caused jobs to be sent to batch servers that were in offline state.	The engine now properly checks the availability of processing resources. In addition, a wide range of Exceptions regarding the connection to the batch processor will now result in a retry for the job affected.

3.43 CORE-18246 (Taggers)

Summary	New Behavior
Uninstalling taggers blocked other operations unnecessarily.	The index can be used during the removal of taggers.

3.44 CORE-18257 (Storage Migration)

Summary	New Behavior
We introduced a tagging into the storage migration, which tags all documents that were successfully copied to AWS S3; however, this tagging was not respected when canceling and restarting the initial copy-job.	When restarting the storage migration copy job, potential tagging of previous runs are now respected.

3.45 CORE-18304 (File Storage Initialization)

Summary	New Behavior
For read-only storage handlers, existence of directories was checked upon start up of an engine, even though an engine never needs to access these directories.	Existence of directories for file-based storage handlers is only performed for those directories which are actively needed, and lazily for read-only handlers.

3.46 CORE-18313 (Daylight Saving Time Offset)

Summary	New Behavior
When the **Use Unified Time Zone and Date Pattern for Conversion and Export** feature was active, dates in the near native view did not respect daylight saving time.	For each document to be rendered in the near native view, the document's sent/delivery date is retrieved in order to compute the time zone offset for daylight saving time, if appropriate.

3.47 CORE-18424 (Republishing)

Summary	New Behavior
On republish, conversions and redactions were invalidated due to the possibility of changes in natives.	On republish, we now try to identify whether a native or image has been changed. If changed, or if the state cannot be easily determined, we still invalidate; otherwise, conversions and redactions are kept.

3.48 CORE-18467 (Word Map)

Summary	New Behavior
Manual interventions temporarily set checkpoints aside and this could lead to a situation where replaying the words of the checkpoints overwrote existing word map entries without being noticed.	Additional sanity check was added to prevent overwriting existing words in a word map with other words from replaying a checkpoint. Note that this situation can only happen by manual intervention of the transaction recovery and in such cases the engine start up will be aborted.

3.49 CORE-18489 (CSV Merge)

Summary	New Behavior
A recently introduced storage type caused the CSV merge wizard to block, warning the target had no writable storage.	The CSV wizard now handles all types of writable storage.

3.50 CORE-18530 (Search)

Summary	New Behavior
Queries with proximity operators where one of the arguments contained an apostrophe (for instance O'Reilly /3 William) caused a parse error in Axcelerate 5, but worked fine in other search interfaces.	Queries with proximity operators where one of the arguments contains an apostrophe now work fine in all search interfaces.

3.51 CORE-18552 (Exception Handling)

Summary	New Behavior
Erroneous exception handling could cause retrainer service to become unusable, blocking engine save operations.	Exception handling has been fixed.

3.52 CORE-18587 (Production)

Summary	New Behavior
A batch server was turned off while a production job was running. The production job failed and persisted the exception. The engine - as intended - rescheduled the job on a different batch server. Although the job was restarted it still remembered the initial exception and would throw it over and over again.	The production job does not memorize its exception state upon restart anymore.

3.53 CORE-18629 (Index Configuration Update)

Summary	New Behavior
When updating the index configuration while an engine was running, an inconsistent state could be created when there were checkpoints present that modified the wordmap upon replay.	When updating the index configuration while an engine is running, and there are checkpoints present that would lead to an inconsistent configuration state, the operation is aborted. The user is asked to save the engine and then perform the update of the index configuration again.

3.54 CORE-18670 (Viewer Panel)

Summary	New Behavior
The system handled occasional connection issues in a bad way; which caused latency of about 2 minutes in the **Viewer** panel. In the worst case, this state became permanent until the APP servers were restarted.	Failure recovery is robust with respect to occasional connection failures and repairs itself.

3.55 CORE-18705 (Search Query Editor)

Summary	New Behavior
Searching for multiple numeric identifiers could cause an engine crash if all documents of the corresponding numeric identifier prefix had been removed.	In this special case, the engine does not crash anymore.

3.56 CORE-18858 (Session Expiration)

Summary	New Behavior
Sessions could not expire if the **Key Performance Indicator** feature was active. This feature is active by default.	Sessions can expire.

3.57 FOUND-3331 (Crawls)

Summary	New Behavior
Notes parser configuration contained wrong expression for **Reply Notice** forms.	Notes parser configuration contains correct expression for **Reply Notice** forms.

3.58 FOUND-6207 (Datamodel Node Expansion)

Summary	New Behavior
Sometimes opening the datamodel node in NAT took several seconds. This happened after a restart of Masterservice and when the internal cache was outdated.	The slow database query loading the relevant information was replaced by a more efficient one. The node will expand faster now even if data is loaded from the database instead of the cache.

3.59 FOUND-9321 (Crawls)

Summary	New Behavior
HTML documents were filtered if a MIME type filter existed for text/xml. Additionally, these documents created two filter events: one for "unsplit attachment" and one for the document itself.	HTML documents are not filtered if a MIME type filter exists for text/xml. Filter events were also fixed.

3.60 FOUND-10594 (Crawls)

Summary	New Behavior
meta_rfc5322_thread_path could contain wrong entries if elements contain $ characters.	**meta_rfc5322_thread_path** is correctly built even if elements contain $ characters.

3.61 FOUND-10682 (Crawls)

Summary	New Behavior
Data source could crash if the recipient table of an Outlook mail is corrupt.	Data source does not crash if the recipient table of an Outlook mail is corrupt.

3.62 FOUND-10729 (Crawls)

Summary	New Behavior
Microsoft XPS documents got file extension PPT in native storage.	Microsoft XPS documents get file extension XPS in native storage.

3.63 FOUND-10744 (Publish)

Summary	New Behavior
DXL-files (or NSF-entries) containing colon characters in item names produced invalid XML which could not be published.	Colon characters in item names in DXL-files (or NSF-entries) are encoded in the XML and can be published.

3.64 FOUND-10764 (Tomcat Update)

Summary	New Behavior
Tomcat version used in CORE was 8.0.33.	Tomcat version used in CORE is now 8.0.39 which fixes (among others) CVE-2016-3092.

3.65 FOUND-10892 (Crawls)

Summary	New Behavior
Corrupt PST folder structures could cause JobExceptions which could ultimately lead to aborted crawls.	Corrupt PST folder structures are tracked as Archive errors and do not abort the crawl.

4 New Feature Descriptions for Axcelerate 5.10.0

4.1 Ad-Hoc Batching (AXC-19919)

Behavior before change

The **Review Workflow** panel was used to batch documents from the **Analysis** page. This required prior project setup where at least one review workflow was added to the universe and often the default batching templates needed to be customized to match project specifications.

Behavior after change

Documents can now be batched from the **All Documents** universe using the new **Batch Documents** wizard, located on the **Analysis** page in the **Actions** menu for the **Results** list. The wizard allows users to add review workflows and configure custom batch settings on-the-fly. If prior project setup was conducted, users also have the option to select existing review workflows and preconfigured batching templates.
If documents are missing from the universe, the user is warned those documents will be skipped from the batching.

Known Limitations

Ad hoc batching can only be used with the **All Documents** universe. The original batching functionality remains and can be used to batch documents from custom universes.

4.2 Application-Level User Security (AXC-12625)

Behavior before change

Access to Axcelerate 5 features was controlled by assigning user groups to one or multiple of the following roles:

- Reviewer
- Power Reviewer
- Power User
- Case Manager

Behavior after change

In addition to the existing role-based security, the following feature can be added or removed individually for individual user groups: **Download native files**.

4.3 Axcelerate 5 Application Cloning (AXC-36956)

Behavior before change

For a new Axcelerate 5 matter, no application-sided settings were made.

Behavior after change

For a new Axcelerate 5 matter cloned from an existing one, the following application-sided settings are copied into the new matter:

- Fields and values, including **Review State** field values
- Tagging arrangements, including the corresponding fields and values, tagging rules, quick tags, and the mapping to review states

Caution: The security for arrangements is not copied, i.e., it remains set to **Everyone** to make sure all users in the new matter have access. If an arrangement in the new application requires security, adjust the user group access on the **Administration > Arrangements** page.

Changes to Default Configuration

The axcelerate.standaloneSystemTemplate-Axc5-v2.config.xml has a new property enabled by default (litigation_ngTemplatable).

Known Limitations

The matter to clone must be in the same pod and must be both running and enabled for cloning (new internal application configuration/enabled by default for axcelerate.standaloneSystemTemplate-Axc5-v2.config). Cloning an offline matter or a matter existing in a different pod is not possible.
All settings are cloned, i.e., you cannot pick and choose which settings are copied to the new matter.

4.4 Default Arrangements Enabled for Corresponding Review States (AXC-36238)

Behavior before change

The default tagging arrangements were not enabled for any review state.

Behavior after change

Two default associations will exist in new Axcelerate Review & Analysis matters:

- The **1st Level** tagging arrangement will be enabled for the **First Pass Review Complete** review state.
- The **2nd Level** tagging arrangement will be enabled for the **Second Pass Review Complete** review state.

Changes to Default Configuration

The default **1st Level** and **2nd Level** tagging arrangement are enabled for the corresponding **First / Second Pass Review Complete** review state.
Existing matters are not affected, this change is only used for newly created Axcelerate 5 matters.

4.5 Links Added to Login Page (AXC-36238)

Behavior before change

The Axcelerate 5 login page contained two links: **Reset password** and **Support terms**.

Behavior after change

The Axcelerate 5 login page contains up to five links: Reset password, Support terms, Training & Certification, Mind over Matters blog, and Online Help. Via configuration these links can be changed.

Changes to Default Configuration

New settings in the webtier.user.properties can be used to control the label and target for up to five links on the login page of Axcelerate 5:

- axcng.navigation.root.branding.section1.label=Passwort zurücksetzen
- axcng.navigation.root.branding.section1.url=http://zurueck.recomind.de
- axcng.navigation.root.branding.section2.label=Nutzungsbebedingungen
- axcng.navigation.root.branding.section2.url=https://terms.recommind.de
- axcng.navigation.root.branding.section3.label=AXCTRAX @ YouTube
- axcng.navigation.root.branding.section3.
  url=https://www.youtube.com/channel/UC6NHYGswQju66MOgRrbMiXA
- axcng.navigation.root.branding.section4.label=
- axcng.navigation.root.branding.section4.url=/
- axcng.navigation.root.branding.section5.label=LaunchPad
- axcng.navigation.root.branding.section5.url=
  https://launchpad.rmcloud.recommind.com/

Per default, the three new links reference the Training & Certification center, Mind over Matters blog, and Axcelerate 5 online help portal.

Known Limitations

The total width for all five links is limited. Exceeding the width limit will cause the links to wrap and not display well.

4.6 Oracle Outside In Upgrade to 8.5.3 (FOUND-9845)

Behavior before change

Outside In 8.5.1 was used.

Behavior after change

Outside In 8.5.3 will be used on new crawls for new Axcelerate 5 projects. This version supports new formats for data loading, including: Microsoft Word, Excel and PowerPoint 2016; Microsoft Outlook 2011 for Mac (OLM and EML); Corel WordPerfect; Corel Quattro Pro; Corel Presentations; Corel Draw X7; AutoCAD 2015.
Outside In 8.5.1 will be used for new and old data sources of existing projects as a default, to avoid unexpected changes in hash values.
Administrators have the ability to change the data source configuration of updated projects to support the new data formats.

5 Issues Resolved in Axcelerate 5.10.0

5.1 FOUND-7953 (Error Messaging)

Summary	New Behavior
Last Crawl Details dialog in CORE Administration showed an error message without any reference to the error reason. This occurred when the index engine was not running.	The **Last Crawl Details** dialog now shows more information on the error reason when the index engine is not running.

5.2 AXC-7895 (Conversion)

Summary	New Behavior
Vertical axis titles not properly rotated during conversion to PDF.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.3 AXC-7898 (Conversion)

Summary	New Behavior
Shifted gridlines during conversion to PDF.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.4 AXC-21532 (Conversion)

Summary	New Behavior
Huge HTML tables could not be converted.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.5 AXC-27882 (Production Workflow Deletion)

Summary	New Behavior
A production workflow could not be deleted after a production was exported through CORE Administration.	The deletion of a production workflow does not lead to an exception when production sets related to the workflow were exported through CORE Administration.

5.6 AXC-34830 (Conversion)

Summary	New Behavior
The conversion to TIFF performance was slow for certain documents.	The conversion to TIFF performance has been improved.

5.7 AXC-34831 (Production)

Summary	New Behavior
Produced email shows incorrect time offset.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.8 AXC-36149 (Conversion)

Summary	New Behavior
Some documents could not be produced to TIFF successfully due to "IPC pipe closed" conversion issues.	These documents are now produced successfully.

5.9 AXC-36649 (SQE Search)

Summary	New Behavior
Some SQE searches time out after 60 seconds.	The performance of the SQE search has been improved by reducing the number of requests from the browser and by combining searches.

5.10 AXC-36884 (Engine Load)

Summary	New Behavior
Opening the **Review Workflows** page generated substantial load on the engine.	Opening the **Review Workflows** page no longer generates a substantial load on the engine.

5.11 AXC-36937 (Conversion)

Summary	New Behavior
PDF export misses some parts for a html file.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.12 AXC-37237 (Mouse Pointer)

Summary	New Behavior
When hovering over the disabled icon for previous or next document in the **Viewer** panel, the mouse pointer turned into the hand pointer. This indicated the button was clickable, but it was actually not.	When hovering over the disabled icon for previous or next document in the **Viewer** panel, the mouse pointer now stays in the default state.

5.13 AXC-37294 (Batch Display Options)

Summary	New Behavior
Check box for **Show Associated Results outside of batch** was shown as disabled in the **Batch Display** fly-in, although the function itself was enabled by default for batch review.	Check box for **Show Associated Results outside of batch** is now shown as enabled in the **Batch Display** fly-in, as the function is enabled by default for batch review.

5.14 AXC-37311 (Processing Page)

Summary	New Behavior
When using the **# Users** filter in the **System Properties** Smart Filter group on the Axcelerate 5 **Processing** page, it was not possible to filter for matters with zero users because the input form only accepted values greater than zero.	It is now possible to filter to matters that have zero users.

5.15 AXC-37313 (Two-Screen Mode)

Summary	New Behavior
In FireFox, when users opened the **Analysis** page in two-screen mode, the second screen was locked.	In FireFox, the second screen is no longer locked when users open the **Analysis** page in two-screen mode.

5.16 AXC-37327 (Quick Tags)

Summary	New Behavior
When users tried to bulk tag documents using the **Last Applied** quick tag, there was no error message shown to indicate the action was not allowed.	An error message is now shown when users try to bulk tag documents using the **Last Applied** quick tag.

5.17 AXC-37351 (Login Page)

Summary	New Behavior
In the case of a fatal error, the Axcelerate 5 error page was grayed out and users were not shown a link to log in again.	In the case of a fatal error, the Axcelerate 5 error page displays correctly.

5.18 AXC-37458 (Production)

Summary	New Behavior
Production snapshot was incorrectly flagged as "Stale."	Stale icon is now only rendered if last snapshot date is a valid date and older than 7 days.

5.19 AXC-37464 (Production)

Summary	New Behavior
Projects in S3 storage deduced file extensions of exported production native files based on the file's MIME type, which is not always unique and is inconsistent with how file-based projects deduce file extensions.	The exported file extensions are correct: new records stored in S3 explicitly provide fast access to the extension, and old ones fetch the correct extension from the index.

5.20 AXC-37635 (Smart Filters)

Summary	New Behavior
If a filter value contained a *NOT* prefix, e.g., *Not Privileged*, and the filter was set to **Exclude** that value, the **Current Criteria** panel showed only the value's display name, e.g., *Privileged*, and did not show the *NOT* prefix of the value's display name.	If a filter is set to **Exclude** a value containing a *NOT* prefix, the value's full display name, e.g., *Not Privileged*, is correctly shown in the **Current Criteria** panel.

5.21 AXC-37789 (Production)

Summary	New Behavior
When a production workflow was exported via the **Productions** tab of CORE Administration, the corresponding default production rules were not deleted as expected.	Default production rules are always deleted when their corresponding production workflow is deleted.

5.22 AXC-37955 (Date Values)

Summary	New Behavior
Date values presented with a time zone, e.g., on the **Analysis** page, reports and document history always showed the daylight saving time (DST) variant of the user's preference time zone. The date values, however, were converted to the user's time zone respecting DST. This resulted in wrong time values (shifted by DST offset) to be displayed for dates not within the DST interval of the user's time zone.	The displayed time zone now matches the DST variant of the date to be displayed, i.e., dates within the DST interval of the user's preference time zone will show the DST variant and dates outside of the DST interval will display the regular time zone variant.

5.23 CORE-16015 (Crawls)

Summary	New Behavior
Incorrect Oracle Outside In PDF conversion.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.24 CORE-16646 (OCR)

Summary	New Behavior
If OCR was executed for an incomplete attachment family and then new documents were indexed and the engine was flushed, the documents for which OCR was executed could have been flagged as MISSING_PARENT and/or MISSING_ROOT by attachment processing.	When inserting new documents and flushing the engine, the attachment processing no longer affects documents for which OCR had been executed previously, in particular these are not incorrectly flagged as MISSING_PARENT and/or MISSING_ROOT.

5.25 CORE-17317 (Field Value Creation)

Summary	New Behavior
Creating a field value with properties via the CORE API happened in two steps and a downtime of the engine between them could leave the field value created without its properties.	Creating a field value with properties via the CORE API is now one transaction, i.e., a downtime of the engine during that operation does not leave the field value created without its properties.

5.26 CORE-17394 (JDBC Crawls)

Summary	New Behavior
Incremental JDBC crawls failed if all chunk entries (document IDs) had occurred in previous chunks.	Incremental JDBC crawls continue after chunks that contain only entries that have occurred in previous chunks.

5.27 CORE-17406 (DXL parser)

Summary	New Behavior
DXL parser sometimes caused HTML entities to be shown as document content.	DXL parser does not cause HTML entities to be shown as document content.

5.28 CORE-17474 (Field-based Search)

Summary	New Behavior
Field-based searches were extended by automatically detected phrases, resulting in matches containing these phrases anywhere (not necessarily in the specified field).	Field-based searches are no longer extended by automatically detected phrases.

5.29 CORE-17481 (Storage Handlers)

Summary	New Behavior
The storage verification did not work properly after a read-only storage handler had been erroneously upgraded to SQL.	The storage verification ignores any read-only handler that has been upgraded to SQL and works properly.

5.30 CORE-17498 (Engine Saves)

Summary	New Behavior
High server peak memory usage could cause engines to become unresponsive.	Internal exception handling has been improved such that low memory conditions cannot cause engines to become unresponsive anymore.

5.31 CORE-17574 (CSV Load)

Summary	New Behavior
If the content of the field **rm_attachmentprimary** sent by a CSV load was invalid or contradicted the value of **rm_attachmentroot**, then this broke features such as family sorting.	If the content of the field **rm_attachmentprimary** sent by a CSV load is invalid or contradicts the value of **rm_attachmentprimary**, then it is automatically corrected at insertion time. Any correction other than adding a missing value leads to the document being tagged as **Content modification** in the **Exception Class** Smart Filter and as **Resolved inconsistent or invalid attachment primary information** in the **Exception Type** Smart Filter.

5.32 CORE-17833 (Box Native Files)

Summary	New Behavior
Box native files were always downloaded during re-crawls despite their unchanged modification date.	Box native files are only downloaded during re-crawls if necessary.

5.33 CORE-17933 (Native Conversion)

Summary	New Behavior
The conversion of some legacy Microsoft Word and PowerPoint documents failed with the error message, "The document-file is corrupt because it contains no storage object."	The affected legacy Microsoft Word and PowerPoint documents can be converted correctly.

5.34 CORE-17944 (Metadata Display)

Summary	New Behavior
The date 1.1.1970 00:00 was displayed as an empty string in the metadata section of the **HTML-View** that is shown, e.g., in CORE Administration. All other dates including 31.12.1969 23:59 and 1.1.1970 00:01 were displayed correctly.	The date 1.1.1970 00:00 is now displayed correctly in the metadata section of the **HTML-View**.

5.35 CORE-18010 (Application Deletion)

Summary	New Behavior
A missing database connection could result in inconsistent project relations in CORE Administration when deleting an Axcelerate Ingestion or Axcelerate Review & Analysis application: the involved entities, e.g., engines, were deleted, but still shown in CORE Administration.	A failing database connection when clearing a case in CORE Administration will not lead to broken relations. The removed entities are deleted and the user is informed to delete the database entries manually. Leftover entries do not harm further operations.

5.36 CORE-18028 (Index Engine)

Summary	New Behavior
Write-locking an index engine followed by a save, blocked field value modifications.	Field value modifications are processed even if the index engine is write-locked. Additionally, save jobs can be canceled while waiting for the "write unlock."

5.37 FOUND-6877 (Crawls)

Summary	New Behavior
Crawling MSG files resulted in Oracle Outside In error.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.38 FOUND-7526 (Conversion)

Summary	New Behavior
Oracle Outside In failed to convert certain .pdf and .doc files.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.39 FOUND-7953 (Error Messaging)

Summary	New Behavior
Last Crawl Details dialog in CORE Administration showed an error message without any reference to the error reason. This occurred when the index engine was not running.	The **Last Crawl Details** dialog now shows more information on the error reason when the index engine is not running.

5.40 FOUND-8803 (Crawls)

Summary	New Behavior
Documents created when loading Microsoft Project files could contain temporary file names.	Documents created when loading Microsoft Project files contain original file names.

5.41 FOUND-8809 (Crawls)

Summary	New Behavior
Oracle Outside In added white space in text strings for ICS files.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.42 FOUND-8811 (Crawls)

Summary	New Behavior
Oracle Outside In added \n characters in text strings for VCF files.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.43 FOUND-9554 (Exception Handling)

Summary	New Behavior
Documents with multiple paragraphs caused postprocessing exceptions if an old (not property-) postprocessor was configured.	Documents with multiple paragraphs do not cause postprocessing exceptions if old (not property-) postprocessors are configured.

5.44 FOUND-9562 (Crawls)

Summary	New Behavior
Oracle Outside In erroneously reported PDF file as corrupt during crawl.	Fixed by upgrade to Oracle Outside In 8.5.3 (see FOUND-9845).

5.45 FOUND-9975 (Crawls)

Summary	New Behavior
Large binary Outlook HTML mail bodies could cause OutOfMemory errors if the mail did not contain an RTF body.	Large binary Outlook HTML mail bodies do not cause OutOfMemory errors.

5.46 FOUND-10208 (Crawls)

Summary	New Behavior
With Microsoft Outlook 2013 the crawler crashed if a PST contained a corrupt entry that could not be exported as a native MSG file.	PSTs that contain an entry that cannot be exported as a native MSG file are not processed and are marked as corrupt.

5.47 FOUND-10217 (Crawls)

Summary	New Behavior
Embedding names that contained illegal file system characters following the last dot caused exceptions.	Embedding names containing illegal file system characters following the last dot do not cause exceptions.

5.48 FOUND-10288 (Crawls)

Summary	New Behavior
MBOX files with mixed newline characters were not processed correctly.	MBOX files with mixed newline characters are processed correctly.

5.49 FOUND-10336 (Reporting)

Summary	New Behavior
A fix user name and password were used to generate reports in Axcelerate Ingestion.	Information from the current user's session is used to generate the reports.

5.50 FOUND-10563 (Index Engine Selection)

Summary	New Behavior
LaunchPad automatically selected index engine hosts based on CPU load.	LaunchPad automatically selects index engine hosts based on free RAM.

5.51 FOUND-10588 (Lotus Notes NSF File Crawls)

Summary	New Behavior
The default CORE configuration was not compatible with Lotus Notes Domino in case RMI over TLS was enabled.	The default CORE configuration is compatible with Lotus Notes Domino in case RMI over TLS is enabled.

5.52 FOUND-10593 (Crawls)

Summary	New Behavior
In certain scenarios the number of waiting jobs was very high, leading to high memory consumption.	An optimization was implemented that notably decreases the number of scenarios where this can happen.

5.53 FOUND-10623 (Crawls)

Summary	New Behavior
There was a performance regression while extracting native files for PST entries and attached MAPI messages.	Performance for extracting native files for PST entries and attached MAPI messages is restored.

6 New Feature Descriptions for Axcelerate 5.9.2

6.1 Microsoft SharePoint 2013/Online Connector (Reference ID: CORE-14693)

Behavior before change

Ther was no connector supporting Microsoft SharePoint 2013 or Online for on premise installations.

Behavior after change

The new SharePoint 2013/Online connector supports Microsoft SharePoint 2013 and SharePoint Office 365 for on premise installations.
With this connector, all content types, attachments and librairy documents are indexed. In addition, indexing of containers, user profiles and hidden items, or items tagged as *NoCrawl* is possible.
The connector offers various data scope filters, such as content type and list template type inclusion/exclusion. Incremental crawls with this connector support deletion.

Known Limitations

- Load balanced URLs are not supported.
- SharePoint farm level is only supported for the User Profile crawl.
- Native files can only be copied at crawl time.
- SharePoint webs and sub webs located on different ports are not supported.

6.2 Microsoft Exchange 2013/Online Connector (Reference ID: CORE-13085)

Behavior before change

There was no connector supporting Microsoft Exchange 2013/Online for on premise installations.

Behavior after change

With the new Microsoft Exchange 2013/Online connector, Microsoft Exchange 2013 and Office 365 for Exchange mailboxes can be crawled for on premise installations. The connector features date range filtering, mailbox inclusion/exclusion, email folder inclusion/exclusion and item type inclusion/exclusion filtering. In addition, it can crawl journal mailboxes for Exchange 2013 OnPremise.

Known Limitations

- Exchange address resolution is not supported.
- Distribution list expansion is not supported.
- Exchange 2013 OnPremise: only one journal mailbox per data source can be crawled.
- No auto discovery support. User must specify Exchange web service URL.
- No public folder support.
- No search folder support.
- Calendar entries accepted into the calendar (ICS files) will not contain any attachments in the native file itself as this is not supported by Microsoft. However, attachments are maintained within the CORE index, are displayed in the user interface and can be exported. This may be an issue for productions that only contain native files. These files should therefore be produced as images.

7 Issues Resolved in Axcelerate 5.9.2

7.1 AXC-34830 (Conversion)

Summary	New Behavior
The conversion to TIFF performance was slow for certain documents.	The conversion to TIFF performance has been improved.

7.2 AXC-36149 (TIFF production)

Summary	New Behavior
Some documents could not be produced to TIFF successfully due to "IPC pipe closed" conversion issues.	These documents are now produced successfully.

7.3 AXC-36649 (Search Query Editor)

Summary	New Behavior
Some Search Query Editor searches timed out after 60 seconds.	The performance of the Search Query Editor search has been improved by reducing the number of requests from the browser and by combining searches.

7.4 AXC-38148 (User preferences)

Summary	New Behavior
Saving user preferences could lead to an error if some unknown parameters were sent to the server.	Unknown parameters will be ignored so that saving user preferences will not fail.

7.5 AXC-38305 (Tomcat security)

Summary	New Behavior
There were several security issues in Tomcat 7.0.63.	Tomcat is updated to 7.0.73 which contains fixes for these security issues.

7.6 AXC-38437 (Associated Results Searches)

Summary	New Behavior
Search requests for associated results are executed twice in parallel when navigating through the Associated Results lists.	Search requests for associated results are always executed once at a time.

7.7 CORE-9061 (CSV load)

Summary	New Behavior
CSV data load could produce nested tags inside an "rm_numeric_identifier" field.	The XML parser is fixed and produces correct data.

7.8 CORE-9390 (XML document structure)

Summary	New Behavior
Document XML structure could become invalid during publishing, due to wrong place of <rm_id.profile> elements.	The internal <rm_id.profile> elements are now always added to the end of the document's XML structure.

7.9 CORE-14923 (SSL for SQL connections)

Summary	New Behavior
SSL could only be configured if the system had the non-standard configuration "Use no SQL connection pool". This meant that SSL was not possible for SQL-managed storages.	The tool used to configure the SQL connection pool accepts a parameter for SSL and configures the connection pool and its consumers accordingly.

7.10 CORE-17561 (OCR)

Summary	New Behavior
Due to the resource management in distributed computing environments (OCR resource != conversion resource) the engine job manager occasionally started all document jobs for a bulk OCR operation at once. Most of those jobs had to wait in a blocked state within the engine and wait for their respective resource. So, pausing a bulk operation was not possible, since all document jobs were already started and had to be finished first.	The engine prevents starting new document jobs for a macro job when a certain threshold of blocked jobs for that macro job is exceeded.

7.11 CORE-17729 (Export)

Summary	New Behavior
In meta engine projects, the CSV export from some sub-engines could fail.	Likelihood of failure is reduced and the error handling is improved.

7.12 CORE-17845 (Conversion)

Summary	New Behavior
On-the-fly conversions had a long timeout which caused batch servers to receive and hold too many jobs.	Timeout for on-the-fly jobs is back to two minutes.

7.13 CORE-18022 (TIFF conversion, OCR)

Summary	New Behavior
Axcelerate Classic conversion and Bulk OCR did not cover exception documents.	Axcelerate Classic conversion and Bulk OCR cover exception documents. Only documents with these exception classes are excluded from Axcelerate Classic conversion and Bulk OCR: File format, Archive, Password protected/encrypted.

7.14 CORE-18231 (Batch servers)

Summary	New Behavior
A bug in the job grid caused the engine to ignore offline states of batch servers. This caused jobs to be sent to batch servers that were in offline state.	The engine now properly checks the availability of processing resources. In addition, a wide range of exceptions regarding the connection to the batch processor will now result in a retry for the job affected.

7.15 FOUND-10217 (Embeddings)

Summary	New Behavior
Embedding names containing illegal file system characters following the last dot caused exceptions.	Embedding names containing illegal file system characters following the last dot don't cause exceptions.

7.16 FOUND-10593 (Box connector)

Summary	New Behavior
In certain scenarios the number of waiting jobs was very high, leading to high memory consumption.	An optimization is implemented that notably decreases the number of scenarios where this can happen.

7.17 FOUND-10623 (Performance)

Summary	New Behavior
There was a performance regression while extracting native files for PST entries and attached MAPI messages.	Performance for extracting native files for PST entries and attached MAPI messages is restored.

7.18 FOUND-10682 (Datasource)

Summary	New Behavior
Datasource could crash if the recipient table of an Outlook mail was corrupt.	Datasource doesn't crash if the recipient table of an Outlook mail is corrupt.

7.19 FOUND-10764 (Tomcat security)

Summary	New Behavior
Tomcat version used in CORE was 8.0.33.	Tomcat version used in CORE is now 8.0.39 which fixes (among others) CVE-2016-3092.

8 New Feature Descriptions for Axcelerate 5.9.1

8.1 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.

8.2 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.

8.3 "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.

8.4 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.

8.5 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".

8.6 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.

8.7 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.

9 Issues Resolved in Axcelerate 5.9.1

9.1 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.

9.2 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.

9.3 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.

9.4 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.

9.5 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.

9.6 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.

9.7 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.

9.8 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.

9.9 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.

9.10 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.

9.11 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.

9.12 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.

9.13 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).

9.14 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.

9.15 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.

9.16 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.

9.17 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.

9.18 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.

9.19 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.

9.20 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.

9.21 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.

9.22 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.

9.23 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.

9.24 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.

9.25 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.

9.26 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.

9.27 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.

9.28 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.

9.29 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.

9.30 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.

9.31 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.

9.32 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.

9.33 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.

9.34 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.

9.35 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.

9.36 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.

9.37 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.

9.38 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.

9.39 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.

9.40 FOUND-7160 (Log messages)

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

9.41 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.

9.42 FOUND-9588 (0 byte file handling)

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

9.43 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.

9.44 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.

9.45 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.

9.46 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.

9.47 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.

9.48 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.

9.49 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.

9.50 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.

9.51 FOUND-10214 (SSO login)

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

9.52 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.

9.53 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.

9.54 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.

10 New Feature Descriptions for Axcelerate 5.9.0

10.1 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.

10.2 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.

10.3 "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.

10.4 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.

10.5 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.

10.6 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.

10.7 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.

11 Issues Resolved in Axcelerate 5.9.0

11.1 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.

11.2 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.

11.3 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.

11.4 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.

11.5 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.

11.6 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.

11.7 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.

11.8 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.

11.9 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.

11.10 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.

11.11 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.

11.12 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.

11.13 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.

11.14 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.

11.15 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.

11.16 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.

11.17 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.

11.18 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.

11.19 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.

11.20 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.

11.21 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.

11.22 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.

11.23 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.

11.24 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.

11.25 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.

11.26 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.

11.27 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.

11.28 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.

11.29 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.

11.30 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.

11.31 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.

11.32 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.

11.33 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.

11.34 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.

11.35 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.

11.36 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.

11.37 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.

11.38 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.

11.39 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.

11.40 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.

11.41 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.

11.42 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.

11.43 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.

11.44 FOUND-7160 (Log messages)

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

11.45 FOUND-9588 (0 byte file handling)

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

11.46 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.

11.47 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.

11.48 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.

11.49 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.

11.50 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.

11.51 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.

12 New Feature Descriptions in version 5.7.0 Update 1

12.1 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.

12.2 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.

12.3 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.

12.4 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.

13 Issues Resolved in 5.7.0 Update 1

13.1 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.

13.2 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.

13.3 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.

13.4 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.

13.5 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.

13.6 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.

13.7 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.

13.8 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."

13.9 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.

13.10 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.

13.11 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.

13.12 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.

13.13 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.

13.14 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.

13.15 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.

13.16 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.

13.17 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.

13.18 FOUND-9633 (Large cache)

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

13.19 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.

13.20 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.

13.21 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.

13.22 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.

13.23 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.

14 New Feature Descriptions for Axcelerate 5.8.0

14.1 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.

14.2 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.

14.3 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.

14.4 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.

14.5 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.

14.6 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. .

14.7 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.

15 Issues Resolved in Axcelerate 5.8.0

15.1 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".

15.2 AXC-32122 (Robustness and stability)

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

15.3 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.

15.4 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.

15.5 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.

15.6 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.

15.7 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.

15.8 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.

15.9 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.

15.10 AXC-36567 (Viewer)

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

15.11 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.

15.12 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."

15.13 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.

15.14 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 fFlter and the **Assignee** Smart Filter on the **Assignments** page now show an **Apply** link rather than document counts.

15.15 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.

15.16 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.

15.17 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.

15.18 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.

15.19 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.

15.20 CORE-15862 (LiveLink connector)

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

15.21 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"

15.22 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.

15.23 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.

15.24 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.

15.25 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.

15.26 CORE-16303 (SharePoint XML)

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

15.27 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'.

15.28 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.

15.29 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.

15.30 CORE-16492 (Robustness and stability)

Summary	New Behavior
An out-of-memory of physical RAM caused a server down-time.	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.

15.31 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.

15.32 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.

15.33 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.

15.34 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.

15.35 CORE-16641 (CertificateUtilities)

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

15.36 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.

15.37 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.

15.38 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.

15.39 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.

15.40 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.

15.41 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.

15.42 FOUND-9633 (Large cache)

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

15.43 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.

15.44 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.

15.45 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.

15.46 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.

15.47 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.

15.48 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.

16 Issues Resolved in Axcelerate 5.7.2

16.1 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.

16.2 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.

16.3 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).

16.4 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.

17 Issues Resolved in Axcelerate 5.7.1

17.1 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.

18 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 : Axcelerate@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.

19 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