

REST Ingest API

Overview

The standard IngestDocument and security model objects have been simplified to make this API easier to work with in a javascript only environment as opposed to porting over the full Java based object model. All index manipulation methods are based on a sessionId that is returned from the connect method. a few notes about sessions:

- Session are NOT thread safe. API users must ensure that multiple threads do not make concurrent calls against the same sessionId.
- By default sessions expire after a 5 minute timeout of no activity. The timeout can be configured on the SessionConfig input.
- Sessions use system resources. Do not create a new session for every document/method call.

Required Modules

These features require that the **rest** module be included when you run createproject to create the project directories.

REST API Parameters

The master list of QueryRequest REST parameters is on the [HTTP REST APIs](#) page.

[View incoming links.](#) >>

[Attivio Platform](#) , [Attivio Platform 5.6.2 Release Notes](#) , [Content Security Quick Start](#) , [REST API Tutorial](#)

- [Overview](#)
- [Endpoint URIs](#)
- [JSON Document List](#)
- [JSON SessionConfig](#)
- [JSON Principal](#)
- [JSON Real Time Field Update](#)
- [JSON Delete By Query](#)
- [JSON Bulk Update By IDs](#)
- [JSON Bulk Update by Query](#)

Endpoint URIs

URL	METHOD	Input	Output
http://yourAieHost:17000/rest/ingestApi/getSessionCount	GET/POST		Number of session registered with the system
http://yourAieHost:17000/rest/ingestApi/getSessionIds	GET/POST		A JSON array consisting of the registered session IDs
http://yourAieHost:17000/rest/ingestApi/connect	GET/POST	SessionConfig object or null for default session configuration	JSON string session ID to be used with all future methods in place of <code>{sessionId}</code>
http://yourAieHost:17000/rest/ingestApi/disconnect/{sessionId}	GET/POST		Disconnect a client.
http://yourAieHost:17000/rest/ingestApi/isCompleted/{sessionId}	GET/POST		Returns true if all sent in documents and messages have completed processing
http://yourAieHost:17000/rest/ingestApi/commit/{sessionId}	GET/POST		Commit all changes to the index and make all recently fed documents searchable
http://yourAieHost:17000/rest/ingestApi/optimize/{sessionId}	GET/POST		Optimize the index.
http://yourAieHost:17000/rest/ingestApi/refresh/{sessionId}	GET/POST		Make any real time field changes visible in search results
http://yourAieHost:17000/rest/ingestApi/delete/{sessionId}	POST	List of document id's	HTTP 200
http://yourAieHost:17000/rest/ingestApi/feedDocuments/{sessionId}	POST	List of documents	HTTP 200
http://yourAieHost:17000/rest/ingestApi/feedPrincipals/{sessionId}	POST	List of principals	HTTP 200

http://yourAieHost:17000/rest/ingestApi/deletePrincipals/{sessionId}	POST	List of principals	HTTP 200
http://yourAieHost:17000/rest/ingestApi/getDocumentsSent/{sessionId}	GET/POST		The number of documents sent to this session
http://yourAieHost:17000/rest/ingestApi/getDocumentsFed/{sessionId}	GET/POST		The number of documents fed to AIE from this session
http://yourAieHost:17000/rest/ingestApi/sendQueuedDocuments/{sessionId}	GET/POST		Send the documents in the queue
http://yourAieHost:17000/rest/ingestApi/getDocumentsQueued/{sessionId}	GET/POST		The number of documents that have been sent but not yet fed to AIE.
http://yourAieHost:17000/rest/ingestApi/startMessageGroup/{sessionId}	GET/POST		HTTP 200
http://yourAieHost:17000/rest/ingestApi/endMessageGroup/{sessionId}	GET/POST		HTTP 200
http://yourAieHost:17000/rest/ingestApi/deleteByQuery/{sessionId}	POST	Delete by query	HTTP 200
http://yourAieHost:17000/rest/ingestApi/updateRealTimeField/{sessionId}	POST	Real time field update	HTTP 200
http://yourAieHost:17000/rest/ingestApi/bulkUpdateByIds/{sessionId}	POST	Bulk update by Id	HTTP 200
http://yourAieHost:17000/rest/ingestApi/bulkUpdateByQuery/{sessionId}	POST	Bulk update by query	HTTP 200
http://yourAieHost:17000/rest/ingestApi/getLastActivity/{sessionId}	GET/POST		The Date that the session was last used

JSON Document List

```
[ {
  "fields" : {
    "title" : [ "this is the first title" ],
    "date" : [ "1969-12-31T19:00:10.000-0500" ]
  },
  "id" : "1",
  "zone" : "z1",
  "correlationId" : null,
  "mode" : "ADD",
  "permissions" : [ {
    "principal" : {
      "name" : "myGroup1",
      "realm" : "myRealm",
      "groupMemberships" : null,
      "type" : "group"
    },
    "readable" : true
  } ]
}, {
  "fields" : {
    "title" : [ "this is the second title" ],
    "date" : [ "1969-12-31T19:00:20.000-0500" ]
  },
  "id" : "2",
  "zone" : null,
  "correlationId" : "c2",
  "mode" : "ADD",
  "permissions" : [ {
    "principal" : {
      "name" : "Anonymous",
      "realm" : "Anonymous",
      "type" : "user"
    },
    "readable" : true
  } ]
} ]
```

Parameter	Required	Default	Description
id	✓		Document ID
zone		default	Zone for the document to be indexed in.
correlationId			CorrelationId used for matching up document results if the document ID changes
mode		ADD	Mode to use for processing this document. One of: CREATE, ADD, DELETE, PARTIAL. ADD implies a create or replace where as CREATE requires the user to ensure that no document with the same ID already exists.
fields	✓		Name to List of values for the field. Strings, Dates, and number types are supported.
permissions			A list of permissions for the document if content security is enabled
permissions.principal			The user or group who has permission to read the document
permissions.principal.name	✓		The id of the user who has access to the document
permissions.principal.realm	✓		The realm of the user
permissions.principal.type		user	one of user or group
permissions.principal.readable		true	The document will be readable by the user if set to true. By default all documents without explicit read permissions are hidden from the user. Setting readable=false is only necessary to implement deny permissions. For example, a document might be readable by everyone in Engineering except some specific user or sub group.

JSON SessionConfig

```
{
  "sessionTimeout" : -1,
  "documentBatchSize" : 1,
  "orderedCommits" : true,
  "commitInterval" : 1000,
  "commitOnDisconnect" : true,
  "ingestWorkflowName" : "fileIngest",
  "messageResultListenerClass" : "com.acme.myListener",
  "messageResultListenerProperties" : {
    "foo" : "bar"
  }
}
```

Parameter	Required	Default	Description
sessionTimeout		10 minutes	Number of milliseconds between any two posts that will cause the session to be automatically closed. A value of -1 will cause the session to never expire. Any call using a sessionId will cause the session timer to be reset.
documentBatchSize		1	The number of documents to send to AIE in a single batch.
orderedCommits		true	Ordered commits will cause the system to wait for all previously fed content to be indexed before processing the commit.
commitInterval		0	Issue a commit every x number of documents submitted. Set to 0 to disable automatic commits.
commitOnDisconnect		true	<i>(Added in Attivio Platform 5.6.2)</i> Issue a commit when calling disconnect for this session
ingestWorkflowName		ingest	The workflow to use for all ingestion documents and messages
messageResultListenerClass		System default - LoggingResultListener	The class to use for processing message results.
messageResultListenerProperties			Key to Value map of properties for the custom message result listener class to use for configuration.

JSON Principal

```
{
  "name" : "myUsername",
  "id" : "myId",
  "realm" : "myRealm",
  "type" : "user",
  "groupMemberships" : [ {
    "name" : "myGroup1",
    "realm" : "myRealm"
  }, {
    "name" : "myGroup2",
    "realm" : "myRealm"
  } ]
}
{
  "name" : "myOtherUsername",
  "realm" : "myOtherRealm",
  "type" : "user"
}
{
  "name" : "myAlias",
  "realm" : "aliases",
  "type" : "group",
  "principalAliases" : [ {
    "name" : "myUsername",
    "id" : "myId",
    "realm" : "myRealm"
  }, {
    "name" : "myOtherUsername",
    "realm" : "myOtherRealm"
  } ]
}
```

Parameter	Required	Default	Description
name	✓		Name of the user/group
realm	✓		Realm of the user/group
id		value of 'name' parameter	Id of user/group
type		user	One of user or group
groupMemberships			The groups that the user/group is a member of
groupMemberships.name	✓		The group name
groupMemberships.realm	✓		The group realm
groupMemberships.id		value of 'groupMemberships.name' parameter	The group id
principalAliases			The principals to be included in the alias
principalAliases.name	✓		The principal name
principalAliases.realm	✓		The principal realm
principalAliases.id		value of 'principalAliases.name' parameter	The principal id

JSON Real Time Field Update

```
{
  "id" : "doc1",
  "zone" : "zone1",
  "fieldName" : "field1",
  "values" : [ "val1", "val2" ]
}
```

Parameter	Required	Default	Required
id	✓		The ID of the document to update
zone		default	The zone the document is in
fieldName	✓		The name of the real time field to update
values	✓		The new values for the field. Set to null to clear the real time field.

JSON Delete By Query

```
{
  "query" : "title:foo",
  "queryLanguage" : "advanced",
  "queryInitWorkflow" : "queryInit"
}
```

Parameter	Required	Default	Required
query	✓		The query to use for matchign documents to delete.
querylanguage		advanced	The query language to use for parsing the query. One of <code>simple</code> or <code>advanced</code>
queryInitWorkflow		queryInit	The name of the workflow to proces the query with before sending to the ingestion workflow. Genreally speaking this workflow is responsible for query parsing.

JSON Bulk Update By IDs

```
{
  "fields" : {
    "myRealTimeField" : [ "val1", "val2" ]
  },
  "documentIds" : [ "doc1", "doc2" ]
}
```

Parameter	Required	Default	Required
fields	✓		Name to List of values for the field. Strings, Dates, and number types are supported.
documentIds	✓		List of document ids to update with the new field values.



Bulk Updates by ID and Partition Routing

If your index has [partition routing](#) configured to operate on a field other than the document ID, you must specify a value for the routing field when performing bulk updates by document ID. Failure to include a value for the routing field will cause Attivio to route documents to index partitions based on document-ID hash values (the default non-routing behavior).

As an alternative, you can perform a bulk update by query, since the selection-by-query mechanism will ensure that updates are routed to the correct index partitions for the existing documents. You can mimic a bulk update by ID by specifying a query of this form, replacing the various `<ID_*>` placeholders with your document IDs:

```
.id:OR(<ID_1>, <ID_2>, <ID_3>, ... ,<ID_N>)
```

JSON Bulk Update by Query

```

{
  "fields" : {
    "myRealTimeField" : [ "val1", "val2" ]
  },
  "query" : "title:foo",
  "queryLanguage" : "advanced",
  "queryInitWorkflow" : "queryInit"
}

```

Parameter	Required	Default	Required
fields	✓		Name to List of values for the field. Strings, Dates, and number types are supported.
query	✓		The query to use for matchign documents to delete.
querylanguage		advanced	The query language to use for parsing the query. One of <code>simple</code> or <code>advanced</code>
queryInitWork flow		queryInit	The name of the workflow to proces the query with before sending to the ingestion workflow. Genreally speaking this workflow is responsible for query parsing.