

Changelog yuuvvis® Momentum - Final - 2023 autumn

Schlüssel	Zusammenfassung	Beschreibung
AI-1683	Implement stricter matching in transformer based ee pipeline	<p>Due to the limitations of our data (we do not possess precise labels and positions of fields we're trying to extract) we must match values from metadata to what is in the document. This is a process that can result in errors that are not related to code (everything is working as expected but just can't handle certain edge cases), especially when using all matchers simultaneously as can be seen in BUG AI-1659. Since matchers themselves can't be fixed, these issues can be remedied by improving the quality of the dataset we're using for training, by implementing stricter document selection.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• Implement logic for stricter matching.<ul style="list-style-type: none">◦ For a document to be considered valid (from a training dataset standpoint) it must have matched values for each field that is being extracted.<ul style="list-style-type: none">▪ If a value is missing from metadata the document will be discarded▪ If a value can be found in metadata but that value can't be matched in the document itself, that document will be discarded◦ All field values must be unique in regards to other metadata values<ul style="list-style-type: none">▪ If total and base amount are equal the document will be discarded.
BUD-2368	New MetricService is used in the tenant management API	<p>The old MericService in the tenant management API should not be used anymore.</p> <p>The tenant-management service asks the metricservice https://git.optimal-systems.org/enaio/red/tenant-management-api/-/blob/develop/src/main/java/com/os/services/management/dao/APITMDao.java#L243</p>
BUD-2452	Switch to Core-API-Model (Java Components)	<p>As a programmer, you are supported to use a future-friendly domain model. Core-API-Domain Models are future-friendly.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• It is checked if only in API-Web an adjustment has to be made or if further adjustments are necessary.• API-Web is adapted for use of the Core-API-Domain Models
BUD-2454	Tenant Management API supports Keycloak version 21	<p>As an administrator, I want to use the Tenant Management API service with the released Keycloak version 21 to be sure to have a secured system running.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• All existing endpoints are working well accessing Keycloak 21.• Unit-tests are adapted
BUD-2595	Adaptation of the client to the new endpoint for getting tasks	<p>Due to adaption of endpoints the tasks of user shall be taken from /api/bpm/inbox in api-web</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• the tasks of user shall be taken from /api/bpm/inbo
BUD-2623	API web supports an additional search operator	<p>As a programmer, I am supported by a search operator to force the search for the field to be a full-text search, even if it contains wildcards. Thus, these searches are case-insensitive and wildcard-capable.</p> <p>Acceptance criteria</p> <p>API web provides:</p> <pre>SELECT * FROM system:object WHERE yaField CONTAINS(the value)</pre>
BUD-2647	Keycloak is upgraded to version 22	<p>As a system operator I want to use keycloak in version 22.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• The Upgrade to keycloak version 22 is successfully implemented,
BUD-2648	Keycloak themes are adopted for the docker image with version 22	<p>As a system operator I want to use keycloak in version 22.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• The keycloak themes are adopted for the docker image with version 22.

ERA-8472	Tasks are shown in an overview	<p>As workflow administrator, I would like to see the tasks that currently exist in the system, in order to be able to manage them.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • Users with tenant management rights can see all tasks in the tenant where they are currently logged in • Following metadata is shown in overview <ul style="list-style-type: none"> ◦ ID ◦ Name ◦ Editor (resolved to readable name) ◦ Owner (resolved to readable name) ◦ Created ◦ Ended
ERA-8938	Process-initiator is shown in human-readable format in process-overview	<p>As workflow administrator, I would like to see who has started a process, in human-readable format so that I can understand the process-overview page better.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • Started by field in process-overview table does not show IDs but the name of user in the same format as in client: Lastname, Firstname (username) • The filter field allows search and selection of a user by its name, user's ID is used to execute the search in backend <ul style="list-style-type: none"> ◦ current implementation hides admin users in search results
ERA-9221	Tasks endpoint provides all tasks in the tenant to process administrators	<p>As an administrative API user of bpm-engine, I would like to see all tasks in the system (including those that are for other users), in order to have a complete overview of system behavior.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • The endpoint returns a list of tasks from the same tenant as the authenticated user. • A user without admin rights can only see his/her tasks (user is candidate or assignee) • A user with admin rights can see all tasks in the tenant. • An admin user can filter for tasks visible to a given user using the parameter <i>visibleTo</i> • Query parameter <i>isCompleted</i> allows filtering by task status (completed or not) <ul style="list-style-type: none"> ◦ by default (if parameter <i>isCompleted</i> is not specified), both running and completed tasks are returned • Field <i>endTime</i> is added to response
ERA-9344	User-service can use different databases for different tenants	<p>As a system integrator, I would like to be able to split data of different tenants to different relational databases, in order to better control the scaling of the system.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • Tenants are mapped to different databases • A default database instance can be defined, for all tenants that do not require particular database connection • The configuration can be either defined only for user-service, or reused from core-service • The same format for db-connections is used as for core-service: https://yuuvidevelop.atlassian.net/wiki/spaces/YM22WI/pages/338625772/application-dbs.yml
ERA-9380	bpm-engine uses Flowable 6.8	<p>As an operator of yuuvim Momentum, I would like that bpm-engine gets the newest features available in Flowable engine.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • Flowable version 6.8.x is used in the bpm-engine (currently it is 6.7.2)
ERA-9433	Endpoint for getting tasks in api-web supports additional filter parameters	<p>As a developer, I would like to have more filter options when getting tasks from backend, to be able to show only relevant tasks to the user.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • GET <code>api/bpm/tasks</code> supports additional filter request parameters: <ul style="list-style-type: none"> ◦ <code>processInstanceId</code> ◦ <code>processDefinitionId</code> ◦ <code>nameLikeIgnoreCase</code> ◦ <code>businessKeyLike</code> ◦ <code>createdBefore</code> ◦ <code>createdAfter</code> ◦ <code>completedBefore</code> ◦ <code>completedAfter</code> ◦ <code>isCompleted</code> ◦ <code>visibleTo</code> • GET <code>api/bpm/inbox</code> calls <code>getTasks</code> with queryParameter <code>isCompleted=false</code>, <code>visibleTo='ownUserId'</code>

ERA-9514	bpm-engine supports Keycloak v22	<p>As a system operator I want to use keycloak in version 22.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • it is possible to connect bpm-engine directly with keycloak • in Autumn 23 we use Keycloak22 • so, bpm-engine shall be capable of getting the data from new version of Keycloak
OKTO-5598	Automatic extraction of header data when importing e-mails	<p>As a user of the system, I would like it to automatically recognise when an email is imported that it is an email and assign the appropriate type and extract the metadata so that no further action is required.</p> <p>Currently, emails are not treated separately and behave like other documents. The mime type is determined and the text of the email body and the first level attachments are extracted.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • Emails can be imported via the API • EML and MSG files are accepted as e-mails • there is no conversion of the formats into another format • it is recognised during import that it is an e-mail • an object type for e-mails is available • when an email has been recognised, the metadata is mapped to the fields of the email object type • if errors occur, the e-mail is tagged with a defined error value • if the e-mail object type is already supplied, the fields are not filled, this is to ensure that the fields are not extracted again if they have already been processed externally <p>not observed:</p> <ul style="list-style-type: none"> • When updating an email object, the email metadata is NOT changed by this automatic mechanism, i.e. the mechanism only takes effect during the first import, not during the update. • no separate objects are created for recognised attachments
OKTO-5603	Support of multiple Elasticsearch Instances	<p>As a system integrator, I need the possibility to connect several Elasticsearch instances with a yMomentum system in order to achieve better scaling of the system and separation of the clients.</p> <p>It should also be possible to distribute different clients to elasticsearch instances with different performance.</p> <p>In order to realize high performance while the workload increases, it is now possible to configure multiple Elasticsearch instances.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • multiple instances of elasticsearch can be used in the system. • Individual instances can each be assigned to one or more clients • there is a default instance for clients not specifically assigned • The feature is documented • Cross-tenant currently only works for the default index
OKTO-5627	classification attributes are also available for reference fields	<p>It is currently possible to add additional <classification>attributes to property types and object types in the schema. This is currently not possible for property references.</p> <p>As a developer, I would also like to be able to specify <classification>-elements for property references in order to enrich the property references with additional information in the respective context.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • one or more <classification>-elements can also be used for property references in the schema • the feature is documented
OKTO-5648	Optionally store user attributes in Redis cache	<p>As one of the thousands of users of a Momentum platform, I don't want to experience any delays in the Momentum system's workflows in my daily work and I always want to be able to log in without any problems, even if I have a lot of roles.</p> <p>For a low risk of failures, I as the operator of a Momentum system want my cluster network to be loaded as little as possible and the available working memory to be allocated sparingly, even if thousands of users are performing their daily work with the system at the same time.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • the authentication service determines the user attributes via <i>user.info</i> hook and copies them into a Redis cache • The transfer method of the user attributes is adjustable for the purpose of downward compatibility. • the connection parameters of the <i>redis</i> jump profile are applied • if activated, the data connection to the Redis cache is only established in the following cases and then remains permanently maintained <ul style="list-style-type: none"> ◦ the Authentication Service restarts ◦ the connection to the Redis cache breaks. • Users with a role list larger than 8KB can log on to the system normally

OKTO-5649	Security library optionally reads user attributes from a Redis cache	<p>As one of the thousands of users of a Momentum platform, I don't want to experience any delays in the Momentum system's workflows in my daily work and I always want to be able to log in without any problems, even if I have a lot of roles.</p> <p>For a low risk of failures, I as the operator of a Momentum system want my cluster network to be loaded as little as possible and the available working memory to be allocated sparingly, even if thousands of users are performing their daily work with the system at the same time.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • Security library has a new version • the user attributes, including the user roles, are only transmitted to the services that also process these attributes • Each of these services receives the attributes via a data connection that is only established once, when it is started or, if necessary, when the connection is interrupted. • the user attributes are loaded from the Redis cache if they are missing in the incoming JWT • even with many HTTP connections open at the same time, the memory requirements of the services remain moderate <ul style="list-style-type: none"> ◦ Background: for each open HTTP connection, the request headers are kept in memory until the connection ends.
OKTO-5651	Extension option for overwriting endpoints in the API gateway	<p>As a developer, I would like to redirect API endpoints to project-specific services for projects in order to implement more appropriate call control.</p> <p>In this case, it is a matter of providing an interceptor for retrieving document files and additionally implementing the endpoints for retrieving parts of multi-page document files.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • endpoints in the API gateway can be configured to be redirected • additional endpoints can be hooked into the API • the feature is documented
OKTO-5659	Subject-specific schema version	<p>Currently, the schema version is assigned by the system. When importing into different systems without the staging mechanism via GIT, it can happen that the same version of the schema has different automatically assigned version numbers.</p> <p>As a developer, I would like to define which logical version of the schema it is, independently of the system, via a technical version number, in order to be able to transport solutions better and to be able to react to schema changes in a defined way.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • There is the possibility to include a technical version in the schema. • this is not changed by the system when it is imported
OKTO-5679	Repository supports predicate to select the repository to use	<p>As a system integrator, I would like to control the selection of the repository to be used via a predicate in order to enable a better distribution of the binary data to the individual repositories.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • the repository to be used can be selected via a predicate • the feature is documented
OKTO-5680	Offer SUM() function in SQL search	<p>As a developer, I would like to be able to sum over selected integer fields to more easily determine, for example, how much memory objects are using.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • The sum function is available and documented • Provide a function in SQL that can be used to sum over integer fields. • The following query should work: <code>SELECT SUM(appEnaioarchive:totalSize) FROM appEnaioarchive:compound.</code> • SUM() function should conform to SQL-92 standard • tests exist for this function
OKTO-5682	Endpoint for updating the user attribute cache	<p>As a developer, I need an endpoint that I can use to trigger the update of the cache for user attributes so that the information can be updated.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none"> • the endpoint is available and can be used by other services • the endpoint is documented

OKTO-5685	Setting the S3 Object Lock Mode	<p>As a system integrator, I would like to be able to configure the ObjectLock used in S3 to better respond to the service provider's requirements.</p> <p>Currently, when documents are created by Yuuvis Momentum with ObjectLock, it automatically sets the Object Lock mode to COMPLIANCE. This means that we cannot remove the document until the retention effectively expires. However, as a service provider we have various customers and they may also decide that they no longer need our services, so we need a way to delete these files anyway.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• the ServiceLock mode can be configured• the default behavior of the system remains the same without configuration changes
OKTO-5696	Securing the correct version in versioned S3 drivers	<p>As a system integrator, I would also like to ensure that no false files can be foisted on the system when breaking into the system. so as to ensure the immutability of the files.</p> <p>Both Amazon and Dell recommend for applications that use the S3 interface with Object Lock to additionally persist the version ID, because only with this information (object ID and version ID) an unchanged output of the document content can be guaranteed. Since the creation of objects via yMomentum always follows a defined pattern, this is always version 1.</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• The data already stored in the system can still be used unchanged• when accessing the S3 data, version 1 is always explicitly accessed• if data is accessed that is more accessible an explicit error is logged and passed on to the caller
OKTO-5709	Updating the elasticsearch version and the relevant plugins	<p>As a system integrator, would like to use a current version of elasticsearch so that I can benefit from security patches</p> <p>Acceptance criteria:</p> <ul style="list-style-type: none">• the elasticsearch version 7.17 is available as image for installation• the required plugins are installed in the image
